



ARTICLE

Food Price Prediction in Nigeria: A Comparative Analysis of Linear Regression and Machine Learning Models to Analyze Subsector Price Interdependencies

Mohammad Shahfaraz Khan¹ , Amir Ahmad Dar^{2*} , Chinyere Perpetua Okechukwu² , Mohammed Wamique Hisam³ , Imran Azad¹ , Aseel Smerat⁴ , Murtaza M. Junaid Farooque³

¹ College of Economics and Business Administration, University of Technology and Applied Sciences-Salalah, Salalah 211, Oman

² Department of Statistics, Lovely Professional University, Phagwara 144411, India

³ College of Commerce and Business Administration, Dhofar University, Salalah 211, Oman

⁴ Faculty of Educational Sciences, Al-Ahliyya Amman University, Amman 19328, Jordan

ABSTRACT

This study investigates the effectiveness of predictive performance of multiple modeling techniques in forecasting Nigeria's Food Price Index (FPI) using monthly data from the Food and Agriculture Organization (FAO) Food Price Index spanning January 1990 to January 2025. Five major food subsector indices—Meat, Dairy, Cereals, Oils, and Sugar—serve as predictors. Initial regression on the raw series produced inflated R^2 values, prompting diagnostic tests that revealed severe non-stationarity and multicollinearity. These issues were addressed through a two-step transformation: natural logarithmic conversion to stabilize variance, followed by first-order differencing to achieve stationarity and eliminate false trends. Post-transformation diagnostics confirmed full correction, with all variables stationary and multicollinearity reduced to acceptable levels. Using standardized log-differenced data, five models were evaluated: Multiple Linear Regression, Artificial Neural Network, Random Forest, Support Vector Machine, and Long Short-Term Memory (LSTM). On the transformed scale, the Support Vector Machine was the

*CORRESPONDING AUTHOR:

Amir Ahmad Dar, Department of Statistics, Lovely Professional University, Phagwara 144411, India; Email: sagaramir200@gmail.com

ARTICLE INFO

Received: 10 October 2025 | Revised: 5 December 2025 | Accepted: 21 December 2025 | Published Online: 12 June 2026

DOI: <https://doi.org/10.36956/rwae.v7i2.2820>

CITATION

Khan, M.S., Dar, A.A., Okechukwu, C.P., et al., 2026. Food Price Prediction in Nigeria: A Comparative Analysis of Linear Regression and Machine Learning Models to Analyze Subsector Price Interdependencies. *Research on World Agricultural Economy*. 7(2): 745–772.

DOI: <https://doi.org/10.36956/rwae.v7i2.2820>

COPYRIGHT

Copyright © 2026 by the author(s). Published by Nan Yang Academy of Sciences Pte. Ltd. This is an open access article under the Creative Commons Attribution-NonCommercial 4.0 International (CC BY-NC 4.0) License (<https://creativecommons.org/licenses/by-nc/4.0/>).

strongest performer; however, after back-transforming predictions to the original scale through recursive exponential reconstruction, Random Forest achieved the highest accuracy ($R^2 = 0.9342$, MAPE = 0.65%). LSTM models performed poorly, indicating a structural mismatch with differenced series lacking persistent trends. The study provides three methodological insights: (1) top performance on transformed data does not guarantee accuracy after inverse conversion; (2) log-differencing is highly effective for resolving non-stationarity, heteroscedasticity, and multicollinearity while preserving economic interpretability; and (3) ensemble averaging enhances recursive multi-step forecasting stability. For practical policy applications requiring multi-month food price projections, Random Forest is recommended, with Multiple Linear Regression serving as a transparent alternative for coefficient interpretation. These findings support the design of more informed food security interventions and market stabilization strategies in Nigeria.

Keywords: Food Price Forecasting; Log-Differencing; Multiple Linear Regression; Random Forest; Support Vector Machine; LSTM

1. Introduction

Accurate forecasting is crucial for efficient policy intervention, agricultural planning, and household welfare protection because food price volatility in Nigeria poses significant risks to food security and macroeconomic stability^[1]. The nation saw sharp increases in food prices between 2007–2012 and 2020–2022, which disproportionately impacted low-income households, which spend more than 60% of their income on food^[2]. Even though predicting price movements is crucial, Nigeria's current forecasting efforts are still limited and frequently rely on univariate models or concentrate on individual commodities while ignoring the interconnected structure of food subsectors. The ability of policymakers to react proactively to market shocks is undermined by this gap. Nigeria's recent food price crisis highlights the importance of this research for policy. Between January 2023 and December 2024, an additional 14.2 million Nigerians fell into poverty due to food inflation, which reached historic levels of 40.87% in June 2024^[3]. The removal of fuel subsidies in May 2023, the devaluation of the naira, ongoing insecurity that hindered agricultural production in northern states, and external shocks like the conflict between Russia and Ukraine and supply chain disruptions all contributed to this crisis^[2]. Hence, precise and timely food price forecasts become more than just an academic exercise; they are a crucial tool for social protection targeting, humanitarian response planning, and policy intervention.

In order to overcome these limitations, this study examines how well modern machine learning algorithms and conventional econometric techniques forecast Nigeria's Food Price Index (FPI), with a focus on capturing ripple effects among important subsectors, including meat, dairy, cereals, oils, and sugar. By identifying the most accurate and reliable forecasting methods, this comparative method seeks to provide policymakers, humanitarian organizations, and market regulators with useful information^[4]. This study fills three significant gaps in the current literature. First, there is little data comparing the effectiveness of machine learning algorithms and conventional econometric models in Nigerian food price forecasting. Second, there is still a lack of research and application in predictive modeling regarding the cross-subsector price interdependencies that collectively impact the dynamics of aggregate food prices. Third, previous research has not sufficiently addressed methodological issues like non-stationarity, multicollinearity, and non-linear relationships. Thus, the following questions are examined in this study:

- How well can machine learning models and Multiple Linear Regression (MLR) predict Nigeria's Food Price Index when subsector interdependencies are taken into account?
- After taking multicollinearity and data stationarity into account, which modeling technique yields the most accurate forecasts?

Agriculture remains central to Nigeria's economy,

providing livelihoods for more than 70% of the population^[3,5]. This makes accurate food price forecasting especially critical for strengthening food security, improving supply chain efficiency, and guiding agricultural investment strategies. The Food Consumer Price Index (FCPI), which tracks changes in consumer food prices, is closely tied to domestic agricultural productivity and therefore serves as a key indicator for understanding short-term and structural dynamics in Nigeria's food markets.

By offering a methodical comparison of statistical and machine learning models to assess their applicability for Nigeria's intricate and dynamic food markets—particularly in capturing non-linear behavior and cross-subsector interactions—this study contributes to the body of literature^[6]. The study also takes into account global explanatory factors, recognizing their significant impact on domestic food prices^[7]. To ensure a thorough and multifaceted evaluation of predictive accuracy, model performance is assessed using Root Mean Square Error (RMSE), Mean Absolute Error (MAE), Symmetric Mean Absolute Percentage Error (SMAPE), and the Coefficient of Determination (R^2)^[8-10]. Evidence demonstrating the superior performance of machine learning models, especially Random Forest and related ensemble methods, in agricultural yield and price forecasting and the pressing need for dependable, data-driven tools to address Nigeria's ongoing food security challenges are the driving forces behind this work^[9,11]. The study's main objective is to produce reliable, empirically supported forecasting models that can aid in the development of policies, market stabilization strategies, and plans for food assistance in Nigeria.

1.1. Multiple Linear Regression (MLR)

Linear regression analysis is a statistical method used to estimate the relationship between one or more independent variables and a single dependent variable. A Multiple Linear Regression method was applied to determine significant factors from potential explanatory variables. The general form of a multiple linear regression model is given by:

$$Y = \beta_0 + \beta_1 X_1 + \beta_2 X_2 + \dots + \beta_k X_k + \varepsilon_i \quad (1)$$

where Y = Food Price Index; β_0 is the intercept

and $\beta_1, \beta_2, \dots, \beta_k$ are coefficients of the variables $X_1, X_2, \dots, X_k$ which are independent variables, and ε_i is the error term. In this study, we have five predictor variables which include meat, dairy, cereals, oils and sugar.

Estimation of Model Parameters

The Maximum Likelihood method was used to estimate the parameters of the regression model. Analysis of variance was used to test the general significance of the model. The model is significant when at least one predictor significantly influences the dependent variable.

Model Adequacy Checking

The model adequacy diagnostics were checked via residual analysis or residual plots. These include:

- **Linearity:** Linearity indicates that the relationship between the dependent and independent variables should be linear in the parameters.
- **Normality:** Normality of the random error was tested with a plot of residuals against the cumulative probability or a quantile-quantile plot.
- **Homoscedasticity:** Plotting the standardized residuals against time order is used to examine whether the variance of the error term is constant.
- **Multicollinearity:** The decision on multicollinearity was based on the Variance Inflation Factor (VIF). If the value of VIF is less than 10, the collinearity is tolerable, but if it is more than 10, VIF indicates a risk.

Sun et al.^[12] assert that Multiple Linear Regression (MLR) is a statistical technique used to model the relationship between a dependent variable and two or more independent variables by fitting a linear equation to the observed data. In agricultural research, MLR is employed to understand and predict crop yields, optimize resource allocation, and assess the impact of various factors on agricultural productivity. Danbatta et al.^[13] employed MLR to identify factors affecting ginger production, revealing that fertilizer use, disease incidence, education levels, farm size, farmer experience, and weed presence are significant predictors. Accordingly, he con-

cluded that the model explained 85% of the variance in ginger production.

Again, Eticha^[14] investigated factors affecting coffee production in Southwest Ethiopia using MLR and asserted that coffee is the major source of income for smallholder farmers and is a leading export crop for Ethiopia; hence, identifying factors affecting coffee production is paramount to improving coffee productivity in the country. Their results revealed that factors such as education level, types of coffee, farmers' total income, and farm size are determinants of coffee production. Therefore, they concluded that farmers should develop their education level to maximize their coffee production and that producing farmers should use different types of coffee seeds to maximize productivity. They also opined that farmers' income and farm size are important factors that increase coffee production.

Wu et al.^[15] employed a linear regression algorithm to construct a model and indicated that it could play an important role in predicting soil moisture content (SMC) of cultivated soils over a large area for agricultural irrigation and remote monitoring of crop yields. Their results indicated that the inversion model established by the MLR algorithm (the predicted R^2 was 0.83 and the RMSE was 0.0078) was feasible and highly accurate, hence linear regression plays an important role in predicting SMC of cultivated soils over a large area for agricultural irrigation and remote monitoring of crop yields.

Istifaroh and Harini^[16] affirmed that agriculture is an important sector supporting national food availability. Soil and climatic conditions in Indonesia are suitable for agriculture, and a study was carried out to determine which commodities are suitable for specific land characteristics and how land characteristics influence agricultural productivity. Their results showed that eight variables of land characteristics significantly affect agricultural productivity by 93%. The independent variables were tested using the Simultaneous F-Test, which showed a combined effect on agricultural productivity.

Tule et al.^[17] investigated the predictability of agricultural commodity prices in Nigeria's inflation forecast, considering twelve major agricultural commodities. Their predictability was evaluated singly and jointly for

both food and headline inflation using various methodological approaches. They obtained consistent results for both in-sample and out-of-sample forecasts.

1.2. Artificial Neural Network (ANN)

ANNs are especially versatile, serving as universal approximators and proving effective in forecasting a variety of time-series data, both linear and non-linear^[18]. Nonetheless, the particular features of the dataset, such as its linearity, stationarity, and the existence of intricate interdependencies, continue to influence the choice of the best model. Because artificial neural networks can model the intricate, non-linear relationships found in economic time series data, they are being used more and more in agricultural price forecasting^[19]. For agricultural price prediction, these models have proven to perform better than traditional statistical techniques like linear regression and Box-Jenkins approaches, especially when handling complex data dependencies and high seasonality^[20,21]. For instance, by capturing both trend and seasonality components, hybrid models that combine traditional methods with neural networks have demonstrated improved accuracy in predicting monthly rainfall and agricultural commodity prices^[11,13]. Additionally, the application of deep learning architectures, such as Long Short-Term Memory networks, has demonstrated significant improvements in agricultural commodity price prediction, frequently outperforming conventional time series models (e.g., ARIMA) due to their ability to manage volatile data and long-term dependencies^[10]. Advanced deep learning models, such as NBEATX and TransformerX, have outperformed recurrent neural networks in agricultural price forecasting by successfully utilizing exogenous variables to increase prediction accuracy and generalize across various datasets^[10].

Artificial Neural Networks (ANNs) are computational models inspired by the structure and function of biological neural networks and are composed of interconnected nodes or artificial neurons, organized in layers that process and transmit signals to each other. The main objective of ANN is to develop a system to perform various computational tasks faster than traditional systems. The idea of ANNs is based on the belief that the

way the human brain makes the right connections can be imitated using silicon and wires as living neurons and dendrites. It is known that the human brain is composed of 86 billion nerve cells called neurons. They are connected to other thousand cells by Axons. Stimuli from the external environment or inputs from sensory organs are accepted by dendrites. These inputs create electric impulses, which quickly travel through the neural network. A neuron can then send the message to other neurons to handle the issue, or it cannot send it forward.

1.3. Random Forest

Random Forest (RF) Regression has evolved as a powerful and versatile machine learning tool, significantly applied in recent years across various domains. Random Forest (RF) has demonstrated exceptional efficacy as a supervised learning method for both regression and classification applications. In the realm of water resource management, Tyralis et al.^[22] explored the implementation of RF algorithms, noting that while Breiman's basic formulation is routinely employed, its advanced variations remain underutilized despite their potential to solve more complicated real-world situations. In the banking industry, Malysheva et al.^[23] employed RF to anticipate average ATM cash withdrawals—a key task for improving cash delivery logistics and maintaining uninterrupted ATM services. According to their investigation, RF could anticipate average daily withdrawals with almost 79.8% accuracy, proving its usefulness. Zhang et al.^[24] employed the RF algorithm to estimate forest aboveground biomass (AGB) using various satellite-derived data, showing that ensemble models like RF considerably enhance estimation accuracy across low and high AGB levels. When Doz et al.^[25] used RF regression in education to predict student performance based on demographic characteristics, they found that socioeconomic position, school type, and geography are important predictors, whereas gender and ethnicity have little bearing. For the purpose of estimating apartment prices in Nicosia, Dimopoulos et al.^[26] compared RF with conventional linear regression and found that RF was more accurate, especially when dealing with a variety of independent variables.

1.4. Support Vector Machine (SVM)

The idea of Support Vector Machines (SVMs) was initially developed by Vapnik^[27]. At first, it was created to sort objects into two groups, but later it was also adapted to predict numbers, a method called Support Vector Regression (SVR). Simply put, the algorithm takes the training data and maps it into a particular area where it looks for a straight line (or a boundary) that best separates the two groups, like drawing a line between two types of items. If it's hard to separate the groups in the original space, SVM will move the data into a higher-dimensional space where finding a clear dividing line becomes easier. This ability to go into more complex environments to find a solution is what makes SVMs. This ability to move into more complex spaces to find a solution is what makes SVMs different from many other machine learning methods. In their paper "Quadratic hyper-surface kernel-free large margin distribution machine-based regression and its least-square form," He et al.^[28] contend that ϵ -support vector regression (ϵ -SVR) reduces the margin between expected and actual values, which is essential for generalization performance. They emphasize that margin distribution is more crucial since merely reducing structural risk does not ensure a well-margined distribution. The authors suggest two reliable methods that use an optimal margin distribution: quadratic hyper-surface kernel-free large margin distribution machine-based regression (QLDMR) and its least-squares equivalent (QLSLDMR). Experimental results indicate that these models, integrating the kernel-free technique with optimal margin distribution, exhibit superior robustness, generalization, and efficiency. Support Vector Machines (SVMs) are a top machine learning algorithm for predicting active compounds and molecular properties, especially in chemoinformatics and drug discovery, according to Rodríguez-Pérez and Bajorath^[29]. Unlike many machine learning techniques that employ lower-dimensional representations, SVMs are renowned for their ability to operate in high-dimensional feature spaces. Their use in tasks like compound classification, ranking, and regression modeling is made possible by this feature. SVMs are still prized in the field for their resilience and adaptability despite the emergence

of deep learning, and the study delves deeper into their advantages and disadvantages in contemporary computational chemistry.

1.5. Long Short-Term Memory (LSTM)

Long Short-Term Memory networks are a type of RNN that is made to better handle and model sequential data. They are mostly known for being able to capture long-range dependencies, which is something that traditional RNNs often can't do because of problems like vanishing gradients. Gers et al.^[30] asserted that LSTMs could address numerous time series tasks that remained unsolvable by feed-forward networks employing fixed-size time windows. In their work, they find that LSTM's superiority does not carry over to certain simpler time series prediction tasks solvable by time window approaches, hence suggesting the use of STM only when simpler traditional approaches fail.

Hochreiter and Schmidhuber^[31] addressed a serious limitation of traditional RNNs. Learning to store information over long time intervals is problematic in standard RNNs, where error signals fade or decay over time during backpropagation. Hochreiter and Schmidhuber had earlier identified this problem^[32]; this prevented RNNs from capturing long-term dependency. In their paper, they presented an adaptive, gradient-based approach that was novel and efficient, called Long Short-Term Memory (LSTM). By introducing the mechanism of truncating the gradients only when it is innocuous, LSTM maintains a constant error flow through what they call constant error carousels, learning to bridge time lags in excess of 1000 discrete time steps. They emphasized the fact that LSTM operates locally in space and time and has a per-time-step and per-weight computational complexity in $[0, 1]$. Using artificial data sets with local, distributed, real-valued, and noisy pattern representations, their experiments showed that not only does LSTM learn much faster but also yields far more successful runs. They also provided evidence that LSTM can solve complex tasks containing long-time-lag dependencies that had hitherto been beyond the reach of the recurrent network algorithms. In Van et al.'s^[33] wide review of LSTM's formulation and training, relevant applications claimed that LSTM has revolutionized both

machine learning and neurocomputing areas, and according to several online sources, this model improved Google's speech recognition, considerably improved machine translations on Google Translate, and the answers of Amazon's Alexa. Besides, they confirmed that this neural system is also used by Facebook, where more than 4 billion LSTM-based translations are processed every day as of 2017. Curiously, it was said that until the emergence of the LSTM, RNNs had presented a somewhat discrete performance; once again, the success of the latter kind of recurrent network lies in a positive feature: it could handle the problem of exploding/vanishing gradients. For example, Zhang et al.^[24] asserted that RNNs learn features and long-term dependencies from sequential and time-series data. They maintained that the RNNs comprise a stack of nonlinear units where at least one connection between units forms a directed cycle. Consequently, they reiterated that well-trained RNNs can model any dynamical system; however, training RNNs is mostly plagued by issues in learning long-term dependencies. They did a survey on RNNs and several new advances underscoring the basics and recent advances in the application areas.

2. Materials and Methods

The data utilized for this study were sourced from the Food and Agriculture Organization (FAO) via the FAO Food Price Index. The dataset comprises historical monthly records of the global Food Price Index and its major subsector indices—including, but not limited to, Meat, Dairy, Cereals, Oils, and Sugar—covering the period from 1990 to March 2025, spanning over thirty-five years of continuous price monitoring.

Nigeria's deep integration into international commodity markets, especially for cereals (dominated by imports of wheat and rice), vegetable oils (largely imported), and sugar, justifies the choice of Nigerian food price dynamics, even though it depends on global FAO indices. Although domestic meat and dairy production offers some protection against global shocks, it is still linked to the cost of feed grains (cereals) and imported intermediate inputs. As a result, this dataset structure makes it possible to evaluate how changes in inter-

national subsector prices affect the overall inflation of food prices in Nigeria's import-dependent food system throughout the duration of the study. This volatility—consistent with the documented food price crises of 2007–2008, 2011, and 2020–2022—underscores the need for reliable forecasting models.

2.1. Theoretical Frame Work

This study's forecasting approach is based on three interrelated theoretical frameworks: price transmission theory, machine learning generalization theory, and time series econometrics.

2.1.1. Price Transmission Theory

Price changes in subsector commodities in integrated food markets are transmitted to aggregate food price indices via a number of channels^[34]. Price shocks from international commodity markets are transmitted vertically through domestic supply chains to retail prices, with market structure, transaction costs, and policy interventions influencing the speed of transmission. Horizontal transmission happens between related food categories through complementarity (such as bread and butter) or substitution (such as consumers switching between meat and dairy proteins). Cross-commodity transmission results from shared production inputs; for instance, rising cereal prices drive up the cost of livestock feed, which in turn drives up the price of meat and dairy products. Modeling the overall food price is justified by these mechanisms.

2.1.2. Time Series Econometrics

Stochastic trends frequently appear in economic time series, making estimation more difficult. Regressing one non-stationary series on another can result in spurious correlations and artificially high R^2 values, as demonstrated by Newbold and Harvey^[35]. Therefore, achieving stationarity—that is, a constant mean, variance, and autocovariance—is necessary for reliable inference. When unit roots are present, the Box-Jenkins method^[36] suggests differencing. In accordance with economic theory, which emphasizes proportional rather than absolute price movements, applying a log transformation prior to differencing stabilizes variance (addressing heteroscedasticity common in price series where

volatility scales with price level) and produces first differences that approximate percentage changes.

2.1.3. Machine Learning

Controlling the bias-variance tradeoff is critical to model performance. While non-parametric models (Random Forests, Neural Networks, and SVMs) reduce bias but may increase variance, linear models reduce variance but may experience bias under non-linear conditions. By averaging, ensemble techniques such as Random Forests reduce variance^[37]. SVMs use high-dimensional mappings to capture non-linear patterns^[38]. Long-term dependencies are learned by LSTM networks^[31], but their benefit is diminished when underlying trends are eliminated through differencing.

The rest of our methodological steps follow through the **Figure 1**.

2.2. Data Preprocessing, Transformation and Partitioning

Stationarity Testing

We used a two-stage transformation to handle heteroscedasticity and non-stationarity while guaranteeing meaningful economic interpretation. To stabilize variance, all variables were first converted to natural logarithms. Second, the log-transformed series was subjected to first-order differencing, which produced log returns that roughly corresponded to percentage changes. This method, which is common in financial econometrics^[39], guarantees stationarity and yields coefficients that can be understood as elasticities. Stationarity was assessed using: the Augmented Dickey-Fuller (ADF) test and the Kwiatkowski–Phillips–Schmidt–Shin (KPSS) test. The ADF regression takes the form:

$$\Delta Y_t = \alpha + \beta t + \gamma Y_{t-1} + \sum_{i=1}^k \delta_i \Delta Y_{t-i} + \varepsilon_t \quad (2)$$

The hypotheses tested are:

$$H_{0,\gamma} = 0 \quad (\text{unit root; non-stationary}) \quad (3)$$

$$H_{1,\gamma} < 0 \quad (\text{stationary}) \quad (4)$$

We selected Lag length k using the Akaike Information Criterion (AIC). Rejection of H_0 at the 5% level indicates stationarity.




Figure 1. Methodological Flowchart: Food Price Prediction in Negeria: A Comprehensive Analysis Pipeline.

KPSS Test

Our KPSS statistic was computed thus:

$$KPSS = \frac{1}{T^2} \sum_{t=1}^T \frac{S_t^2}{\sigma^2} \quad (5)$$

where S_t is the cumulative sum of residuals from regressing Y_t on a constant, and σ^2 is a consistent estimate of the long-run variance. The dataset is found to be stationary when the ADF rejects a unit root and the KPSS fails to reject stationarity.

Log Transformation

All variables were transformed using natural logarithms to stabilize variance and convert multiplicative relationships into additive form using:

$$Z_t = \ln(X_t) \quad (6)$$

First-Order Differencing

First differencing of the log-transformed variables was computed in order to remove trends using:

$$\Delta Z_t = Z_t - Z_{t-1} = \ln \left(\frac{X_t}{X_{t-1}} \right) \quad (7)$$

The transformation provides a stationary series of 422 usable observations.

Multicollinearity Assessment

The decision on multicollinearity is based on the variance inflation factor (VIF). If the value of VIF is less than ten, the collinearity is tolerable, but if it is more than 10, VIF is a risk. Variance Inflation Factors (VIF) were used to diagnose multicollinearity:

$$VIF_j = \frac{1}{1 - R_j^2} \quad (8)$$

where R_j^2 is from regressing predictor X_j on all other predictors. VIF values above 10 indicate severe multicollinearity.

Standardization

We standardized Log-differenced variables as

$$X_t^* = \frac{X_t - \mu}{\sigma} \quad (9)$$

where μ and σ are the mean and standard deviation computed from the training set only to avoid information leakage.

Data Partitioning

The dataset was divided into training and testing subsets using a chronological split for time series forecasting. The training set covered 408 monthly observations from January 1990 to December 2023 and included a variety of economic cycles, such as the food crisis of 2007–2008 and the pandemic of 2020–2022. The testing set included 14 monthly observations from January 2024 to March 2025, enabling an out-of-sample assessment for short-term forecasting. In order to avoid temporal leakage, cross-validation was avoided, and temporal order-respecting validation techniques were advised. The train-test split approach was chosen for a clean assessment of model performance.

2.3. Model Specification

Multiple Linear Regression (MLR)

Multiple Linear Regression was estimated using Ordinary Least Squares on the standardized log-differenced training data. The model specification is given by:

$$Y_t = \beta_0 + \beta_1 X_{1,t} + \beta_2 X_{2,t} + \beta_3 X_{3,t} + \beta_4 X_{4,t} + \beta_5 X_{5,t} + \varepsilon_t, \quad (10)$$

where Y_t denotes the log-differenced Food Price Index, $X_{j,t}$ represent the five standardized subsector indices (Meat, Dairy, Cereals, Oils, and Sugar), β_j denote the corresponding coefficients interpreted as semi-elasticities (the percentage point change in Y_t associated with a one percentage point change in $X_{j,t}$), and ε_t is an error term assumed to be independently and identically distributed with zero mean and constant variance.

Parameter estimation was performed through maximum likelihood, assuming normally distributed errors, equivalent to Ordinary Least Squares. Model diagnostics involved: (1) Durbin–Watson test for autocorrelation in residuals; (2) Shapiro–Wilk test for residual normality; (3) Breusch–Pagan test for heteroscedasticity; (4) ANOVA for overall model significance; and (5) All calculations were executed using the `statsmodels` library (version 0.13) in Python 3.9.

Machine Learning Models

In this work, we evaluated four machine learning models, each with the following specifications: The ReLU activation, L2 regularization, and dropout are used to control overfitting in the ANN's feedforward design, which has three hidden layers with 64, 32, and 16 neurons, respectively. The Adam optimizer and Mean Squared Error (MSE) loss function are used in its training. Using early stopping based on validation loss for optimal training, the model is constructed in TensorFlow 2.10 using Keras.

The Random Forest regression model uses settings that avoid overfitting, like limiting maximum depth and enforcing minimum samples at leaves, to aggregate predictions from 100 decision trees. Using `scikit-learn` 1.0.2, it evaluates features using Gini significance and bootstrap sampling. These models demonstrate a delicate balancing act between generalization and complexity.

Support Vector Regression (SVR) employed a Radial Basis Function (RBF) kernel to model non-linear relationships. The RBF kernel is defined as:

$$K(x_i, x_j) = \exp(-\gamma \|x_i - x_j\|^2), \quad (11)$$

where the hyperparameters were set to $C = 1.0$, $\epsilon = 0.1$, and $\gamma = \text{scale}$. The ϵ -insensitive loss function encourages sparsity in support vectors. The model was implemented using the SVR class in `scikit-learn`.

The Long Short-Term Memory (LSTM) network consisted of two LSTM layers with 50 units each and dropout rates of 0.2, followed by a dense layer with 25 units and a single-output layer. A lookback window of 12 time steps was used to capture annual seasonality. The model was trained using the Adam optimizer with a learning rate of 0.001 for up to 100 epochs, with early stopping applied based on validation loss. The implementation, done in TensorFlow 2.10, resulted in a reduced effective training sample from 408 to 396 sequences due to the sequence construction process.

2.4. Model Evaluation Metrics

Model performance was assessed using four complementary metrics, each providing a distinct perspective on predictive accuracy.

Root Mean Squared Error (RMSE)

$$\text{RMSE} = \sqrt{\frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_i)^2} \quad (12)$$

RMSE measures average prediction error in the original units of the dependent variable; lower RMSE values indicate better performance.

Mean Absolute Error (MAE)

$$\text{MAE} = \frac{1}{n} \sum_{i=1}^n |y_i - \hat{y}_i| \quad (13)$$

MAE represents the mean absolute deviation between predicted and actual values. Unlike RMSE, it provides a linear penalty for errors and is more robust against outliers.

Mean Squared Error (MSE)

$$\text{MSE} = \frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_i)^2 \quad (14)$$

MSE serves as the squared version of RMSE and is commonly used in optimization during training.

Coefficient of Determination (R^2)

$$R^2 = 1 - \frac{SS_{\text{res}}}{SS_{\text{tot}}} = 1 - \frac{\sum_{i=1}^n (y_i - \hat{y}_i)^2}{\sum_{i=1}^n (y_i - \bar{y})^2} \quad (15)$$

The coefficient of determination, R^2 , captures the proportion of variance in the dependent variable explained by the model. It ranges from $-\infty$ to 1, with values above 0.70 generally considered good and above 0.90 excellent.

To assess how well the model predicts new data, the metrics were calculated solely using the test set. Predictions were then converted back from the transformed scale to the original Food Price Index scale so that the accuracy accurately represents forecast performance in the real world.

2.5. Inverse Transformation and Original-Scale Evaluation

To interpret and evaluate forecasting performance on the original Food Price Index (FPI) scale, log-

differenced predictions were transformed back using the formula:

$$\text{FPI}_t = \text{FPI}_{t-1} \times \exp(\Delta \log(\text{FPI}_t)), \quad (16)$$

where FPI_{t-1} is the previous period's actual Food Price Index value and $\Delta \log(\text{FPI}_t)$ denotes the model's predicted log difference. This recursive process compounds predictions over time, with the initial condition set to the last observed training value (December 2023).

Evaluation on the original FPI scale used the same metrics (RMSE, MAE, MSE, R^2) and additionally incorporated Mean Absolute Percentage Error (MAPE), defined as:

$$\text{MAPE} = \frac{100}{n} \sum_{i=1}^n \left| \frac{y_i - \hat{y}_i}{y_i} \right|. \quad (17)$$

MAPE expresses average forecast error as a percentage of actual values, enabling intuitive interpretation and alignment with standards in forecasting literature. Values below 10% are typically considered accurate, while values below 5% indicate high accuracy.

MAPE calculates the average forecasting error as a percentage of actual values; values less than 10% are considered accurate, and values less than 5% are considered very accurate. A comprehensive performance evaluation is made possible by the use of both log-differenced and original data, emphasizing that the best outcomes on converted data could not accurately represent the efficacy in practical applications relevant to market and policy reforms

2.6. Statistical Software and Reproducibility

All analyses were performed using Python 3.9 with major libraries, including pandas for data manipulation, numpy for numerical computing, matplotlib and seaborn for visualization, scipy for statistical tests, statsmodels for econometric modeling, scikit-learn for machine learning, and TensorFlow/Keras for deep learning. Random seeds were set for all stochastic operations (`np.random.seed(42)`, `tf.random.set_seed(42)`, `random.seed(42)`) to ensure reproducibility. Complete codes are available upon request.

3. Results and Analysis

The dataset comprises historical monthly records of the global Food Price Index and its major subsector indices—covering the period from 1990 to March 2025, spanning over thirty-five years—and including, but not limited to, Meat, Dairy, Cereals, Oils, and Sugar. **Table 1** shows that the dataset comprises 423 observations across six variables: Food Price Index, Meat, Dairy, Cereals, Oils, and Sugar. The mean Food Price Index is 88.30 with a standard deviation of 26.92, indicating

moderate variation in food prices over the study period. Among the subsectors, Oils recorded the highest average price (Mean = 92.24), while Sugar had the lowest (Mean = 84.58). Oils also showed the greatest variability (Std. Dev = 39.05), followed by Sugar and Cereals, suggesting significant fluctuations in their prices. The minimum and maximum values highlight the range of price changes: the Food Price Index ranged from 50.80 to 160.20, while Oils spanned from 35.83 to 251.80—the widest range among all categories.

Table 1. Descriptive Analysis of the Subsectors of the Food Price Index.

Statistic	FPI	Meat	Dairy	Cereals	Oils	Sugar
Count	423	423	423	423	423	423
Mean	88.2986	86.6452	87.4500	89.2633	92.2365	84.5823
Std. Deviation	26.9230	17.8556	34.2761	31.8386	39.0546	32.9749
Minimum	50.8000	52.3000	36.8000	48.6000	35.8300	31.8000
Maximum	160.2000	125.4000	158.2000	173.5000	251.8000	183.2000

The Food Price Index (blue line) and its major subsector components—meat (orange line), dairy (green line), cereals (red line), oils (purple line), and sugar (brown line)—from 1990 to 2024 are depicted in **Figure 2**. Food prices are comparatively stable at lower values, particularly between 1990 and 2004, but they also show a notable price increase between 2007 and 2012 and between 2020 and 2022. This trend coincides with the global economic shocks or food crises that occurred

during the COVID-19 pandemic. Meat, dairy, cereals, oils, and sugar are among the subsectors that move together, but they also exhibit independent patterns. For example, oil shows a notable peak in 2021–2022, indicating a sharp increase in vegetable oil prices; similar patterns are seen in the series of cereals and sugar. Once more, cereals, sugar, and oils exhibit trends that closely resemble the Food Price Index, suggesting that they have a significant impact on the FPI in Nigeria.




Figure 2. Food Price and Independent Variables Over Time

All variables have strong positive correlations, with correlation coefficients ranging from 0.69 to 0.94, according to the raw data correlation matrix of **Figure 3**. Strong correlations between cereals and oils ($r = 0.92$) and dairy ($r = 0.91$) indicate significant interdependencies caused by similar production inputs, shared distribution networks, and coordinated reactions to macroeconomic shocks. Although these high correlations made sense economically, they also raised questions about

multicollinearity, which were resolved by data transformation, as explained in the next section.

3.1. Stationarity Tests and Data Transformation

All variables were non-stationary in their initial form, according to initial stationarity diagnostics using the Augmented Dickey-Fuller (ADF) and Kwiatkowski-Phillips-Schmidt-Shin (KPSS) tests as seen in **Table 2**.




Figure 3. Correlation Matrix of Raw Data.

Table 2. Stationarity Testing on Raw Data (ADF and KPSS).

Variable	ADF Statistic	ADF (<i>p</i> -Value)	KPSS Statistic	KPSS (<i>p</i> -Value)	Status
Food Price Index	−1.7355	0.4129	2.2916	0.0100	Non-stationary
Meat	−0.6533	0.8585	2.1695	0.0100	Non-stationary
Dairy	−1.2365	0.6577	2.4754	0.0100	Non-stationary
Cereals	−2.0801	0.2526	2.0932	0.0100	Non-stationary
Oils	−2.1810	0.2132	1.9435	0.0100	Non-stationary
Sugar	−2.3362	0.1606	1.5450	0.0100	Non-stationary

While the KPSS test rejected the null hypothesis of stationarity for all variables (all p -values < 0.05), the ADF test was unable to reject the null hypothesis of a unit root for all six series (all p -values > 0.05). These findings support the existence of stochastic trends in the data, which, if ignored, would result in spurious regres-

sion (high R^2 values reflect common trends rather than true causal relationships)^[40].

We developed a two-stage transformation protocol to handle non-stationarity while guaranteeing economically significant interpretation. To stabilize variance and reduce the impact of extreme values, we ap-

plied the natural logarithms first to every variable. In economic time series analysis, this transformation is a standard practice, as it converts absolute changes to proportional changes. However, stationarity tests on the

log-transformed data, as shown in **Table 3**, still revealed that non-stationarity persisted (ADF p -values > 0.05 for most variables), irrespective of the transformation improving some statistical characteristics.

Table 3. Stationarity Test Results on Log Transformed Data.

Variable	ADF Stat	ADF (p -Value)	KPSS Stat	KPSS (p -Value)	Status
Food Price Index	-1.3005	0.6289	2.3437	0.0100	Non-stationary
Meat	-0.8175	0.8139	2.0462	0.0100	Non-stationary
Dairy	-1.4348	0.5655	2.6497	0.0100	Non-stationary
Cereals	-1.7899	0.3855	2.2549	0.0100	Non-stationary
Oils	-2.0880	0.2494	2.1111	0.0100	Non-stationary
Sugar	-2.1532	0.2236	1.5979	0.0100	Non-stationary

This is evidently seen in **Figure 4**.

Consequently, the log-transformed series was then subjected to first-order differencing, generating log returns that approximate percentage changes between successive periods:

$$\Delta \log(X_t) = \log(X_t) - \log(X_{t-1}) = \log\left(\frac{X_t}{X_{t-1}}\right). \quad (18)$$

In addition to addressing non-stationarity, this transformation yields economically interpretable coefficients that represent elasticities or semi-elasticities.

The log-differenced series, as shown in **Figure 5**, now exhibits stationary behavior and fluctuates around zero. Trends were effectively eliminated by the transformation, while co-movements and volatility patterns were retained.

To further confirm the stationarity of the log-differenced dataset, a post-transformation stationarity test was carried out, and the results in **Table 4** provided strong evidence that all series achieved stationarity. The ADF test rejected the null hypothesis of a unit root for all variables with extremely low p -values (all $p < 0.0001$), and the test statistics ranged from -5.49 (Oils) to -13.65 (Sugar). In the same vein, the KPSS test failed to reject the null hypothesis of stationarity for all series (all p -values = 0.10), with test statistics well below critical values at the 5% significance level. Concerns regarding spurious regression are eliminated, and the suitability of further econometric modeling is confirmed by the strong confirmation of stationarity provided by the concordance between the two tests.



Figure 4. Food Price Index and Subsectors—Log Transformed Data.




Figure 5. Food Price Index and Subsectors—Log-differenced series.

Table 4. Stationarity Test Results on Log-Differenced Transformed Data.

Variable	ADF Stat	ADF (<i>p</i> -Value)	KPSS Stat	KPSS (<i>p</i> -Value)	Status
Food Price Index	−9.7067	1.04×10^{-16}	0.0607	0.10	Stationary
Meat	−5.6639	9.25×10^{-7}	0.0894	0.10	Stationary
Dairy	−6.5884	7.22×10^{-9}	0.0314	0.10	Stationary
Cereals	−13.1487	1.38×10^{-24}	0.0501	0.10	Stationary
Oils	−5.4912	2.17×10^{-6}	0.0357	0.10	Stationary
Sugar	−13.6493	1.60×10^{-25}	0.0629	0.10	Stationary

3.2. Multicollinearity

To address multicollinearity, three stages of multicollinearity diagnostics were carried out using Variance Inflation Factors (VIF): raw data, log-transformed data, and log-differenced data as shown in **Table 5**. With VIF values ranging from 20.44 (sugar) to 98.50 (Cereals), preliminary analysis of the raw data indicated significant

multicollinearity issues. Severe multicollinearity, which inflates standard errors and destabilizes coefficient estimates, is indicated by values greater than 10. Log transformation unexpectedly made the issue worse, with VIF values sharply rising—Cereals' VIF rose to 1929.58. This happened because the logarithmic transformation compresses the scales of the original series while maintaining their co-trending structure.

Table 5. Variance Inflation Factor (VIF) for Different Transformations.

Variable	VIF (Raw Data)	VIF (Log Transformed)	VIF (Log-Diff.)
Meat	30.49	565.54	1.02
Dairy	46.27	736.05	1.08
Cereals	98.50	1929.58	1.26
Oils	46.52	837.98	1.30
Sugar	20.44	371.30	1.04

First-order differencing of the log-transformed data significantly reduced multicollinearity. Post-differencing VIF values varied from 1.02 for Meat to 1.30 for Oils, remaining significantly below the conservative threshold of 5 and nearing the optimal value of 1 as seen in **Table 5**, column 4. This significant enhancement indicates that differencing successfully eliminates the

shared stochastic trends that cause high multicollinearity, thereby isolating the unusual variations in each subsector that hold distinct predictive information. The low VIF values in the final specification indicate that the estimated regression coefficients are stable, interpretable, and not inflated by redundant information among predictors. **Table 5** is graphically represented in **Figure 6**.




Figure 6. VIF Transformation Comparison.

To verify the diminished multicollinearity, a heatmap of the correlation matrix among the dependent variables is presented in **Figure 7**. A correlation analysis was conducted on 'Oils', 'Sugar', 'Dairy', and 'Meat'. **Figure 7** illustrates a significant positive reduction in correlations among the majority of variables when compared to the heatmap of the correlation matrix of the raw data in **Figure 3**.




Figure 7. A heatmap of the correlation matrix of Log-difference series.

3.3. Model Performance and Evaluation

After data preprocessing, the dataset was divided into training (January 1990–December 2023; 408 observations) and testing (January 2024–March 2025; 14 observations) subsets. Five modeling approaches were implemented and assessed: Multiple Linear Regres-

sion (MLR), Artificial Neural Network (ANN), Random Forest (RF), Support Vector Machine (SVM), and Long Short-Term Memory (LSTM) networks. All models were trained on standardized log-differenced data and assessed using Root Mean Squared Error (RMSE), Mean Absolute Error (MAE), Mean Squared Error (MSE), and

R-squared (R^2) metrics.

3.3.1. Multiple Linear Regression Fitting on All Variables

Ordinary Least Squares (OLS) was used to estimate the multiple linear regression (MLR) model on log-differenced, standardized training data. All subsector predictors exhibited statistically significant relationships with the Food Price Index (FPI) ($p < 0.0001$), as shown in **Table 6**. The table shows that the Food Price Index is influenced by factors which include meat, dairy, cereals, oils, and sugar because each has p -values less than 0.05. The results of the multiple linear regression analysis on the Food Consumer Price Index (FCPI) indicate that all independent variables—Meat, Dairy, Cereals, Oils, and Sugar—have statistically significant positive relationships with FCPI, as evidenced by their extremely low p -values ($p < 0.0001$). Among the predictors, Meat ($\beta = 0.336$), Oils ($\beta = 0.3993$) and Cereals ($\beta = 0.4896$) exert the largest effects on changes in food prices, followed by Dairy ($\beta = 0.2031$), and Sugar ($\beta = 0.2060$). Keeping other subsectors constant, the esti-

mated coefficients indicate that a one-percentage-point increase in a subsector's log return results in a proportionate change in the Food Price Index. The t -values for all significant variables are exceptionally high, indicating strong evidence against the null hypothesis. Additionally, the 95% confidence intervals for these variables are very narrow and stable, reflecting high precision in the estimation of their coefficients. These findings suggest that price increases in these food sub-sectors are strong drivers of overall changes in Nigeria's FCPI, highlighting the need to monitor and possibly regulate these key food categories to stabilize food prices.

The regression model's overall significance in **Table 7**, ($F = 6549.98, p < 0.0001$), was confirmed by analysis of variance. The five subsector predictors account for 98.79% of the variance in the log-differenced Food Price Index, with an R^2 of 0.9879 and an Adjusted R^2 of 0.9877. The small difference (0.0002) between R^2 and Adjusted R^2 indicates that every predictor included makes a significant contribution to model fit without overfitting.

Table 6. Coefficient Significance Tests for Multiple Linear Regression.

Variable	Coefficient	Std. Error	t -Statistic	p -Value	CI Lower	CI Upper
Intercept	-0.0000	0.0055	-0.0000	1.0000	-0.0108	0.0108
Meat	0.3442	0.0056	61.9281	0.0000	0.3333	0.3551
Dairy	0.2031	0.0057	35.6549	0.0000	0.1919	0.2143
Cereals	0.4896	0.0062	78.5943	0.0000	0.4774	0.5019
Oils	0.3993	0.0063	63.2622	0.0000	0.3869	0.4117
Sugar	0.2060	0.0056	36.7660	0.0000	0.1949	0.2170

Table 7. ANOVA Table for Multiple Linear Regression.

Source	DF	Sum Squares	Mean Square	F-Statistic	p -Value
Regression	5	403.0526	80.6105	6549.9778	0.0000
Residual	402	4.9474	0.0123	-	-
Total	407	408.0000	-	-	-

Note: R^2 : 0.9879; Adjusted R^2 : 0.9877.

3.3.2. MLR Model Diagnostic

The MLR model's assumptions were validated through diagnostic testing. The Durbin-Watson statistic of 1.5122 indicates no autocorrelation in residuals. The Breusch-Pagan test supported the null hypothesis of homoscedasticity with a statistic of 7.4281 ($p = 0.1907$), confirming constant error variance. However, the Shapiro-Wilk test rejected the null hypothesis of normality, with a statistic of 0.9383 ($p < 0.0001$). Despite

this, the large sample size ($n = 408$) allows for the Central Limit Theorem to ensure asymptotic normality of coefficient estimates, making the normality deviation typical in financial time series data and not compromising regression estimates.

The model's adequacy is visually confirmed by the diagnostic plots in **Figure 8**. The central part of the Q-Q plot closely resembles the theoretical normal distribution, but the tails show minor deviations from normal-

ity, which are consistent with the Shapiro-Wilk test results. The linearity and homoscedasticity assumptions are confirmed by the residuals versus fitted values plot, which displays no discernible pattern. Homoscedastic-

ity is further supported by the scale-location plot's comparatively constant spread. With a slight negative skewness, the residuals histogram resembles a normal distribution.




Figure 8. MLR Diagnostic.

ANN

The ANN model used L2 weight penalties ($\lambda = 0.001$), dropout regularization (rate = 0.2), and ReLU activation functions in a feedforward architecture with three hidden layers (64, 32, and 16 neurons, respectively). The model converged after about 80–100 epochs, after being trained over 200 epochs with early stopping (patience = 20). The ANN explained 96.17% of the variance in log-differenced food prices on the test set with RMSE = 0.0020, MAE = 0.0018, and $R^2 = 0.9617$.

Random Forest

The Random Forest regressor was set up with 100 trees, a minimum of five samples per split, a maximum depth of ten, and square-root feature selection at every node. These hyperparameters achieved a balance between the ability to generalize and the complexity of the model. The RF model explained 89.76% of the vari-

ance with test set performance of RMSE = 0.0033, MAE = 0.0027, and $R^2 = 0.8976$.

The feature importance from the Random Forest model, as shown in **Figure 9**, highlights Cereals and Oils as the most influential predictors, with importance scores of 0.375 and 0.325, respectively, while Meat, Dairy, and Sugar contribute minimally, indicating their limited role in the model's predictions. This suggests that cereals and oils drive the random forest model's outcomes more than the other subsectors.

Support Vector Machine (SVM)

For the SVM machine learning model, a Radial Basis Function (RBF) kernel with regularization parameters $C = 1.0$, $\epsilon = 0.1$, and automatic gamma scaling was used, and it achieved RMSE = 0.0015, MAE = 0.0012, and $R^2 = 0.9779$, accounting for 97.79% of the variance. The RBF kernel successfully captures subtle non-linearities in

the log returns, as evidenced by the SVM's marginal superiority over MLR ($\Delta R^2 = 0.0028$). The detection of intricate relationships between subsector price changes that are impossible to capture with straightforward lin-

ear combinations was made possible by the kernel's capacity to implicitly map data into higher-dimensional feature spaces. This supports the body of research showing that SVM models perform best with small sample sizes.




Figure 9. Feature Importance(RF).

Long Short-Term Memory (LSTM)

Using a 12-month lookback for possible seasonal patterns, the LSTM architecture used two layers (50 units each) with dropout (0.2) and a dense output layer. The model performed poorly, with RMSE = 0.0114, MAE = 0.0089, and $R^2 = -0.2404$, suggesting lower accuracy compared to mean predictions, despite being trained for 100 epochs with early stopping. This deficiency is explained by the fact that although differencing eliminates trends, decreasing its effectiveness for stationary fluctuations, LSTM is designed for long-term dependence in trending data.

3.4. Making Predictions with the Different Models under Study

The model rankings and performance metrics differ significantly between the log-differenced and raw scales.

Performance on Log-Differenced Scale (Predicting Percentage Change)

On the log-differenced scale, SVM, MLR, and ANN demonstrated excellent performance, indicating that

they were highly effective at predicting monthly percentage changes in the Food Price Index, whereas the LSTM model performed very poorly, even yielding a negative coefficient of determination (R^2), which indicates performance worse than a naive mean-based benchmark.

Table 8 reports the performance metrics of our models. The Support Vector Machine (SVM) model achieves the strongest predictive accuracy, recording the lowest RMSE (0.001515988), lowest MAE (0.001170473), and the smallest MSE (2.30×10^{-6}), along with the highest R^2 value (0.977872671). The Multiple Linear Regression (MLR) model follows closely, performing competitively with similarly low error values and an R^2 of 0.975091613. The Artificial Neural Network (ANN) ranks next, with moderate error levels and an R^2 of 0.961668026, indicating solid but slightly weaker performance relative to SVM and MLR. The Random Forest model shows noticeably higher errors (RMSE = 0.003261001; MAE = 0.002744329) and a lower R^2 of 0.897614264, reflecting reduced predictive precision. In contrast, the LSTM model performs the worst, exhibiting substantially larger errors (RMSE = 0.011350446; MSE = 1.28833×10^{-4}) and a negative R^2

(−0.240405187), indicating that it does not generalize well and performs worse than a simple baseline model.

The Model comparison bar chart of **Figure 10** visually confirms these rankings across all evaluation metrics.

Table 8. Ranked Model Performance on Log-Differenced Scale.

Rank	Model	RMSE	MAE	MSE	R^2
1	SVM	0.001515988	0.001170473	2.30×10^{-6}	0.977872671
2	MLR	0.001608437	0.001306828	2.59×10^{-6}	0.975091613
3	ANN	0.001995316	0.001760463	3.98×10^{-6}	0.961668026
4	Random Forest	0.003261001	0.002744329	1.06×10^{-5}	0.897614264
5	LSTM	0.011350446	0.008933783	1.28833×10^{-4}	−0.240405187




Figure 10. Model Comparison of the log-difference data.

Support Vector Machines (SVM) and Multiple Linear Regression (MLR) both obtain above 97% variance explained for Nigerian food price indices, according to performance analysis. This is because SVM uses an RBF kernel to leverage non-linear relationships; hence, it performs marginally better than MLR, whereas MLR verifies almost linear relationships after transformation. Furthermore, Artificial Neural Networks (ANN) perform well for complicated patterns ($R^2 = 0.9617$), although the benefits of greater complexity are negligible. Random Forest underperforms ($R^2 = 0.8976$) because of inefficiencies in partitioning approaches compared to gradient-based or kernel methods. Finally, LSTM models exhibit a mismatch for the interpretation of differenced data, where their temporal memory is ineffectual.

3.5. Predictions on Raw Data

Log-differenced predictions were converted back to the original Food Price Index (FPI) scale using the recursive formula:

$$FPI_t = FPI_{t-1} \times \exp(\Delta \log(FPI_t))$$

This approach helps in interpretation and practical application. The reconstructed predictions and actual values for the 14-month test period (January 2024–February 2025) are shown in **Table 9**. **Figures 11–14** illustrate how SVM, MLR, and ANN closely follow real price movements, while Random Forest exhibits a moderate lag and LSTM diverges significantly.

Table 9. Model Performance on Raw Data.

Rank	Model	R ² _Raw	RMSE_Raw	MAE_Raw
1	Random Forest	0.9342	0.9063	0.8118
2	MLR	0.9081	1.0714	0.9647
3	ANN	0.8708	1.2705	1.1905
4	SVM	0.8356	1.4328	1.4150
5	LSTM	-0.6888	4.5929	3.9060




Figure 11. Predictions Original Scale.




Figure 12. Actual vs. Predicted (Raw data).




Figure 13. Predictions of Eaw data (original data series by Different Models).




Figure 14. Model Comparison After Reconversion.

3.6. Actual Versus Predicted for Different Models at Different Scaling

The predictions of our log-differenced series are displayed in **Figure 15**, which was the scale employed to train the model. It helps examine consistency prior to reversion and evaluate performance in the changed space.

The prediction errors for each model at each time point are included in **Table 10**. It helps determine which

models work best by displaying the degree to which each model deviates from the actual value.

The same predictions are shown in **Tables 11** and **12**. Following the conversion of the log-differenced predictions back to the original scale, this table displays the actual observed values with forecasts from various models (MLR, ANN, RF, SVM, and LSTM). This is used to compare how well each model's predictions match the actual data.




Figure 15. Predictions of log-differenced series.

Table 10. Step 6: Error Analysis–Original Scale (Log-Differenced).

Date	MLR_Error	ANN_Error	RF_Error	SVM_Error	LSTM_Error
2024-02-01	0.1199	0.3062	0.4419	0.0008	0.0218
2024-03-01	0.1642	0.2773	0.1718	-0.0106	1.7092
2024-04-01	0.0580	0.0123	0.3650	-0.1531	2.1732
2024-05-01	-0.1386	-0.3580	0.2375	-0.2619	3.5273
2024-06-01	-0.0127	-0.3159	0.4314	-0.2078	3.9509
2024-07-01	0.1709	-0.1136	0.9975	-0.0989	3.6678
2024-08-01	0.2872	-0.2414	1.1790	0.0532	4.2832
2024-09-01	0.3544	-0.0478	1.5624	-0.1482	6.9423
2024-10-01	0.5755	0.3787	2.2701	-0.0245	8.8912
2024-11-01	1.0776	0.7601	2.2711	0.4772	9.1904
2024-12-01	1.0684	0.8885	2.2450	0.4209	8.3628
2025-01-01	0.9444	0.6704	1.7659	0.1781	5.5504
2025-02-01	1.1379	0.8795	2.4864	0.2728	7.0690
2025-03-01	1.4459	1.0324	3.0053	0.5252	6.5290

Table 11. Predictions Table (Raw Scale: After Reconversion).

Date	Actual	MLR_Pred	ANN_Pred	RF_Pred	SVM_Pred	LSTM_Pred
2024-02-01	117.4487	118.7462	118.5576	118.4203	118.8667	118.8455
2024-03-01	118.9467	120.2175	120.1029	120.2098	120.3944	118.6538
2024-04-01	119.2498	120.6317	120.6779	120.3210	120.8453	118.4909
2024-05-01	120.4892	122.0850	122.3070	121.7043	122.2097	118.3748
2024-06-01	120.9907	122.4651	122.7720	122.0156	122.6625	118.4537
2024-07-01	120.8818	122.1691	122.4570	121.3326	122.4422	118.6300
2024-08-01	121.6681	122.8472	123.3822	121.9447	123.0841	118.8030
2024-09-01	124.5897	125.7360	126.1430	124.5134	126.2447	119.0686
2024-10-01	126.9398	127.8908	128.0900	126.1757	128.4980	119.4746
2024-11-01	127.7120	128.1641	128.4855	126.9562	128.7718	119.9533
2024-12-01	127.4096	127.8674	128.0494	126.6766	128.5227	120.4849
2025-01-01	125.1045	125.6599	125.9372	124.8285	126.4354	120.9983
2025-02-01	127.1206	127.5045	127.7661	126.1398	128.3801	121.5018
2025-03-01	127.1000	127.1720	127.5905	125.5937	128.1038	122.0275

Table 12. Predictions Table–Log-Differenced.

Date	Actual	MLR	ANN	RF	SVM	LSTM
2024-02-01	117.4487	117.3288	117.1426	117.0068	117.4479	117.4269
2024-03-01	118.9467	118.7826	118.6694	118.7750	118.9574	117.2375
2024-04-01	119.2498	119.1918	119.2375	118.8848	119.4029	117.0766
2024-05-01	120.4892	120.6278	120.8472	120.2516	120.7510	116.9619
2024-06-01	120.9907	121.0034	121.3066	120.5592	121.1984	117.0398
2024-07-01	120.8818	120.7109	120.9954	119.8844	120.9808	117.2140
2024-08-01	121.6681	121.3809	121.9096	120.4892	121.6149	117.3849
2024-09-01	124.5897	124.2353	124.6374	123.0273	124.7379	117.6474
2024-10-01	126.9398	126.3643	126.5612	124.6697	126.9643	118.0486
2024-11-01	127.7120	126.6343	126.9519	125.4409	127.2348	118.5216
2024-12-01	127.4096	126.3412	126.5211	125.1646	126.9887	119.0468
2025-01-01	125.1045	124.1600	124.4340	123.3386	124.9263	119.5541
2025-02-01	127.1206	125.9826	126.2411	124.6342	126.8477	120.0516
2025-03-01	127.1000	125.6541	126.0676	124.0947	126.5748	120.5710

4. Conclusions, Recommendations and Future Directions

4.1. Findings and Conclusion

Significant violations of the classical regression assumptions were resolved by the very successful log-differencing transformation. The residuals were homoscedastic (Breusch–Pagan $p = 0.1907$), stationarity was attained (ADF $p < 0.0001$; KPSS $p > 0.10$), multicollinearity was removed (VIF decreased from 98.50 to < 1.31), and no first-order autocorrelation was found (Durbin–Watson = 1.5122). The large sample size guaranteed that the inference remained valid even though the Shapiro–Wilk test indicated non-normality, which is typical in financial data.

Model performance proved highly sensitive to the evaluation scale. SVM had the highest predictive accuracy ($R^2 = 0.9779$) on the log-differenced scale, closely followed by MLR and ANN. However, Random Forest outperformed both MLR and SVM when predictions were converted back to the original Food Price Index scale using the recursive exponential transformation ($R^2 = 0.9342$; RMSE = 0.9063). This reversal happened because Random Forest’s ensemble averaging generates more stable multi-step forecasts than SVM’s occasionally sharp deviations, and recursive reversion increases single-step prediction errors.

These findings demonstrate a fundamental trade-off between ensemble-based cumulative stability and kernel-based single-step precision. SVM is very good at predicting stationary log returns, but when converted

back to the original scale, its tiny high-magnitude errors are amplified exponentially. The variance-reducing ensemble structure of Random Forest, on the other hand, produces smoother, self-correcting forecasts that reduce long-horizon error accumulation, which is essential for policy-oriented forecasting on the original scale.

Multiple Linear Regression ranked second on both the original and transformed scales, indicating remarkable reliability. It is well-suited for applications where comprehending sectoral contributions to price changes is just as important as predictive accuracy due to its robust and reliable performance and interpretable coefficients. Further evidence that log-differencing linearizes the underlying relationships and favors simpler models comes from the small R^2 difference between SVM and MLR on the transformed scale ($\Delta R^2 = 0.0028$).

With MLR coefficients of 0.336 and 0.308, respectively, which jointly account for 64.4% of the variation, variable significance analysis revealed that meat and cereals are the main drivers of food price dynamics in Nigeria. Sugar (0.069) has a negligible direct influence, followed by oils (0.183) and dairy (0.123). The rankings were somewhat rearranged by Random Forest analysis, with dairy (52%) and cereals (41%) emerging as the most important predictors. Different mechanisms of analysis are reflected in the rankings between the MLR and RF techniques. Notably, Sugar’s little contribution indicates that it is heavily impacted by the other sub-sectors, emphasizing the necessity of keeping an eye on sugar prices while concentrating on Meat and Cereals for successful price stability measures.

Lastly, LSTM had negative R^2 values on both scales

and performed poorly on all assessments. Differencing makes the architecture inappropriate for stationary log-return data because it eliminates trends and long-term dependencies, which are exactly the patterns LSTMs are meant to learn. The inconsistency was made worse by the long retrospective window, which added noise instead of helpful temporal structure. This shows that

when temporal memory is not useful, LSTMs should not be applied to differenced economic series.

4.2. Recommendation

Based on our findings and the results obtained, we recommend the following **Table 13**.

Table 13. Model Performance Comparison Across Evaluation Scales.

Model	R ² (Log-Diff)	Rank (Log-Diff)	R ² (Original)	Rank (Original)	MAPE (%)	Recommendation
Random Forest	0.8976	4	0.9342	1	0.65	Primary model
MLR	0.9751	2	0.9081	2	0.78	Best for interpretation
ANN	0.9617	3	0.8708	3	0.96	Alternative model
SVM	0.9779	1	0.8356	4	1.14	Short-term only
LSTM	−0.2404	5	−0.6888	5	3.15	Not recommended

With Mean Absolute Percentage Errors (MAPE) significantly below 1% on original-scale predictions, our analysis shows that the best-performing models attained remarkable accuracy:

- In comparison to standard benchmarks for commodity price forecasting, our top three models—Random Forest (0.65%), MLR (0.78%), and ANN (0.96%)—achieved Mean Absolute Percentage Errors (MAPE) below 1%. While conventional ARIMA models obtain 2–3% MAPE, naive approaches typically reach 3–5%. These models' remarkable accuracy is essential for a number of applications: they can help predict price spikes for policy interventions, support NGOs in allocating resources for humanitarian needs, stabilize agricultural markets, and safeguard household welfare through social programs. Over the course of a 14-month testing period, the Random Forest and MLR models demonstrated strong accuracy in capturing Nigeria's food price patterns, offering decision-makers insightful information. According to Makridakis et al.^[41], MAPE values below 5% are regarded as acceptable and below 2% as excellent. These results greatly exceed industry benchmarks. Our models provide unprecedented precision when compared to naive forecasting (3–5% MAPE) and ARIMA models (2–3% MAPE), allowing for dependable short-term forecasting for:

- **Timing of Policy Intervention:** For proac-

tive reserve allocation or import adjustments, price spikes should be anticipated six to twelve months in advance.

- **Humanitarian Response Planning:** Enabling non-governmental organizations and aid agencies to distribute resources prior to the escalation of crises.
- **Agricultural Market Stabilization:** Setting off early warning systems for emergency procurement or price support mechanisms.
- **Household Welfare Protection:** Modifying social assistance payments to account for rising food prices.
- The less than one percent MAPE for Random Forest and MLR shows that these models accurately capture the key features of Nigeria's food price system to give decision-makers useful information. The model's robustness under various market situations is further validated by the consistent accuracy over the 14-month test period (January 2024–March 2025) in **Table 11**.

- Our analysis reveals a trade-off between multi-step stability and single-step precision in forecasting models. By transferring data into higher-dimensional spaces, the SVM's RBF kernel performs exceptionally well in individual log return predictions ($R^2 = 0.9779$), which results in greater deviations during reconversion. On the other hand,

Random Forest's ensemble method averages predictions from 100 decision trees, performing better in cumulative accuracy on the original scale (MAPE = 0.65% versus SVM's 1.14%), hence providing more consistent results with lower volatility. This suggests that projections should be assessed using the original measurement scales because end users may be misled if converted metrics are the only ones used. Even while ensemble approaches perform poorly in single-step forecasts, their robustness is advantageous in recursive forecasting.

- The following forecasting model recommendation is based on a thorough evaluation: Because of its accuracy and robustness, Random Forest ($R^2 = 0.9342$, MAPE = 0.65%) is ideal for operational forecasting, making it appropriate for humanitarian responses and policy planning. For policy analysis that requires transparent coefficient estimates, Multiple Linear Regression ($R^2 = 0.9081$, MAPE = 0.78%) is favored, sacrificing minimal accuracy for increased interpretability. As an alternative or for ensemble averaging, Artificial Neural Network ($R^2 = 0.8708$, MAPE = 0.96%) offers a useful third viewpoint despite having lesser accuracy. Although Support Vector Machine performs well in log terms ($R^2 = 0.9779$), its low precision ($R^2 = 0.8356$, MAPE = 1.14%) makes it impractical for original-scale forecasts. Long Short-Term Memory is inadequate for real-world applications using differenced data due to its low original-scale performance (negative R^2 and 3.15% MAPE), high processing costs, and lack of interpretability.
- The issue of multicollinearity among food subsector variables should be resolved by applying techniques such as VIF and other transformations to enhance the stability and interpretability of the regression model.
- Accurate forecasting needs good quality data. Government agencies, researchers, and market watchers should focus on collecting food price data more often (like every week) and from different regions. This will help the models work better and also give a clearer picture of what's happening in different

parts of the country, and this may take care of multicollinearity.

- Based on reliable predictive models, stakeholders should build automated food price monitoring and early warning systems to proactively manage food inflation, mitigate supply shocks, and guide targeted interventions.

4.3. Limitation and Future Direction

Although the study on Nigerian food price forecasting provides insightful information, it has a number of drawbacks. The training dataset might be too small for machine learning models, and the limited testing period of 14 months might not accurately represent long-term price trends. Furthermore, the models' explanatory power was impacted by the exclusion of unpredictable elements such as economic factors, climate conditions, policy actions, security concerns, and global markets. Again, because of structural changes in the food system, the assumptions of stable correlations across food subsector indices may not be true, and the model's performance might be impacted by incomplete hyperparameter tuning. Also, the evaluation of model robustness is limited by the use of a single train-test split, and the absence of investigation into ensemble approaches may hinder the best possible prediction results. These restrictions should be addressed in future studies by exploring More extensive datasets, external factors, and sophisticated modeling approaches should be used in future studies to overcome these constraints.

Using FAO data, our findings and conclusions are unique to Nigeria's Food Price Index. It's possible that our best-performing models' relative superiority over other more intricate models won't hold true for different nations, food indices, or forecasting scenarios. Alternative modeling methodologies may be preferred by different economic structures, data quality levels, or price volatility characteristics. Our conclusions regarding the Food Price Index are specific to Nigeria and FAO data. The relative performance of our models may vary for other countries, food indices, or forecasting scenarios, as different economic structures, data quality, or price volatility could favor alternative approaches. To improve generalization, cross-national studies should be

conducted to assess the suitability of ensemble models for simulating subsector price interdependencies in different contexts.

Author Contributions

Conceptualization, A.A.D. and C.P.O.; methodology, C.P.O.; software, A.A.D.; validation, M.W.H. and I.A.; formal analysis, A.S.; investigation, M.M.J.F.; resources, M.S.K.; data curation, C.P.O.; writing—original draft preparation, A.A.D.; writing—review and editing, C.P.O.; visualization, A.A.D.; supervision, M.S.K.; project administration, I.A. All authors have read and agreed to the published version of the manuscript.

Funding

This work received no external funding.

Institutional Review Board Statement

Not applicable.

Informed Consent Statement

Not applicable.

Data Availability Statement

The data that support the findings of this study are publicly available and were obtained from the Food and Agriculture Organization of the United Nations (FAO). Specifically, the Food Price Index data used in this research were retrieved from the FAO World Food Situation portal at <https://www.fao.org/worldfoodsituation/foodpricesindex>. These datasets are openly accessible and can be downloaded directly from the FAO website. No new primary data were generated in this study.

Conflicts of Interest

The authors declare no conflict of interest.

References

- [1] Akande, E.O., Akanni, E.O., Taiwo, O.F., et al., 2022. Predicting Inflation Component Drivers in Nigeria: A Stacked Ensemble Approach. *SN Business & Economics*. 3(1), 9.
- [2] Nwafor, M., Lodugnon-Harding, J., Tuyishime, C., et al., 2025. Nigeria Food and Agriculture Policy Monitoring Review: FAO Monitoring and Analysing Food and Agricultural Policies Programme. Food and Agriculture Organization: Rome, Italy.
- [3] National Bureau of Statistics, 2021. Labor Force Statistics: Unemployment and Underemployment Report (Q4 2020). National Bureau of Statistics: Abuja, Nigeria.
- [4] Shan, S., 2024. Identification of Socioeconomic Factors Influencing Global Food Price Security Price Security Using Machine Learning. *arXiv preprint*. arXiv:2403.04231. DOI: <https://doi.org/10.48550/arXiv.2403.04231>
- [5] Ikuemonisan, E.S., Akinbola, A.E., 2021. Future Trends in Cassava Production: Indicators and Its Implications for Food Supply in Nigeria. *Asian Journal of Agricultural Extension, Economics & Sociology*. 39(3), 60–74.
- [6] Kupferschmidt, K.L., Requiema, J., Simpson, M., et al., 2024. Food for Thought: How Can Machine Learning Help Better Predict and Understand Changes in Food Prices? *arXiv preprint*. arXiv:2412.06472. DOI: <https://doi.org/10.48550/arXiv.2412.06472>
- [7] Ulussever, T., Ertuğrul, H.M., Kılıç Depren, S., et al., 2023. Estimation of Impacts of Global Factors on World Food Prices: A Comparison of Machine Learning Algorithms and Time Series Econometric Models. *Foods*. 12(4), 873.
- [8] Majhi, S.K., Bano, R., Pradhan, R., et al., 2024. Food Price Index Prediction Using Time Series Models: A Study of Cereals, Millets and Pulses. In *International Conference on Technology Advances for Green Solutions and Sustainable Development*. Springer: Singapore. pp. 216–238.
- [9] Ingio, J.A., Nsang, A.S., Iorliam, A., 2024. Comparing Machine Learning Algorithms for Rice Yield Prediction in Adamawa and Cross Rivers States of Nigeria. *Gazi University Journal of Science Part A: Engineering and Innovation*. 11(3), 481–496.
- [10] Nayak, G.H., Alam, M.W., Singh, K., et al., 2024. Exogenous Variable Driven Deep Learning Models for Improved Price Forecasting of TOP Crops in India. *Scientific Reports*. 14(1), 17203.
- [11] Sanusi, O.I., Safi, S.K., Adeeko, O., et al., 2022. Forecasting Agricultural Commodity Price Using Different Models: A Case Study of Widely Consumed Grains in Nigeria. *Agricultural and Resource Eco-*

- nomics: International Scientific E-Journal. 8(2), 124–140.
- [12] Sun, Y., Wang, X., Zhang, C., et al., 2023. Multiple Regression: Methodology and Applications. *Highlights in Science, Engineering and Technology*. 49, 542–548.
- [13] Danbatta, S.J., Muhammad, A., Varol, A., et al., 2025. Forecasting Monthly Rainfall Using Hybrid Time-Series Models and Monte Carlo Simulation Amidst Security Challenges: A Case Study of Five Districts from Northern Nigeria. *Environment, Development and Sustainability*. 27(6), 13815–13837.
- [14] Eticha, A., 2020. Multiple Linear Regressions on Determinants of Ginger Production in Yeki District, Sheka Zone, South West Ethiopia. *International Journal of Agricultural Science and Food Technology*. 6(2), 151–156.
- [15] Wu, T., Yu, J., Lu, J., et al., 2020. Research on Inversion Model of Cultivated Soil Moisture Content Based on Hyperspectral Imaging Analysis. *Agriculture*. 10(7), 292.
- [16] Istifaroh, E.P., Harini, R., 2021. The Effect of Land Characteristics on Agricultural Productivity Using Multiple Linear Regression Method in Madiun Regency. *E3S Web of Conferences*. 325, 02005.
- [17] Tule, M.K., Salisu, A.A., Chiemeké, C.C., 2019. Can Agricultural Commodity Prices Predict Nigeria's Inflation? *Journal of Commodity Markets*. 16, 100087.
- [18] Das, D., Chakrabarti, S., 2023. An Extreme Learning Machine Approach for Forecasting the Wholesale Price Index of Food Products in India. *Pertanika Journal of Science & Technology*. 31(6), 3179–3198.
- [19] Jin, B., Xu, X., 2025. Wholesale Price Forecasts of Green Grams Using the Neural Network. *Asian Journal of Economics and Banking*. 9(3), 463–490.
- [20] Choong, K.Y., Sudin, S., Raof, R.A.A., et al., 2023. Hybrid Approach for Vegetable Price Forecasting in Electronic Commerce Platform. *International Journal of Artificial Intelligence*. 13(2), 1858–1867.
- [21] Paul, R.K., Yeasin, M., Kumar, P., et al., 2022. Machine Learning Techniques for Forecasting Agricultural Prices: A Case of Brinjal in Odisha, India. *Plos One*. 17(7), e0270553.
- [22] Tyrallis, H., Papacharalampous, G., Langousis, A., 2019. A Brief Review of Random Forests for Water Scientists and Practitioners and Their Recent History in Water Resources. *Water*. 11(5), 910.
- [23] Malysheva, T., Panachev, A., Medvedeva, M., et al., 2019. Application of Random Forest Algorithm to Predict the Average Issued Amounts in ATMs. In *Proceedings of the International Conference of Computational Methods in Sciences and Engineering 2019 (ICCMSE-2019)*, Rhodes, Greece, 1–5 May 2019. DOI: <https://doi.org/10.1063/1.5137947>
- [24] Zhang, Y., Ma, J., Liang, S., et al., 2020. An Evaluation of Eight Machine Learning Regression Algorithms for Forest Aboveground Biomass Estimation from Multiple Satellite Data Products. *Remote Sensing*. 12(24), 4015.
- [25] Doz, D., Cotič, M., Felda, D., 2023. Random Forest Regression in Predicting Students' Achievements and Fuzzy Grades. *Mathematics*. 11(19), 4129.
- [26] Dimopoulos, T., Tyrallis, H., Bakas, N.P., et al., 2018. Accuracy Measurement of Random Forests and Linear Regression for Mass Appraisal Models that Estimate the Prices of Residential Apartments in Nicosia, Cyprus. *Advances in Geosciences*. 45, 377–382.
- [27] Vapnik, V., 2013. *The Nature of Statistical Learning Theory*. Springer: New York, NY, USA.
- [28] He, H., Wang, K., Jiang, Y., et al., 2024. Quadratic Hyper-Surface Kernel-Free Large Margin Distribution Machine-Based Regression and Its Least-Square Form. *Machine Learning: Science and Technology*. 5(2), 025024.
- [29] Rodríguez-Pérez, R., Bajorath, J., 2022. Evolution of Support Vector Machine and Regression Modeling in Chemoinformatics and Drug Discovery. *Journal of Computer-Aided Molecular Design*. 36(5), 355–362.
- [30] Gers, F.A., Schmidhuber, J., Cummins, F., 2000. Learning to Forget: Continual Prediction with LSTM. *Neural Computation*. 12(10), 2451–2471.
- [31] Hochreiter, S., Schmidhuber, J., 1997. Long Short-Term Memory. *Neural Computation*. 9(8), 1735–1780.
- [32] Hochreiter, S., Schmidhuber, J., 1994. Simplifying Neural Nets by Discovering Flat Minima. In *Proceedings of the Advances in Neural Information Processing Systems*, Denver, CO, USA, 1 January 1994; pp. 529–536.
- [33] Van Houdt, G., Mosquera, C., Nápoles, G., 2020. A Review on the Long Short-Term Memory Model. *Artificial Intelligence Review*. 53(8), 5929–5955.
- [34] Fackler, P.L., Goodwin, B.K., 2001. Spatial Price Analysis. In: Gardner, B., Rausser, G. (Eds.). *Handbook of Agricultural Economics*. Elsevier: Amsterdam, Netherlands. pp. 971–1024.
- [35] Newbold, P., Harvey, D.I., 2002. Forecast Combination and Encompassing. In: Clements, M.P., Hendry, D.F. (Eds.). *A Companion to Economic Forecasting*. Wiley: Hoboken, NJ, USA.
- [36] Box, G.E.P., Jenkins, G.M., 1976. *Time Series Analysis: Forecasting and Control*. Holden-Day: San Francisco, CA, USA.
- [37] Breiman, L., 2001. Random Forests. *Machine Learning*. 45(1), 5–32.
- [38] Cortes, C., Vapnik, V., 1995. Support-Vector Net-

- works. *Machine Learning*. 20(3), 273–297.
- [39] Campbell, J.Y., Lo, A.W., MacKinlay, A.C., 1997. *The Econometrics of Financial Markets*. Princeton University Press: Princeton, NJ, USA.
- [40] Granger, C.W.J., Newbold, P., 1974. Spurious Regression in Econometrics. *Journal of Econometrics*. 2(2), 111–120.
- [41] Makridakis, S., Spiliotis, E., Assimakopoulos, V., 2017. *The Accuracy of Machine Learning Forecasting Methods Versus Statistical Ones: Extending the Results of the M3-Competition*. University of Nicosia: Nicosia, Cyprus.