Journal of Mathematics and Statistics Studies
ISSN: 2709-4200
DOI: 10.32996/jmss
Journal Homepage: www.al-kindipublisher.com/index.php/jmss

JMSS
AL-KINDI CENTER FOR RESEARCH AND DEVELOPMENT

| RESEARCH ARTICLE

Forecasting Philippine Rice Prices: Comparison of Traditional Time Series and Machine Learning Models

Christony Duyapat[1]✉
[1]*Saint Mary's University, School of Graduate Studies, Nueva Vizcaya, Philippines*
Corresponding Author: Christony Duyapat **E-mail**: hed-cduyapat@smu.edu.ph

| ABSTRACT

The Philippine rice market is characterized by high volatility and a critical impact on national food security, necessitating the use of accurate forecasting tools for price monitoring and risk mitigation. The primary objective of this study was to determine the optimal forecasting methodology for the nominal wholesale price of regular-milled rice. The analysis utilized 430 monthly observations (January 1990 – October 2025), partitioned into an 80% training set (344 data points) and a 20% out-of-sample test set (86 data points). Eight time series models including traditional methods: Seasonal Naïve, ETS, ARIMA, TBATS, Theta, and Prophet, and machine learning algorithms: Random Forest and XGBoost were evaluated. Performance was assessed using Mean Absolute Error (MAE), Root Mean Squared Error (RMSE), and Mean Absolute Percentage Error (MAPE), supplemented by Ljung–Box and Lilliefors tests for residual diagnostics. Evaluation revealed that the Random Forest model achieved the best predictive accuracy, confirming the superior capability of non-linear models to capture the volatile patterns present in the wholesale price data. Residual diagnostics indicated a fundamental trade-off between structural adequacy and predictive accuracy. The final projection forecasts that wholesale prices will stabilize within the Php 34.00/kg to Php 35.50/kg range through 2027. The general price stability predicted for the next two years suggests policy focus may prioritize long-term supply-side improvements rather than short-term demand controls, unless external shocks occur.

| KEYWORDS

non-stationary, ensemble methods, food security, price volatility, time series forecasting, random forest

| ARTICLE INFORMATION

ACCEPTED: 10 December 2025 **PUBLISHED:** 21 December 2025 **DOI:** 10.32996/jmss.2025.6.6.3

1. Introduction

Rice is one of the staple foods in the Philippines. Its wholesale price plays a central role in economic stability, inflation management, and household food security. The significance of rice extends deeply into the political economy; its price can influence election results and overall economic stability, especially given historical government policy challenges in stabilizing prices or improving production (Intal & Garcia, 2005). Because rice is essential to nearly every Filipino family, changes in its price have immediate social and economic implications. Given the high volatility inherent in Philippine rice prices, stabilization through market mechanisms alone is inadequate, necessitating strategic government intervention for economic stability and household food security (Montano, 2025). Consequently, accurate price forecasting becomes an essential tool. Government agencies, such as the Department of Agriculture, rely on timely monitoring and projections to guide buffer stock decisions, set import allocations, and implement measures aimed at maintaining stable and affordable rice prices.

However, generating accurate forecasts for the wholesale price of rice remains a technically demanding task due to the complex underlying dynamics of the time series itself. Montano (2025) suggests that while rice price demonstrates general annual patterns, its notable volatility and persistence mean that recurring seasonal impacts may be minimal and potentially masked by larger underlying shocks and non-linear trends. These complex characteristics challenge many traditional forecasting models that typically assume stable, linear patterns. To address, this study compares an extensive set of models, including both traditional forecasting methods and modern machine learning (ML) approaches, in terms of performance and diagnostic assessments.

This study aims to offer a methodological contribution by investigating a forecasting model that may serve as a useful input for short-to-medium term policy and planning in the Philippine rice sector. To achieve this, the study seeks to examine the overall structure of the monthly wholesale price series of regular-milled rice from January 1990 to October 2025, to assess and compare the predictive performance of different forecasting models, and to identify the best-performing model to generate point forecasts and 95% confidence intervals for wholesale rice prices through December 2027.

1.1 Literature Review

The wholesale price of a commodity represents the price for large-volume transactions intended for further resale or processing (Philippine Statistics Authority, PSA, n.d.). This price is the spot value received by wholesalers for bulk lots, calculated net of discounts, allowances, or rebates, and is economically composed of the producer price, trade margin, tax mark-up, and distribution cost. PSA computes national data from the monthly average wholesale price of regular-milled rice collected across various provinces. Given that wholesale price data represent an aggregated, time-dependent economic series subject to market and policy shocks, accurate prediction requires models capable of handling high volatility and temporal structure.

Forecasting commodity prices is a mature yet rapidly evolving field. Researchers generally rely on three major methodological traditions: classical statistical models such as Autoregressive Integrated Moving Average (ARIMA) and Error-Trend-Seasonality (ETS); decomposition and complex-seasonality methods such as Theta, TBATS (Trigonometric seasonality, Box-Cox transformation, ARMA errors, Trend and Seasonal), and Prophet; and ML and ensemble methods such as Random Forest (RF) and Extreme Gradient Boosting (XGBoost).

Classical time series models remain foundational in commodity price forecasting. ARIMA and its seasonal extension (SARIMA) explicitly model autocorrelation, differencing, and moving-average components. While these are effective for linear and stationary series (Hyndman & Athanasopoulos, 2021), they address non-stationary trends through the differencing process. However, they often struggle when the underlying data exhibit nonlinearity, structural breaks, or shocks, phenomena common in agricultural price series. In the Philippine context, ARIMA and its extensions were reported to have performed adequately in modeling rice production (Caquilala et al., 2025) and wheat moisture content variations for grain storage systems (Rubillos et al., 2024).

As time series data have grown more complex, decomposition and hybrid approaches have gained prominence. The Theta Method, which decomposes a series into modified trend components, has been shown to perform exceptionally well in major forecasting competitions, including the M-Competitions (Hyndman & Billah, 2001). TBATS, an extension of exponential smoothing using Box–Cox transformations and trigonometric seasonality, was specifically developed to model multiple and non-integer seasonality common in long agricultural series (De Livera et al., 2011). Prophet, introduced by Facebook, has also been widely adopted due to its ability to capture nonlinear trends, accommodate user-defined changepoints, and model multiple seasonal patterns effectively, making it suitable for socio-economic variables subject to structural changes (Taylor & Letham, 2018).

In parallel, ML techniques have emerged as powerful tools for forecasting. Unlike traditional models, ML does not assume linearity or fixed parametric relationships and can detect nonlinear interactions and higher-order patterns. RF, introduced by Breiman (2001), uses ensembles of decorrelated decision trees trained on bootstrap samples, resulting in robust and low-variance predictions. XGBoost, a gradient boosting algorithm developed by Chen and Guestrin (2016), builds trees sequentially to minimize prediction error and has achieved state-of-the-art results in a wide range of forecasting applications. Recent applications in agricultural forecasting show that ML ensembles can outperform traditional models when nonlinearities are present (Li et al., 2024).

Despite these advancements, notable gaps remain in literature. Few studies have conducted comprehensive comparisons spanning simple benchmarks (Seasonal Naïve), classical statistical models (ARIMA, ETS), decomposition-based methods (Theta, TBATS, Prophet), and modern ML models (Random Forest, XGBoost) using long historical datasets of Philippine rice prices. To address this gap, the present study offers a general comparison using a long-term (1990–2025) series of Philippine wholesale rice prices, aiming to identify robust forecasting tools that support informed decision-making in rice market management.

2. Materials and Methods

2.1 Data Source and Preprocessing

This study utilized monthly secondary data on the wholesale price of regular-milled rice in the Philippines from January 1990 to October 2025, covering 430 monthly observations. The data were retrieved from Ricelytics, a platform developed by the Data Analytics Center of the Department of Agriculture – Philippine Rice Research Institute (DA–PhilRice) in partnership with DA–National Rice Program. The variable of interest was the nominal wholesale price of regular-milled rice in Philippine pesos per kilogram (Php/kg). Prices were retained in nominal form because the objective of the study was to assess forecasting accuracy rather than analyze changes in real purchasing power.

2.2 Exploratory Analysis and Preprocessing

Initial analysis was performed to understand the underlying structure and patterns in the wholesale price data. Visual inspection through line plots revealed a long-term upward movement. Monthly boxplots highlighted recurring peaks and dips.

These patterns justified the use of forecasting models that incorporate seasonal adjustments. For model estimation and accuracy assessment, the first 80% of the series, equivalent to 344 monthly observations from January 1990 to August 2018 were used as the training dataset. The remaining 20%, composed of 86 observations from September 2018 to October 2025, served as the test dataset. This data split ensured that all forecasts were assessed strictly on data not used during model estimation, thereby preventing information leakage. No manual transformations were applied to the series; instead, any required differencing or detrending was handled internally by the respective models.

2.3 Time Series Models

The study incorporated eight different models with six traditional time series methods: Seasonal Naïve, ETS, ARIMA, TBATS, Theta, and Prophet and two models based on ML techniques: RF and XGBoost. Each model was chosen to represent a distinct modeling philosophy, ensuring comprehensive evaluation of forecasting performance.

The Seasonal Naïve model served as the baseline and predicts each month's value as equal to the corresponding month of the preceding year. Despite its simplicity, it is effective in highly seasonal series and establishes a minimum standard that more sophisticated models are expected to outperform.

The ETS model was estimated using the *ets()* function, which automatically determines the optimal combination of error type (additive or multiplicative), trend structure (none, additive, multiplicative, or damped), and seasonal pattern. ETS models rely on exponentially decreasing weights, giving greater influence to more recent observations. This makes ETS suitable when recent values are more informative for forecasting than earlier historical data.

The ARIMA model was estimated using the *auto.arima()* function, which identifies the optimal combination of autoregressive terms, differencing, and moving-average components automatically based on Akaike Information Criteria (AIC). The *auto.arima()* function can detect and incorporate both seasonal and non-seasonal structures as needed, selecting parameters that best capture temporal dependencies in the data. This approach models the current value of the series as a function of its past observations and forecast errors, making it suitable for time series with autocorrelation and moderate seasonal patterns.

The TBATS model was included to address potentially complex seasonal structures, nonlinear growth patterns, and variance instability. TBATS incorporates multiple features: trigonometric representations of seasonal components, Box–Cox transformation for variance stabilization, ARMA residual adjustment, and trend damping. Its structure makes it flexible for data where seasonality does not maintain a constant magnitude over time, which is a plausible characteristic in agricultural prices.

The Theta method was included due to its strong empirical forecasting performance and theoretical simplicity. It decomposes the time series into modified components known as theta lines, adjusts curvature through a scaling parameter, and recombines them to form forecasts. This method has shown competitive performance in large forecasting competitions and is especially robust for series with linear growth patterns and smooth long-term evolution.

The Prophet model was implemented to capture multiple seasonal cycles through an additive decomposition approach. Prophet supports both nonlinear trends and structural shifts, making it suitable for economic and market data that can experience abrupt changes due to policy adjustments or external shocks such as supply disruptions.

RF was applied using lagged values of the series as predictors. Because RF is an ensemble of regression trees constructed through bootstrapped samples and random feature splits, it captures nonlinear interactions in historical patterns without requiring explicit statistical assumptions such as stationarity.

Finally, XGBoost was implemented using lag-based predictors similar to RF but uses gradient boosting, where successive trees optimize the residual error from previous trees. XGBoost applies regularization penalties that minimize overfitting, enabling robust forecasting even with high-dimensional predictor sets.

Together, these models offer a wide spectrum of predictive methods, from simple seasonal replication to additive decomposition, autoregressive dynamics, nonlinear decision boundaries, and gradient-based optimization, thereby ensuring meaningful performance comparison.

2.4 Model Evaluation and Diagnostics

Models were evaluated using their predictive accuracy on the test set. Three measures were calculated: Mean Absolute Error (MAE), Root Mean Squared Error (RMSE) and Mean Absolute Percentage Error (MAPE). MAE quantifies the average absolute deviation of forecasts from actual prices and serves as a straightforward indicator of typical forecast errors. RMSE emphasizes larger errors by squaring deviations, which is informative in price forecasting, where occasional spikes or drops may occur. MAPE expresses forecast accuracy relative to actual price magnitudes, expressed as a percentage, thereby enabling intuitive interpretation across different price levels. Forecast performance was compared across models, and the optimal model was selected based on overall lower error magnitudes.

Diagnostic assessment was performed to ensure that residuals satisfied key assumptions. The Ljung–Box test was applied to test for autocorrelation in the residuals. Models that produced nonsignificant results were interpreted as adequately capturing serial dependence. The Lilliefors modification of the Kolmogorov–Smirnov test was used to evaluate normality of residuals, a desirable property particularly for probabilistic forecasting and interval estimation.

2.5 Software and Computational Environment

All procedures were executed in R version 4.4.1 within R Studio environment. Data handling and visualization were performed using *tidyverse* packages including *ggplot2* package. Statistical models were fitted using the *forecast* package, Prophet was implemented through the *prophet* package, and ML models were executed using *randomForest* package and *xgboost* package.

3. Results and Discussion

3.1 Exploratory Analysis of Regular-Milled Rice Wholesale Prices

The initial analysis of the monthly wholesale price series, spanning from 1990 to 2025, confirmed two essential characteristics that guided the selection of appropriate time series models: a non-stationary trend and seasonality.

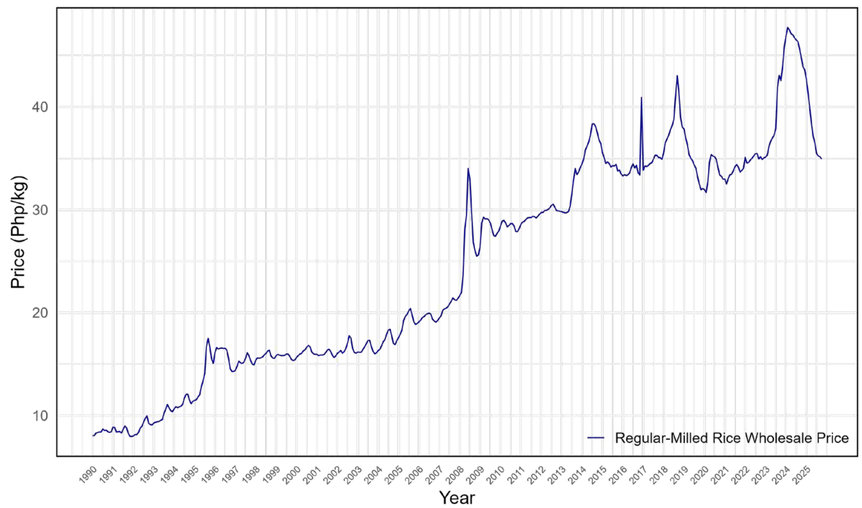

Figure 1. *Monthly Wholesale Price of Regular-Milled Rice in the Philippines (Jan 1990 – Oct 2025)*

Figure 1 shows that the wholesale price of regular-milled rice exhibits a general upward price movement but with fluctuations over the three and a half decades. This is consistent with long-term macroeconomic factors like persistent inflation and rising production costs. The early period (1990 to 2007) showed gradual, relatively stable price growth from approximately Php 8/kg to Php 22/kg. This stability was abruptly broken by the 2008 global price shock, which triggered a massive, near-vertical spike, elevating the price floor to approximately Php 34/kg. The price roughly stabilized at this higher level but maintained increased volatility, culminating in a strong acceleration in the trend visible around 2023, where prices pushed towards the highest recorded values, reaching approximately Php 47/kg. The clear lack of stationarity and the presence of these structural breaks dictated the need for robust modeling techniques capable of adjusting to a changing mean and variance over time.

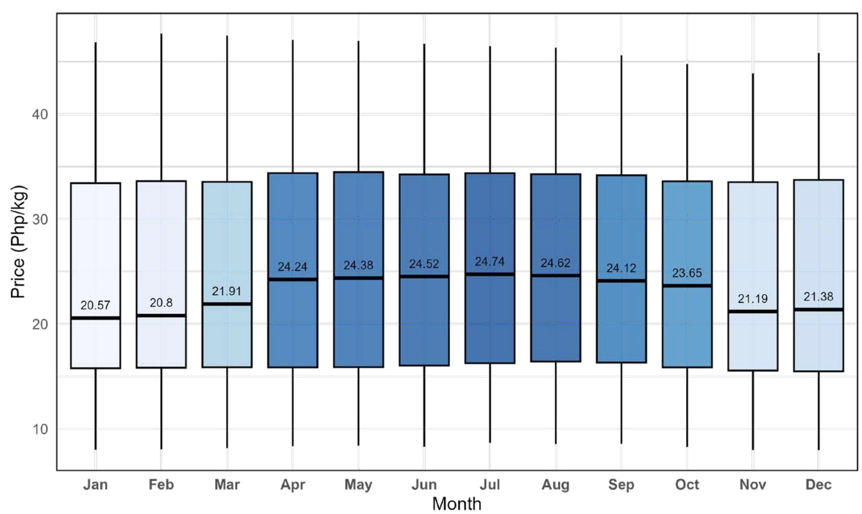

Figure 2. *Boxplot of Monthly Wholesale Price of Regular-Milled Rice*

Figure 2 suggests the existence of a twelve-month seasonal pattern that can be linked to the national rice cultivation cycle. Rice crops in the Philippines are typically cultivated during both the wet and dry seasons, with the Department of Agriculture implementing a crop calendar that defines the planting periods (Cauba et al., 2025). The wet season spans from March 16 to

September 15, while the dry season runs from September 16 to March 15. The planting window for the wet season typically runs from April to August, with peak activity in June-July. The planting window for the dry season runs from September to February, peaking in November-December (Gutierrez et al., 2019).

The overall national pattern observed in Figure 2 shows that the highest median prices from April to October coincide with the lean season, reflecting the general decline in market inventory before the major harvests. Conversely, the lowest median prices are recorded in January (Php 20.57/kg) and February (Php 20.80/kg), which immediately follows the peak wet season harvest and the availability of supply from the main crop. This predictable annual cycle validates the fundamental decision to utilize seasonally-aware time series models to capture the seasonality observed in the regular-milled rice wholesale price.

3.2 Model Fitting

Table 1. Model Specifications and Parameters

Model	Parameters	Notes
1. Seasonal Naïve	No tunable parameters	Baseline model with seasonal frequency = 12
2. ETS	Model: ETS(M,A,A) Alpha = 0.999 Beta = 0.0905 Gamma = 0.0004 Seasonal period = 12	Trend and seasonal components automatically selected by *ets()* using AIC
3. ARIMA	ARIMA(3,1,2) Coefficients: ar1 = 1.6478 ar2 = -0.8620 ar3 = -0.0582 ma1 = -1.7090 ma2 = 0.9309 drift = 0.0943	Model selected automatically by *auto.arima()*; non-seasonal ARIMA(3,1,2)
4. TBATS	Lambda (Box-Cox) = 0.0003 Trend included: TRUE Damping applied: FALSE ARMA errors included: FALSE Seasonal periods: 12	TBATS handles multiple seasonality, damping, and ARMA residuals
5. Theta	Method: Theta Drift: 0.044	Theta model decomposition; robust for trending series
6. Prophet	Trend type: linear Number of changepoints: 25 Seasonality included: monthly, yearly	Prophet decomposes time series into trend, seasonality, and holidays; additive seasonality and automatic changepoints
7. RF	ntree = 500 mtry = 10 nodesize = 3 Number of lagged predictors = 60 Importance computed = TRUE	RF uses lagged predictors for recursive forecasting and captures nonlinear patterns; optimal lag determined through lag-length search; and hyperparameters tuned via grid search
8. XGBoost	nrounds = 1000 max_depth = 12 eta = 0.1 subsample = 0.8 colsample_bytree = 0.9 Objective = reg:squarederror Number of lagged predictors = 72	Gradient-boosted trees with regularization and recursive forecasting using lagged features; optimal lag determined through lag-length search; and hyperparameters tuned via grid search

Table 1 shows the specifications and optimized parameters of all eight forecasting models used in the comparative analysis. The baseline Seasonal Naïve model utilized only the monthly seasonal factor of 12. The automated selection procedure for the ETS model settled on an ETS(M,A,A) structure, combining a Multiplicative (M) error with an Additive (A) trend and Additive (A) seasonality. This combination minimizes AIC and reflects the observed upward trend (Additive Trend) combined with an error term that likely grows proportionally to the series' magnitude (Multiplicative Error). The ARIMA model selection, utilizing the *auto.arima()* function, resulted in the specifications ARIMA(3,1,2) with a seasonal period of 12. This specification indicates a requirement for first-order differencing (d=1) to achieve stationarity, and the resulting structure includes a drift parameter of 0.0943, confirming the presence of a persistent underlying linear trend in the differenced series.

The TBATS model applied a Box–Cox transformation with (λ=0.0003), suggesting the original price scale was nearly optimal for modeling and that variance stabilization was minimally required. The Theta method utilized a small drift parameter of 0.044, which is responsible for projecting the long-term linear trend of the series. The Prophet model automatically identified 25 trend change-points, a high quantity consistent with the multiple structural shifts and periods of pronounced volatility observed historically. This flexible structure allows the model to adapt quickly to changes in the series' mean growth rate.

For the ML models, the RF model was configured with 60 lagged predictors and 500 trees (ntree). The XGBoost model utilized a higher number of 72 lagged predictors and used parameters like max_depth=12 and eta=0.1 for gradient boosting. Both ensemble methods relied on a high number of historical data points to implicitly capture seasonality and non-linear dependencies for their recursive forecasts.

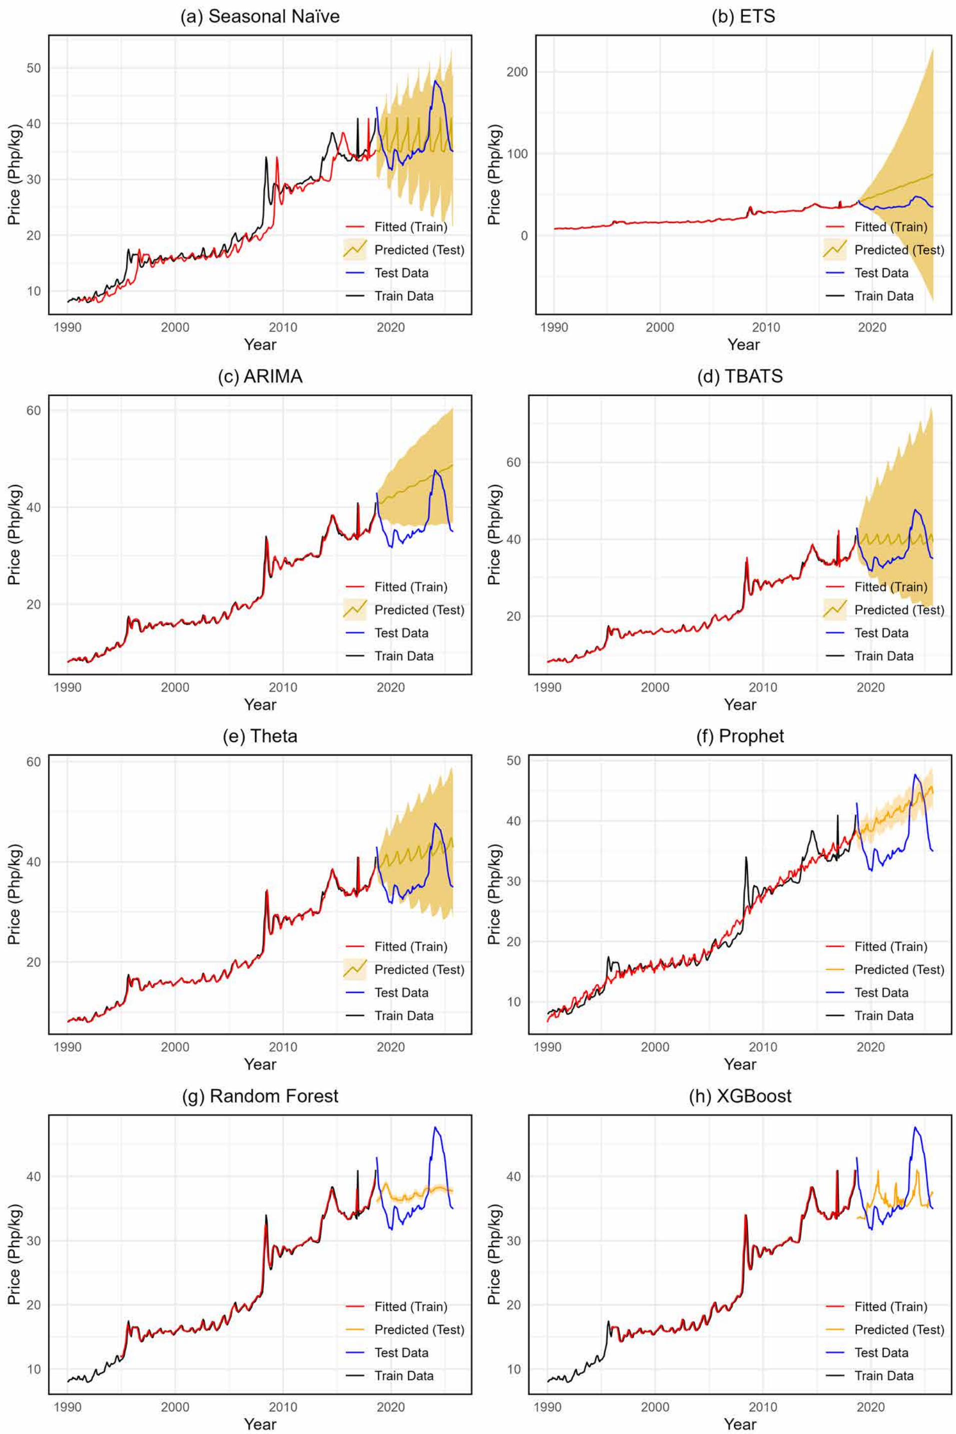

Figure 3. *Visual Comparison of Model Fit and Forecast Paths of the Different Models*

Figure 3 illustrates how each model's fitted values (red line) tracked the training data (black line) and how their predictions (orange line) projected the series into the test period (blue line), along with the 95% confidence interval (shaded orange region).

Extrapolative models like Theta (e), ARIMA (c), and ETS (b) generated predictions that generally maintained the steep upward trend established by the training data. For instance, the ETS model, though showing a pronounced trend, severely underestimated the magnitude of prices until the end of the test set, leading to a massive error.

Conversely, models based on lagged inputs, such as RF (g) and XGBoost (h), display a relatively flat or moderately increasing trajectory for the prediction path. This behavior is characteristic of recursive ML models, which struggle to extrapolate a steep trend that lies significantly outside the range of their training features. This visual projection contrasts sharply with the steep upward path of the actual prices in the test period.

Furthermore, the confidence intervals (CI, shaded orange region) for the traditional models Seasonal Naïve (a), ETS (b), ARIMA (c), TBATS(d), and Theta (e) were visibly wider, reflecting a greater range of uncertainty. This outcome contrasts sharply with the tighter CI displayed by the ML models, RF (g) and XGB (h). This visual comparison of differing trend extrapolation strategies and uncertainty across model classes establishes the context for the numerical accuracy comparison in the subsequent section.

3.3 Model Evaluation and Diagnostic

Table 2 shows the predictive accuracy of the different models assessed using MAE, RMSE, and MAPE on the test period September 2018 to October 2025. It shows that RF has the lowest MAE at 3.60, indicating that it produced the smallest average deviation from actual prices. RF also has the lowest RMSE at 4.30, suggesting strong performance even when larger errors were penalized. In terms of MAPE, RF again registered the lowest at 9.36%, which implies the highest relative accuracy.

Table 2. Comparison of Predictive Accuracy

Model	MAE	RMSE	MAPE
1. Seasonal Naïve	3.88	4.98	9.87
2. ETS	20.22	21.83	54.75
3. ARIMA	7.48	8.34	21.27
4. TBATS	4.94	5.28	13.52
5. Theta	5.45	5.94	15.33
6. Prophet	5.13	5.68	14.43
7. RF	3.60	4.30	9.36
8. XGBoost	3.68	4.63	9.37

The next best performers were the XGBoost model (second lowest across all accuracy metrics: MAE=3.68, RMSE=4.63, MAPE=9.37%) and the Seasonal Naïve model (third lowest across all accuracy metrics: MAE=3.88, RMSE=4.98, and MAPE=9.87%). The high accuracy of RF and the robust performance of the Seasonal Naïve model confirm the necessity of focusing on strong annual periodicity and non-linear patterns when forecasting volatile commodity prices such as the wholesale price of regular-milled rice. This result is consistent with studies emphasizing the superior performance of tree-based models in capturing localized non-linear dependencies over purely linear projections (Montgomery, 2018).

Conversely, the ETS model recorded the highest error values (MAE=20.22, RMSE=21.83, MAPE=54.75%). This performance deficiency suggests that the structural sensitivity of the ETS(M,A,A) model specification was unable to adapt to the abrupt trend shifts and extreme volatility observed in the historical data. This finding supports the literature on the fragility of exponential smoothing decomposition models when facing significant structural breaks in economic time series (Barrow, 2020).

Table 3 summarizes the residual diagnostics of the models assuming a significance level α=0.05. The Ljung–Box test revealed that ARIMA, TBATS, Theta, and RF produced residuals without significant autocorrelation ($p>0.05$), indicating that they sufficiently captured the time dependence in the data. This means these models successfully extracted all predictable linear structure. In contrast, the Seasonal Naïve, ETS, Prophet, and XGBoost models showed significant autocorrelation, suggesting they failed to fully fit the seasonal or trend components. This challenges the structural completeness of these models for this dataset.

Table 3. Residual Diagnostic Test Results

Model	Ljung-Box p-value	Lilliefors p-value	Residuals Interpretation
1. Seasonal Naïve	0.00	1.93×10^{-32}	Autocorrelated and not normally distributed
2. ETS	0.00	1.83×10^{-28}	Autocorrelated and not normally distributed
3. ARIMA	0.80	3.53×10^{-41}	Not autocorrelated and not normally distributed
4. TBATS	0.17	1.54×10^{-24}	Not autocorrelated and not normally distributed
5. Theta	0.08	4.31×10^{-43}	Not autocorrelated and not normally distributed
6. Prophet	0.00	3.30×10^{-16}	Autocorrelated and not normally distributed
7. RF	0.12	1.24×10^{-28}	Not autocorrelated and not normally distributed
8. XGBoost	0.00	0.06	Autocorrelated and normally distributed

Normality testing using the Lilliefors test showed that the residuals for nearly all models were not normally distributed ($p<0.05$). The sole exception was the XGBoost model, which exhibited approximately normal residuals ($p=0.06$), suggesting a distribution of errors that adhered to theoretical statistical assumptions. This contradicts the non-normality found in the majority of time series models tested and is often indicative of effective error regularization within the boosting framework. For the models that failed the normality test, this deviation is likely due to the presence of large, non-linear price shocks such as the 2008 spike that skew the error distribution. This is consistent with literature on the non-normality of residuals following major economic crises or periods of high volatility in commodity markets (Corsi, 2008).

The overall results highlight a practical trade-off between predictive accuracy and diagnostic adequacy. The RF model provided the best predictive accuracy, suggesting its non-linear approach successfully minimized final error. Conversely, models like ARIMA demonstrated structural adequacy but yielded higher errors. Similarly, the XGBoost model, despite achieving near-best accuracy, was not structurally adequate as its residuals showed significant autocorrelation. This situation is often described in forecasting literature as the bias-variance trade-off, where the flexible structure of ML models allows for a lower variance (better accuracy) at the cost of higher diagnostic bias (Yu & Eng, 2020).

3.4 Best Fit Model and its Forecast

The RF model was identified as the best-fit model for forecasting the wholesale price of regular-milled rice in the Philippines, having demonstrated the lowest error across all three primary accuracy metrics in the out-of-sample test. Following its validation, the RF model was re-estimated using the complete dataset (January 1990 – October 2025) to maximize its predictive capacity for the future period. Table 4 presents the point forecast for November 2025 to December 2027, including the 95% confidence interval, which is also visually depicted in Figure 4.

Table 4. RF Model Forecast of Regular-Milled Rice Price with 95% Confidence Interval

Date	Point Forecast	Lower 95%	Upper 95%
Nov-2025	34.62	32.54	36.69
Dec-2025	34.54	31.27	37.80
Jan-2026	34.45	30.60	38.31
Feb-2026	34.26	30.26	38.26
Mar-2026	34.16	30.16	38.16
Apr-2026	34.08	30.09	38.06
May-2026	34.03	30.13	37.93
Jun-2026	34.10	30.02	38.18
Jul-2026	34.15	30.27	38.04
Aug-2026	34.26	30.58	37.94
Sep-2026	34.31	30.76	37.86
Oct-2026	34.36	30.49	38.22
Nov-2026	34.39	30.54	38.24
Dec-2026	34.45	30.61	38.30
Jan-2027	34.47	30.58	38.37
Feb-2027	34.55	30.80	38.31
Mar-2027	34.59	30.97	38.20
Apr-2027	34.62	30.95	38.29
May-2027	34.70	30.91	38.48
Jun-2027	34.90	30.93	38.87
Jul-2027	35.08	31.14	39.03
Aug-2027	35.10	31.55	38.65
Sep-2027	35.04	32.13	37.94
Oct-2027	34.91	32.82	37.00
Nov-2027	34.97	33.01	36.93
Dec-2027	35.02	33.15	36.89

The forecast projects that the wholesale price of regular-milled rice will generally remain within the Php 34.00/kg to Php 35.50/kg range through the end of 2027. The forecasts show a characteristic dampening trend, where the prices slightly decrease from the starting point of Php 34.62/kg in November 2025 to a low of Php 34.03/kg in May 2026, before gradually climbing again. The width of the 95% confidence interval (e.g., from Php 30.26 to Php 38.26 in February 2026) reflects the high degree of price volatility present in the historical series.

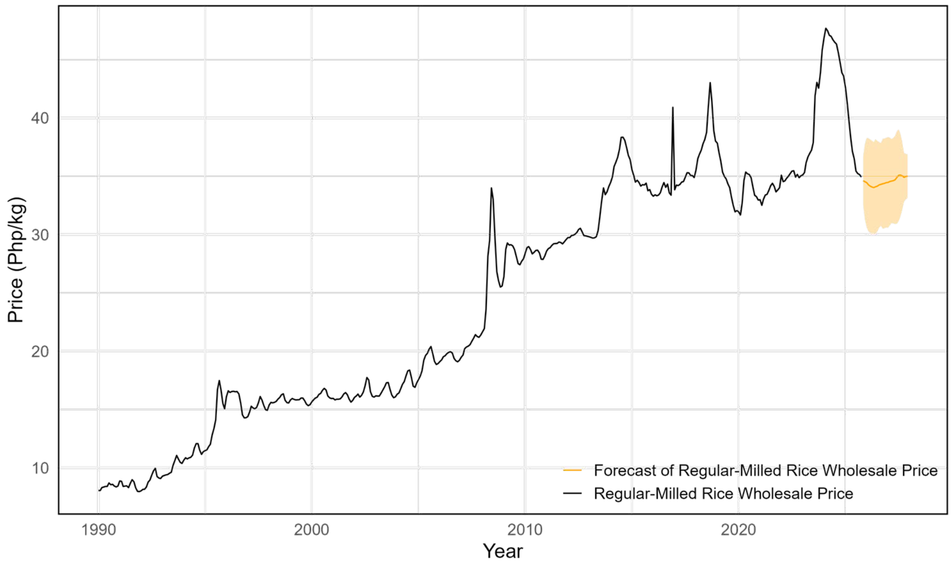

Figure 4. *RF Model Forecast for Nov 2025 to Dec 2027 with 95% Confidence Interval (shaded orange region)*

The general price stability predicted for the next two years suggests that broad market interventions, like implementing price caps, may not be necessary unless external shocks occur. Policy focus may remain on supply-side improvements rather than demand-side controls, as the wholesale price trend is generally contained within the established range. Moreover, the results highlight the importance of selecting forecasting models capable of capturing non-linear dependencies and localized patterns such as seasonal fluctuations, nonlinear trends, and structural changes when modeling the price of rice. The strong performance of the RF model suggests that its ensemble tree structure may be particularly effective for practical forecasting applications in rice price monitoring. Improved forecast accuracy can support policy decisions related to buffer stock management, importation planning, inflation monitoring, and market stabilization efforts. Private sector stakeholders may also benefit from these forecasts for procurement scheduling and risk management.

3.5 Limitations

This study, while providing a rigorous comparative analysis, is subject to several limitations that inform the direction of future research. The analysis relied exclusively on secondary data of nominal prices and a strictly univariate modeling approach, which excludes external variables such as global prices, production volume, weather factors, and other relevant variables. Furthermore, the study used nationally aggregated data, which may mask important regional differences in price dynamics. Future research may build upon this study by exploring multivariate forecasting models that incorporate relevant socioeconomic indicators and specifically investigate hybrid or ensemble frameworks.

4. Conclusions

This study confirmed that the Philippine wholesale price of rice is governed by two dominant forces: a strong, non-stationary trend and a predictable annual seasonal cycle. Models employing non-linear local pattern recognition offered superior predictive accuracy over traditional linear and structural approaches. The RF model was identified as the most effective algorithm, demonstrating the capability to minimize errors despite highly volatile commodity markets characterized by structural breaks. The final projection anticipates that wholesale prices will stabilize within the current range through 2027.

Funding: This research received no external funding.
Conflicts of Interest: The author declares no conflict of interest.
ORCID iD: https://orcid.org/0009-0005-9665-5742

References

[1] Barrow, D., Kourentzes, N., Sandberg, R., & Niklewski, J. (2020). Automatic robust estimation for exponential smoothing: Perspectives from statistics and machine learning. *Expert Systems with Applications, 160*, 113637. https://doi.org/10.1016/j.eswa.2020.113637

[2] Breiman, L. (2001). Random forests. *Machine Learning, 45*, 5–32. https://doi.org/10.1023/A:1010933404324

[3] Caquilala, C. M., Gasmen, P. E., Gonzaga, A. C., Villoria, J. G. P., Añonuevo, L. E., & Bagnol, J. R. V. (2025). Optimizing rice production forecasts in a Philippine province: A comparative analysis of time series models. *AIP Conference Proceedings, 3446*(1). https://doi.org/10.1063/5.0309236

[4] Cauba, A. G., Darvishzadeh, R., Schlund, M., Nelson, A., & Laborte, A. (2025). Estimation of transplanting and harvest dates of rice crops in the Philippines using Sentinel-1 data. *Remote Sensing Applications: Society and Environment, 37*, 101435. https://doi.org/10.1016/j.rsase.2024.101435

[5] Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. https://doi.org/10.1145/2939672.2939785

[6] Corsi, F., Mittnik, S., Pigorsch, C., & Pigorsch, U. (2008). The Volatility of Realized Volatility. *Econometric Reviews, 27*(1–3), 46–78. https://doi.org/10.1080/07474930701853616

[7] De Livera, A. M., Hyndman, R. J., & Snyder, R. D. (2011). Forecasting time series with complex seasonal patterns using exponential smoothing. *Journal of the American Statistical Association, 106*(496), 1513–1527. https://doi.org/10.1198/jasa.2011.tm09771

[8] Hyndman, R. J., & Athanasopoulos, G. (2021). *Forecasting: principles and practice (3rd ed.)*. OTexts: Melbourne, Australia. Retrieved from https://otexts.com/fpp3/

[9] Hyndman, R. J., & Billah, B. (2001). Unmasking the Theta method. *International Journal of Forecasting, 20*, 287–295. https://doi.org/10.1016/S0169-2070(01)00143-1

[10] Intal, P. & Garcia, M. (2005). Rice and Philippine Politics. *Philippine Institute for Development Studies*. https://doi.org/10.62986/dp2005.13

[11] Li, J., Lin, B., Wang, P., Chen, Y., Zeng, X., Liu, X., & Chen, R. (2024). A hierarchical RF-XGBoost model for short-cycle agricultural product sales forecasting. *Foods, 13*(18), 2936. https://doi.org/10.3390/foods13182936

[12] Gutierrez, M. A., Paguirigan, N. M., Raviz, J., Mabalay, M. R., Alosnos, E., Villano, L., Asilo, S., Arocena Jr., A., Maloom, J., & Laborte, A. (2019). The Rice Planting Window in the Philippines: An Analysis Using Multi-Temporal SAR Imagery. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences, XLII-4/W19*, 241–248. https://doi.org/10.5194/isprs-archives-XLII-4-W19-241-2019

[13] Montano, V.E. (2025). Is There a High Rice Price Volatility in the Philippines: Insights from Bayesian Seasonal Autoregressive Distributed Lag (BARDL). *Journal of Economics, Finance And Management Studies*. https://doi.org/10.47191/jefms%2Fv8-i1-40

[14] Montgomery, J.M., & Olivella, S. (2018). Tree-Based Models for Political Science Data. *American Journal of Political Science*. https://doi.org/10.1111/AJPS.12361

[15] Philippine Rice Research Institute (2025). *Rice Prices in the Philippines*. Ricelytics. Retrieved November 21, 2025, from https://ricelytics.philrice.gov.ph/rice_prices/

[16] Philippine Statistics Authority (n.d.). *Metadata on Prices*. Philippine Statistics Authority OpenSTAT. Retrieved November 21, 2025 from https://openstat.psa.gov.ph/Metadata/2M4AWB01

[17] Rubillos, L. E., Mangubat, A. F., Seno, M. A., Caragay, J. C., Salazar, N., Caladcad, J. A., & Villeta, R. (2024). Time series analysis of wheat moisture content variations for grain storage systems. *Journal of Stored Products Research, 108*, 102395. https://doi.org/10.1016/j.jspr.2024.102395

[18] Taylor, S. J., & Letham, B. (2018). Forecasting at scale. *The American Statistician, 72*(1), 37–45. https://doi.org/10.1080/00031305.2017.1380080

[19] Yu, A. C., & Eng, J. (2020). One Algorithm May Not Fit All: How Selection Bias Affects Machine Learning Performance. *RadioGraphics, 40*(7), 1932–1937. https://doi.org/10.1148/rg.2020200040