Short Run Forecasting of Petroleum Prices in Tanzania Using ARIMA and Exponential Smoothing: A Comparative Analysis

Sharifu Nurdin[1] and Bahati Ilembo[2]

[1]Sharifujabiry307@gmail.com
[2]bmilembo@mzumbe.ac.tz
[1,2]Department of Mathematics and Statistics, Mzumbe University, P.O. Box 1, Morogoro, Tanzania

Original Research Article

https://doi.org/10.51867/asarev.2.1.7

ABSTRACT

Tanzania's economy remains highly dependent on imported petroleum, with petrol price fluctuations significantly impacting inflation, transportation costs, and household welfare. The purpose of the study was to conduct a comparative analysis of ARIMA and Exponential Smoothing methods for short-run forecasting of petrol prices in Tanzania to identify the most accurate model for predicting future petroleum price trends. Using monthly petrol price data from 2005 to 2024 obtained from the Bank of Tanzania (BOT), we applied the Box-Jenkins methodology to compare the Autoregressive Integrated Moving Average (ARIMA) and Exponential Smoothing models. The ARIMA (1,1,4) with drift model was identified as optimal based on Akaike's Information Criterion (AIC), Bayesian Information Criterion (BIC), and maximum log-likelihood criteria. This model achieved superior forecasting accuracy with a Mean Absolute Percentage Error (MAPE) of 2.44%, compared to 2.58% for Exponential Smoothing. The model forecasts a steady monthly increase of 0.3-0.5% in petrol prices, projecting prices to reach 3,500.65 TZS/L by September 2026. While the model demonstrates strong predictive performance (Ljung-Box p-value = 0.265), its limitations in anticipating sudden price shocks highlight the need for complementary risk management strategies. These findings provide policymakers and market participants with a reliable tool for budget planning and economic forecasting. The study underscores Tanzania's vulnerability to global oil market volatility and emphasizes the importance of developing strategic fuel reserves and alternative energy sources to enhance energy security.

Keywords: Akaike Information Criteria, ARIMA, Bayesian, Ljung–Box

I. INTRODUCTION

Petrol prices play a pivotal role in shaping global economic trends, significantly influencing inflation rates, transportation costs, and energy security. In 2024, the average price of petrol declined to approximately 3.3 USD per gallon, a decrease from the previous year's average of 4.33 USD per gallon (Smith & Johnson, 2024). This drop in global petrol prices was primarily attributed to reduced crude oil prices and narrower refinery margins, driven by shifts in global supply and demand dynamics (Brown et al., 2024). Research by Kilian (2022) highlights that such fluctuations in oil prices contribute to inflationary pressures, particularly in economies where transportation and food prices are highly sensitive to energy costs. For instance, lower petrol prices can temporarily ease inflation but may also reflect broader economic slowdowns, which can have mixed effects on consumer spending and economic growth (Hamilton, 2019).

The relationship between petrol prices and inflation is well-documented in economic literature. Studies by Baumeister and Hamilton (2019) demonstrate that oil price shocks account for a significant portion of inflation variability in both advanced and emerging economies. Similarly, transportation costs, which are directly tied to petrol prices, have a cascading effect on the prices of goods and services, particularly in sectors that rely heavily on logistics and supply chains (Cologni& Manera, 2008). Furthermore, energy security concerns are exacerbated by price volatility, as nations dependent on oil imports face increased economic vulnerability during periods of rising prices (Sadorsky, 2012).

The decline in petrol prices in 2024, while providing temporary relief to consumers, underscores the need for long-term strategies to mitigate the economic risks associated with oil price fluctuations. Policies aimed at diversifying energy sources, improving energy efficiency, and investing in renewable energy infrastructure are critical to reducing dependence on fossil fuels and enhancing economic resilience (Brown et al., 2024; Kilian, 2022).

In Sub-Saharan Africa, rising petrol prices have significantly contributed to inflation, directly impacting daily living costs across the region. According to regional studies, the price of petrol in Sub-Saharan Africa increased by an average of 10% in 2024, a slower rate compared to the 15% rise observed in 2022 (Oluwaseun & Adebayo, 2024). Despite this moderation, the increase remains a significant burden for households, particularly in oil-importing countries where petrol prices are closely tied to global market fluctuations and local economic conditions (Munyeka, 2023). For example, in Nigeria, petrol prices rose by 15% in 2024, leading to sharp increases in transportation and food costs, which disproportionately affect low-income households (Eze & Onyema, 2024). Similarly, Kenya and South Africa reported petrol price increases of 10% and 12%, respectively, further exacerbating inflationary pressures in these economies (Kiprop & van der Merwe, 2024).

The relationship between petrol prices and inflation in Sub-Saharan Africa is well-documented in academic research. Gelan (2018) finds that petrol price increases have a direct and immediate impact on inflation, as higher transportation costs are passed on to consumers through increased prices for goods and services. This is particularly evident in urban areas, where transportation is a critical component of food distribution and other supply chains. Similarly, Baffes et al. (2015) highlighted that petrol price shocks disproportionately affect oil-importing countries in the region, as they cannot absorb global price fluctuations, leading to higher living costs and reduced household purchasing power. The economic implications of rising petrol prices extend beyond inflation. Studies by Nkengfack and Fotio (2021) showed that higher petrol prices reduce disposable incomes, particularly for low-income households, and can exacerbate poverty and inequality. In response, governments in the region face the dual challenge of mitigating the immediate impact of rising petrol prices while addressing long-term structural issues, such as over-reliance on fossil fuels and inadequate public transportation systems (Munyeka, 2023). Policies aimed at reducing dependence on petrol, such as investing in renewable energy and improving public transportation infrastructure, are critical to building economic resilience and reducing vulnerability to global price fluctuations (Oluwaseun & Adebayo, 2024).

In Tanzania, petrol prices have seen notable fluctuations, closely tied to global oil price trends and the country's reliance on fuel imports. According to the Energy and Water Utilities Regulatory Authority (EWURA, 2024), petrol prices in Tanzania peaked at TZS 2,600 per liter in 2024, reflecting global price fluctuations and the depreciation of the Tanzanian shilling. The Bank of Tanzania (BOT, 2024) reported that fuel imports make up a significant portion of Tanzania's foreign exchange expenditure, with petrol accounting for approximately 25% of total imports in 2024. The price of petrol in Tanzania has increased by around 8-10% annually since 2020, directly affecting transportation costs, food prices, and overall inflation. Research by Mwamunyange (2023) highlights that rising petrol prices disproportionately affect low-income households, as transportation and food costs consume a larger share of their income.

Kato and Mwakatobe (2022) emphasize that inefficiencies in fuel distribution networks contribute to regional price disparities, which further complicating efforts to stabilize petrol prices. Government interventions, such as subsidies and fuel price caps, have been implemented to cushion the impact of these price increases. In 2024, the Tanzanian government allocated approximately TZS 600 billion to subsidize petrol prices to prevent further economic strain on consumers (EWURA, 2024). However, the financial burden of these subsidies has strained the government's budget, and the long-term sustainability of such interventions remains uncertain.

Current literature on petrol price forecasting in Tanzania shows limited recent studies that provide a comparative analysis of ARIMA and Exponential Smoothing models for short-run predictions (1-6 months). While both methods are established in time series forecasting, there is insufficient empirical evidence on their relative performance in Tanzania's petroleum market, particularly under its unique pricing policies and market conditions. Existing research primarily focuses on single-model approaches (either ARIMA or Exponential Smoothing) without direct comparison, long-term forecasting rather than short-term volatility and global or regional markets, with limited Tanzania-specific studies. This gap leaves policymakers, oil industry stakeholders, and economists without clear guidance on which model performs better for short-term petrol price projections in Tanzania. A systematic comparison using error metrics (MAE, RMSE, and MAPE) could enhance forecasting accuracy and support data-driven decision-making in fuel pricing and economic planning.

II. OBJECTIVES

The general objective of the study was to conduct a comparative analysis of ARIMA and Exponential Smoothing methods for short-run forecasting of petrol prices in Tanzania in order to identify the most accurate model for predicting future price trends. Precisely, the study was guided by the following specific objectives;

i. To decompose the historical petrol price time series in Tanzania to identify its underlying components.
ii. To apply the ARIMA model to the historical petrol price data and evaluate its suitability for short-run forecasting in Tanzania.
iii. To apply the Exponential Smoothing model to the historical petrol price data and assess its accuracy in short-run forecasting for Tanzania.

iv. To compare the forecasting accuracy of the ARIMA and Exponential Smoothing models using model selection criteria.
v. To forecast the short-run trends of petrol prices in Tanzania using the best-performing model.

III. MATERIALS AND METHODS

3.1 Data and Data Source

The study used secondary univariate time series data obtained from the official database of the Bank of Tanzania (BOT). The petrol price data for the period from 2005 to 2024 were obtained through (https://www.bot.go.tz/Statistics/externalstatistics?code=PAI2&TypeOption=Prices&variableOption=Petrol%20price).The monthly data providing a total of approximately 237 observations. The data included monthly petrol price records drawn from official reports on economic and energy statistics. This dataset is considered suitable for analyzing trends in petrol prices over time and for forecasting future price trend.

3.2 Autocorrelation and Partial Autocorrelation Functions

i. Autocorrelation Function (ACF)

The Autocorrelation Function (ACF) measures how a time series correlates with its past values at different time lags, helping identify trends, seasonality, and moving average (MA) effects. If ACF values decay slowly, the series is likely to be non-stationary and may need differencing. For an MA(q) model, the ACF drops sharply after lag *q*, indicating the order of the MA component. Meanwhile, the Partial Autocorrelation Function (PACF) isolates the direct correlation between a series and its lagged values, removing indirect effects from intermediate lags. This helps identify an autoregressive (AR) process if PACF cuts off after lag *p*, it suggests an AR(p) model. Altogether, ACF and PACF guide ARIMA modeling by distinguishing between AR and MA structures and determining appropriate lag orders, and it is given as;

$$\rho_i = corr(y_t\,, i) = \frac{cov(y_t, y_{t-i})}{\sqrt{var(y_t).var(y_{t-1})}} = \frac{\gamma_i}{\gamma_0} \ldots\ldots\ldots\ldots\ldots\ldots (1)$$

Where, i =1, 2, 3⋯

ii. Partial Autocorrelation Function (PACF)

The Partial Autocorrelation Function (PACF) isolates the direct relationship between a time series observation and its specific lagged value while eliminating the influence of all intermediate lags. This can be thought like measuring the pure connection between today's stock price and its price five days ago, ignoring what happened on days 1 through 4. Unlike the regular Autocorrelation Function (ACF) that shows cumulative correlation patterns, PACF precisely identifies how many previous periods directly impact the current value, making it invaluable for determining the order (p) in AR(p) models: if PACF shows significant spikes at lags 1 through p but then abruptly drops to zero, it indicates an autoregressive process of order p. This sharp cutoff property allows analysts to distinguish between genuine direct dependencies and indirect correlations, providing critical insights for building accurate ARIMA forecasting models by revealing exactly which historical data points truly matter.

Partial Autocorrelation function at lag i for time series is given as:

$$\emptyset_{11} = corr(Y_{t+1}, Y_t) = \rho_1 \ldots\ldots\ldots\ldots\ldots\ldots (2)$$

$$\emptyset_{ii} = Corr\left(Y_{t+1} - \hat{Y}_{t+i}, \quad Y_t - \hat{Y}_t\right), i \geq 2 \ldots\ldots\ldots\ldots\ldots\ldots (3)$$

3.3 Stationarity Testing

A core requirement of ARIMA modeling is that the time series must be stationary, meaning its statistical properties (mean, variance, and autocorrelation) remain constant over time. Non-stationary data can lead to unreliable forecasts, so testing for stationarity is crucial. We performed the Augmented Dickey Fuller (ADF) test as well as Phillips-Perron (PP) Test.

i. Augmented Dickey-Fuller (ADF) Test

The Augmented Dickey-Fuller (ADF) test is a formal statistical hypothesis test designed to assess whether a given time series contains a unit root, which would indicate non-stationarity. As an enhanced version of the original Dickey-Fuller test, it addresses more complex time series structures by incorporating lagged difference terms in its regression equation to account for potential autocorrelation. The test estimates a regression model where the dependent variable is differenced against its lagged values, potentially including deterministic components like a constant term or linear trend depending on the series characteristics. The key output is a test statistic that is compared against critical values derived from the Dickey-Fuller distribution if this statistic is more negative than the critical value at a chosen significance level (typically 1%, 5%, or 10%), the null hypothesis of a unit root is rejected in favor of stationarity. The test also provides a p-value interpretation, where values below the significance threshold (usually 0.05) suggest sufficient evidence to conclude the series is stationary. Importantly, the ADF test accommodates various data behaviors through different model specifications (none, drift, or trend) and automatically determines the optimal number of lagged differences to include, using information criteria. This makes it particularly valuable for real-world time series analysis where the underlying data-generating process is unknown, though its power decreases when dealing with near-unit-root processes or structural breaks in the series. The methodology of the ADF test applies this statistical procedure by strategically transforming the underlying regression model, represented by the equation:

$$\Delta y_t = \alpha + \beta_t + \gamma y_{t-1} + \epsilon_t \quad (4)$$

Where;

Δyt: First difference of the time series at time t

α (Drift term): Represents a constant intercept, capturing deterministic trends

βt Time trend term): Accounts for linear time-dependent

γy_{t-1} (Key test term): Coefficient of the lagged level of the series.

ϵ_t is the error term

The hypothesis will be:

Null hypothesis (H_0): The time series has a unit root (non-stationary).

$$H_0: \emptyset = 1$$

Alternative hypothesis (H_1): The time series is stationary.

$$H_1: \emptyset \neq 1$$

If the test statistic is smaller than the critical value, H_0 is rejected, confirming stationarity.

ii. Phillips-Perron (PP) Test

The Phillips-Perron (PP) test is a unit root test that, like the ADF test, checks for stationarity in time series data but with a key difference: instead of adding lagged difference terms to correct for serial correlation (as ADF does), the PP test adjusts the test statistic itself to account for autocorrelation and heteroskedasticity using Newey-West standard errors. This makes it more robust to unspecified serial correlation and time-varying volatility in the errors. The test follows the same null hypothesis (presence of a unit root/non-stationarity) and alternative (stationarity) as ADF, but its modified approach often makes it preferable when the autocorrelation structure is unknown or complex. However, its performance can be sensitive to the choice of bandwidth in the variance correction, and it may be less reliable than ADF for some finite samples. The PP test is particularly useful when you suspect your data has autocorrelation patterns that aren't well-captured by simple lag structures. It is represented by the equation (5) underneath.

$$\Delta Y_t = \alpha + \beta_t + y y_{t-1} + \epsilon_t \quad (5)$$

Where:

Δy_t	is the different series (y_t - 1)
α	represents a constant (drift) term
β_t	captures a deterministic time trend (optional)
γ	is the coefficient of interest (testing γ=0 implies a unit root)
ϵ_t	is the error term.

3.3.1 Differencing

Differencing is a fundamental transformation technique in time series analysis that converts non-stationary data into stationary form by computing the difference between consecutive observations (or seasonal lags). This process effectively eliminates trend components and seasonality by removing time-dependent changes in the mean, while preserving the underlying stochastic patterns. First-order differencing ($Y_t'=Y_t - Y_{t-1}$) addresses linear trends, while seasonal differencing ($\Delta_s Y_t' = Y_t - Y_{t-s}$) handles periodic fluctuations. Higher-order differencing may be applied for more complex trends. By stabilizing the mean and reducing variance, differencing enables the reliable application of stationary-dependent models like ARIMA, where the optimal differencing order (d) is typically determined through unit root tests (ADF, PP) or by analyzing autocorrelation decay. If the time series is non-stationary, differencing is applied to stabilize the mean. Equation (6) shows the first-order differencing formula and if stationarity is not achieved, second-order differencing is performed as indicated in equation (7).

$$Y_t' = Y_t - Y_{t-1} \quad (6)$$

:

$$Y_t'' = Y_t' - Y_{t-1}' \quad (7)$$

Further, the appropriate number of differencing steps (d) is determined through stationarity tests and visual inspections of ACF and PACF plots.

3.4 Stationary Linear Time Series Modeling Approaches

3.4.1 Moving Average (MA) Process

This model represents a time series where current values are explained by a linear combination of past error terms. Let $\{u_t\}$ (t = 1,2,3) denote a white noise process comprising independent and identically distributed (iid) random variables with $E(u_t) = 0$ and $Var(u_t) = \sigma^2$. The q-th order moving average process, denoted MA(q), is formally expressed as;

$$y_t = \mu + u_t + \theta_1 u_{t-1} + \theta_2 u_{t-2} + \cdots + \theta q u_{t-q} \ldots\ldots\ldots\ldots (8)$$

This model formulates the current observation y_t as a weighted sum of the current and past q error terms. The coefficients θ_j $(j = 1, q)$ are estimated to quantify the impact of these historical shocks. Crucially, the model exhibits finite memory, only the most recent q errors (u_t to u_{t-q}) influence y_t, while older errors (u_{1-q-1}, u_{1-q-2} etc.) have exactly zero effect. This finite dependence structure gives the MA(q) process its characteristic short-memory property, distinguishing it from infinite-memory models like AR processes. The memory length is precisely controlled by the model of order q, making MA specifications particularly suitable for modeling transient shock effects in time series data.

3.4.2 Autoregressive (AR) Process

An Autoregressive (AR) model predicts the current value of a time series by using a linear combination of its past values plus some random noise. The general form of an AR model of order p (AR(p)) is:

$$y_t = c + \alpha^1 y_t^{-1} + \alpha^2 y_t^{-2} + \ldots + \alpha_p y_t^{-}{}_p + u_t \ldots\ldots\ldots\ldots\ldots\ldots\ldots\ldots\ldots\ldots\ldots\ldots\ldots\ldots (9)$$

Where:

y_t = value at time t (what we're predicting)

c = constant term (like an intercept in regression)

α_1 to α_p = coefficients for each past value (how much weight each lag gets)

y_{t-1} to y_{t-p} = past values (from 1 to p periods back)

u_t = random error term (white noise with mean 0 and constant variance)

The Partial Autocorrelation Function (PACF) is essential for determining the order (p) of an autoregressive (AR) model because it isolates the direct correlation between an observation y_t and its lagged value y_{t-1} while controlling for all intermediate lags (y_{t-1} to y_{t-p}). For an AR(p) process, the PACF exhibits statistically significant spikes at lags 1 through p, representing the direct influence of each historical observation, before abruptly dropping to insignificant levels for lags beyond p. This distinct cutoff point, where the PACF values exit the confidence bounds, precisely indicates the model's order by revealing how many past values directly affect the current observation. Unlike the autocorrelation function (ACF) that shows compounded dependencies, the PACF's ability to measure pure, lag-specific relationships makes it uniquely valuable for AR model identification, as it clearly distinguishes between direct autoregressive effects and indirect correlations propagated through intermediate time steps.

3.4.3 Autoregressive Integrated Moving Average (ARIMA) Model

The ARIMA model is a widely used statistical method for time series forecasting. It combines three key components: Autoregression (AR), Differencing (I), and Moving Average (MA). ARIMA models are generally denoted as ARIMA (p, d, q), where:

p represents the number of autoregressive terms,

d is the number of differencing required to make the series stationary,

q denotes the number of moving average terms

The general ARIMA model equation is given as;

$$Y_t = \sum_{i=1}^{p} \emptyset_i Y_{t-i} + \sum_{j=1}^{q} \theta_j \varepsilon_{t-j} + \varepsilon_t \ldots\ldots\ldots\ldots\ldots\ldots\ldots\ldots\ldots\ldots\ldots\ldots\ldots\ldots (10)$$

Where ϵ_t represents white noise

Where:

- Y_t: value of the time series at time t
- $\emptyset_i$: AR coefficients (lags of Y)
- θ_j: MA coefficients (lags of residuals ε)
- ε_t: white noise error at time t

3.5 Box-Jenkins's methodology

The Box-Jenkins approach, introduced by George Box and Gwily Jenkins (1970), is a widely used methodology for time series forecasting. It focuses on identifying an appropriate model, estimating its parameters, and using the fitted model for forecasting. The strength of the Box-Jenkins approach lies in the ARIMA model, which provides a flexible framework for analyzing time-dependent data. In addition to ARIMA, Exponential Smoothing (ETS models) is another important approach for short-term forecasting. Exponential Smoothing methods, introduced by Charles Holt and Peter Winters, are based on the principle of giving exponentially decreasing weights to past observations. The key advantage of Exponential Smoothing is its ability to capture trends and seasonality while being computationally efficient. In this study, a comparative analysis of ARIMA and Exponential Smoothing methods is conducted to determine the most suitable forecasting technique for petrol prices in Tanzania. Both models are evaluated based on their predictive accuracy and ability to capture underlying patterns in the data.

3.6 Model Identification

The initial step in time series analysis is to plot the data to visually inspect any irregularities, trends, or patterns that may influence model selection (Hyndman & Athanasopoulos, 2018). If the variance appears to increase over time, a transformation such as taking the logarithm or applying a power function may be necessary to stabilize it. Assessing stationarity is crucial since a stationary series ensures consistent statistical properties over time, making forecasting more reliable. This can be determined through graphical analysis or statistical tests like the Augmented Dickey-Fuller (ADF) test. If the series is non-stationary, differencing is applied to remove trends and stabilize the mean. Following this, preliminary values for the Autoregressive (AR) order (p), degree of differencing (d), and Moving Average (MA) order (q), as well as their seasonal components (P, D, Q), are identified. This process involves examining Autocorrelation Function (ACF) and Partial Autocorrelation Function (PACF) plots, which help in selecting an appropriate ARIMA model. Once these parameters are determined, the next stages include model estimation, diagnostic checking, and validation before proceeding with forecasting.

3.6.1 Model Estimation

Once an appropriate ARIMA or Exponential Smoothing (ETS**)** model is selected with specific values for P, d, and Q (for ARIMA) or smoothing parameters (for ETS), the next step is to estimate the model parameters. Maximum Likelihood Estimation (MLE) is used to estimate the coefficients for the chosen model during the identification stage. The best model is chosen based on model selection criteria like the Akaike Information Criterion (AIC) and Bayesian Information Criterion (BIC**)**

i. *Akaike Information Criterion (AIC)*

Kullback *et al*. (1951) developed a measure to capture the information that is lost when approximating reality. Kullback & Leibler measure is a criterion for a good model that minimises the loss of information. Two decades later, Akaike established a relationship between the Kullback & Leibler measure and the maximum likelihood estimation (MLE) method that was used in many statistical analyses for model selection (Akaike, 1974). This criterion, referred to as the Akaike Information Criterion, is generally considered the first model selection criterion that should be used in practice. The AIC is given as: $AIC = -2logL(\Theta) + 2K$ ……………………………………..(11)

where L is the likelihood of the model, K is the number of parameters, and n is the sample size. The model with the lowest AIC/BIC is preferred. The parameter θ is the set of model parameters, L $(\hat{\theta})$ is the likelihood of the candidate model given the data when evaluated at the maximum likelihood estimate of $\theta,$ and k is the number of estimated parameters in the candidate model. Since AIC does not consider the effect of sample size, for small sample sizes, the second-order equation of the Akaike information criterion (AIC_C) is defined as;

$$AICc = -2\log L\left(\hat{\theta}\right) + 2k + \frac{(2k+1)}{(n-k-1)} \quad \ldots\ldots(12)$$

Where

n is the number of observations.

A small sample size is when $n/_k$ Less than 40. As n increases, the third term in AIC_C approaches zero and will therefore, give the same result.

ii. *Bayesian Information Criterion (BIC)*

BIC is another model selection criterion rooted in information theory, but it uses a Bayesian framework. Compared to AIC, BIC imposes a stronger penalty for increasing the number of parameters in the model and is given as;

$$BIC = -2logL(\Theta) + Klog\, n \quad \ldots\ldots(13)$$

Where:

L is the likelihood of the model, K is the number of parameters, and n is the sample size. The model with the lowest AIC/BIC is preferred.

The BIC strongly penalizes the number of involved parameters. High values of AIC means that the observed data does not fit the model, while lower values indicate strong evidence that the observed data fit the models. Similarly, lower values of BIC indicate better fitting of the models.

iii. Diagnostic Checking

After fitting the ARIMA model, it is crucial to check the residuals (the differences between the observed and predicted values) to ensure that the model is adequate. The residuals should resemble white noise, meaning they should have no patterns or serial correlation.

iv. Ljung-Box Test

The Ljung-Box test is used to evaluate whether there are any patterns or correlations in the residuals. The null hypothesis of the test assumes that the residuals are independently distributed with no serial correlation: The Ljung-Box test helps to check whether the errors or residuals in our model have any pattern or correlation. The Ljung-Box testing of hypothesis is defined as;

H_0 = The data are independently distributed (i.e., no autocorrelation up to lag h)

H_1= At least one of the autocorrelations up to lag h is non-zero.

The test statistic given as:

$$Q = n(n+2)\sum_{k=1}^{h}\frac{\hat{\rho}_k^2}{n-k} \quad \ldots (14)$$

Where n is the sample size, $\hat{\rho}_k^2$ is the sample autocorrelation at lag k, and h is the number of lags being tested. Under the null hypothesis, the test statistic Q asymptotically follows a $\mathcal{X}^2_{(h)}$, for the significant levelα, critical region is rejected if;

$$Q > \mathcal{X}^2_{(1-\alpha),h}$$

Where:

$\mathcal{X}^2_{(1-\alpha),h}$ is the $(1-\alpha)$ quantile of Chi-square distribution with h degrees of freedom.

After diagnostic checking, the fitted model is used in forecasting future values if the model is adequate. Otherwise, we need to repeat the selection and estimation method or try with another potential candidate model (Ramasubramanian, 2009).

3.7 Forecasting and Forecasting Accuracy

Once the selected ARIMA model passes the diagnostic checks, it is used for forecasting future values. To evaluate the accuracy of the forecasts, Mean Percentage Error (MPE) and Mean Absolute Percentage Error (MAPE), Mean Absolute Error (MAE) and Root Mean Squared Error (RMSE) were calculated. These accuracy measures are defined as follows:

3.7.1 Mean Percentage Error (MPE)

This is a statistical measure used to evaluate the accuracy of a forecasting or prediction model. It represents the average of the percentage errors between predicted values and actual observed values, taking into account the direction of the errors (over- or under-prediction).

$$MPE = \frac{100}{n}\sum\left(\frac{(p_i - o_j}{o_j}\right) \quad \ldots (15)$$

Where:

p_i is the predicted value for the i^{th} observations, o_i is the observed value for the j^{th} observation and n is the number of non-missing residuals.

3.7.2 Mean Absolute Percentage Error (MAPE)

This is a widely used metric to evaluate the accuracy of forecasting models. Unlike the Mean Percentage Error (MPE), which considers the direction of errors (over- or under-prediction), MAPE measures the average magnitude of errors in absolute terms, making it a robust indicator of prediction accuracy.

$$MAPE = \frac{100}{n}\sum\left|\frac{(p_i - o_j}{o_j}\right| \ldots\ldots\ldots\ldots\ldots\ldots\ldots\ldots\ldots (16)$$

Where:

p_i is the predicted value for the i^{th} observations, o_i is the observed value for the j^{th}observation and n is the number of non-missing residuals.

Mean Absolute Error (MAE) is a fundamental metric used to evaluate the accuracy of a forecasting or prediction model. It measures the average magnitude of errors between predicted and actual values without considering their direction (i.e., over- or under-prediction). Unlike percentage-based metrics (e.g., MAPE), MAE is scale-dependent and expressed in the original units of the data.

$$MAE = \frac{1}{n}\sum_{i=1}^{n}|e_t| \ldots\ldots\ldots\ldots\ldots\ldots\ldots\ldots\ldots (17)$$

Root Mean Squared Error (RMSE) is a widely used metric to measure the accuracy of a predictive model. It calculates the square root of the average squared differences between predicted and actual values. Unlike MAE, RMSE penalizes larger errors more heavily, making it sensitive to outliers.

$$RMSE = \sqrt{\frac{1}{n}\sum_{i=1}^{n}\epsilon_t^2} \ldots\ldots\ldots\ldots\ldots\ldots\ldots\ldots\ldots (18)$$

3.8 Exponential smoothing

In contrast to ARIMA, which treats past observations equally, Exponential Smoothing produces a smoothed time series by applying exponentially decreasing weights to older observations. The technique uses one or more smoothing parameters that need to be estimated, and these parameters determine the weight assigned to each observation (Dimitrov, 2008).

3.8.1 Simple exponential smoothing

Simple Exponential Smoothing used in short-term forecasting is commonly used in a month period. The model assumes the data fluctuated around the static mean value, without a trend or consistent growing pattern. This forecasting method is the most widely used of all forecasting techniques. It requires little computation and is used when the data pattern is approximately horizontal (there is neither cyclic variation nor an obvious trend in the historical data). Simple exponential smoothing (SEM) is the easiest of all exponential smoothing methods. This method is widely used to forecast data which do not indicate a clear trend or a seasonal pattern. This model is represented as:

$$Y_{t+1} = \alpha y_t + \alpha(1-\alpha)y_{t-1} + \alpha(1-a)^2 y_{t-2} + \ldots\ldots\ldots\ldots\ldots (19)$$

Where:

$0 \leq \alpha \leq 1$ is the smoothing parameter. One-step ahead forecast for time $t+1$ is the weighted average of all observations up to time t. Parameter α determines the rate at which eight weights will decline. If α is close to zero, then the fall in weights is sharp, whereas if α is close to 1, then decay in weights is gradual and slow. However, for any value of α, weights assigned to observations fall exponentially when we navigate the series backwards.

3.8.2 Holt's Linear Exponential Smoothing

This method is used when the data shows a trend. Exponential smoothing with the trend is like simple smoothing except both components must be updated in every level-periodic and its trend. The level is a smoothed estimation from the data value at the end of each period. The trend is a smoothed estimation from average growth at the end of each period. Holt's method extends Simple Exponential Smoothing (SES**)** by adding components to handle data with a trend**.** It is ideal when your time series exhibits a trend (upward or downward) but no seasonality**.** This model is represented in equation (20)

$$Lt = \alpha y_t + (1-\alpha)(L_t - 1 + T_t - 1) \ldots\ldots\ldots\ldots\ldots\ldots\ldots (20)$$

In Holt's Linear Exponential Smoothing, Lt represents the current estimated level (baseline value), calculated as a weighted average between the actual observed value at time t(y_t) and the previous one-step-ahead forecast (T_t−1+T_t−1) which combines the prior level and trend. The smoothing weight $\boldsymbol{\alpha}$ controls this balance: if $\alpha = 1$, the model discards past estimates and relies solely on the latest observation (y_t), while $\alpha = 0$ ignores new data, depending entirely on prior forecasts (L_t−1+T_t−1). This ensures adaptive smoothing for trended data.

3.9 Comparative Analysis between ARIMA and Exponential Smoothing Model

When comparing ARIMA and Exponential Smoothing methods for forecasting, a forecasting period of 24 months was utilized. Various performance metrics were applied to assess the effectiveness of the models. These metrics include Mean Absolute Percentage Error (MAPE), Root Mean Squared Error (RMSE), and Bayesian Information Criterion (BIC). MAPE is useful for reporting as it expresses accuracy in percentage terms, making it straightforward to interpret. RMSE, which is minimized during parameter estimation, indicates how well the model fits the data, with larger errors receiving heavier penalties. It also determines the width of the confidence intervals for predictions. . BIC is preferred in model selection because, given enough data, it will favor the true underlying model by penalizing the inclusion of too many parameters. These measures allow us to identify the most effective forecasting model by evaluating how accurately each one predicts future values. By comparing both ARIMA and Exponential Smoothing models using these criteria, we aim to determine which model performs best in forecasting the variable of interest over 24 months. We use the following measures of accuracy to identify the best model

i. MAPE= $\frac{1}{n}\sum_{i=1}^{n}\frac{\epsilon_t}{\alpha_t} * 100$

ii. RMSE= $\sqrt{\frac{1}{n}\sum_{t=1}^{n} et^2}$

iii. BIC $=-2ln\,(L) + ln\,(N)\,k$

In (iii), L is the value of the likelihood function evaluated at the parameter estimates, N is the number of observations, and k is the number of estimated parameters. Minimum values of these accuracy measures provide best-fitting models.

IV. RESULTS AND DISCUSSION

4.1 Trend of Petroleum Price

The data was analyzed using statistical software (R). The figure below presents a time series plot of monthly petrol prices in Tanzania from January 2005 to September 2024, providing a visual representation of price trends over time.

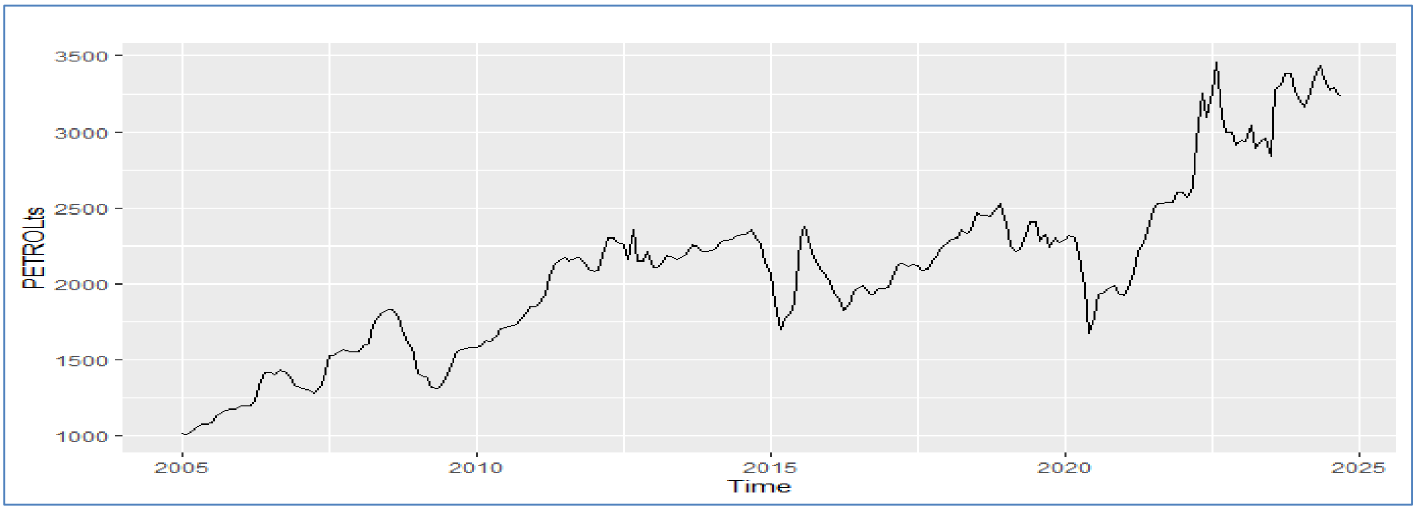

Figure 1: Trend of Petrol Price in Tanzania
Source: Created by authors

Figure 1 shows time series plot of Tanzania's monthly petrol prices (2005–2024), which shows an upward trend indicating an increased price from 1000 to 3500 TZS/L for the respective years. This was due to inflation, currency devaluation, and changes in the world oil price. Notable volatility includes both brief declines (as in 2020 during COVID-19) and significant surges (such as in 2011–2013 and 2022–2023, which are probably related to geopolitical events like the Russia–Ukraine war). Although slight cyclical patterns point to potential seasonality, these effects must be quantified using additional decomposition or statistical tests (such as ADF for stationarity or ARIMA for modeling). The sharp increase beyond 2020 highlights Tanzania's vulnerability to outside shocks and the difficulties in stabilizing fuel prices through policy.

4.2 Analysis of Time Series Data Behavior

The Correlogram

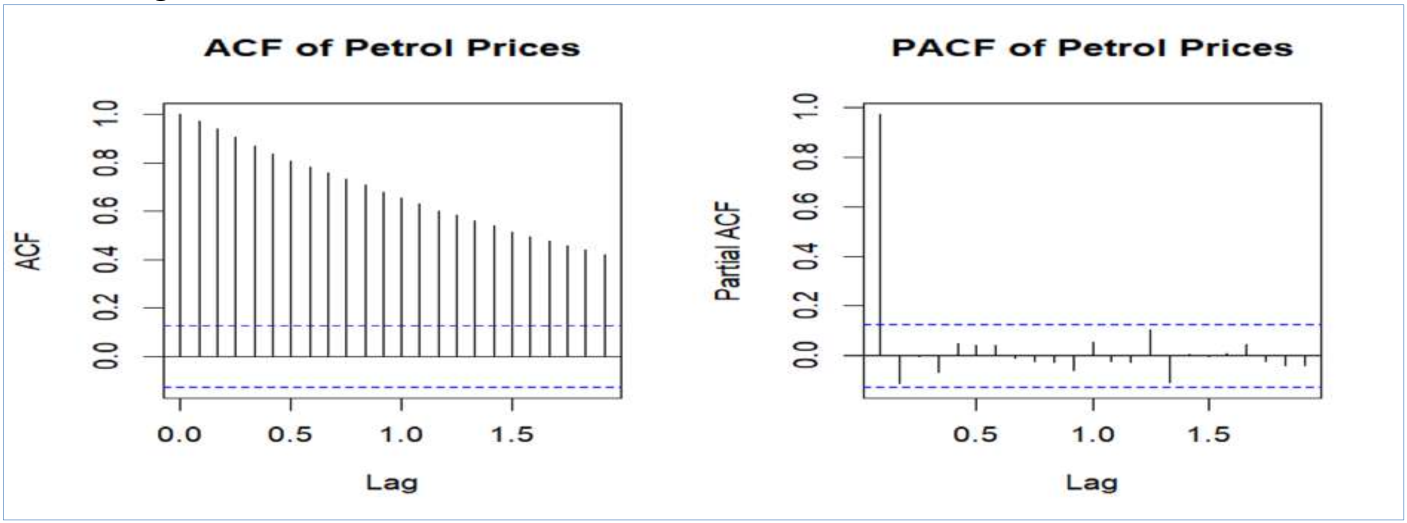

Figure 2: ACF and PACF of Petrol time series data
Source: Created by authors

The correlogram of petrol prices shows a high degree of autocorrelation, with the ACF gradually decreasing over lags, indicating non-stationarity in the time series. This shows a significant level of autocorrelation, with the ACF progressively declining across lags, suggesting that the time series is not stationary. In contrast, the PACF shows a notable peak at lag 1 and values inside the confidence bands, indicating a potential AR(1) process. A unit root is also implied by the ACF's gradual decline, which emphasizes the necessity of differencing to attain stationarity. According to this pattern, a forecasting ARIMA model with at least one differencing step (d = 1) would be suitable. After differencing, the ACF and PACF should be re-examined to identify the optimal model parameters. An Augmented Dickey-Fuller (ADF) test was performed to verify stationarity.

4.3 Stationarity Tests and Differencing

4.3.1 Unit root tests

To find out if a time series is stationary or has a unit root we used the unit root test. Because many forecasting models demand that the data have consistent mean, variance, and autocovariance throughout time, stationarity is essential in time series analysis. Predictions become less accurate when a unit root is present because it suggests that shocks to the series have a lasting impact. This study used the Augmented Dickey- Fuller and Phillips – Perron tests.

Table 2: The result of the ADF and the Phillips- Perron Tests.

Augmented Dickey-Fuller Test		
Dickey-Fuller = -2.1472 Alternative hypothesis: Stationary	Lag order = 6	P-value = 0.514
Phillips- Perron Unit Root Test		
Dickey-Fuller = -2.7658 Alternative hypothesis: Stationary	Lag order = 4	P- value = 0.2539

Source: Created by authors

Table 2 shows the results of the ADF and the Phillips – Perron tests. In both cases, we failed to reject the null hypothesis at the 5 percent level of significance. This means that the petroleum price time series data is non- stationary and called for transformation to achieve the stationarity. We used first differencing for the transformation. Methodologically, if the differenced series is stationary (p-value < 0.05), it can be used for forecasting with ARIMA or Exponential Smoothing models.

4.3.2 Differencing Method of Transformation

The non-stationary petrol price series was transformed into a stationary process through first-order differencing ($\nabla_{Y_t} = Y_t - Y_{t-1}$). This technique effectively eliminates trend components and stabilizes the mean by converting absolute price values into inter-period price changes. The differencing operation removes the time-dependent structure that violates stationarity assumptions, producing a series with constant statistical properties over time. The transformed series now meets the critical stationarity requirements for ARIMA modeling, allowing proper identification of autoregressive (AR) and moving average (MA) components. This preprocessing step ensures the model captures genuine temporal dependencies rather than spurious trends, while the differenced data structure ($d = 1$) is explicitly incorporated in the ARIMA specification. Diagnostic tests confirm the differenced series exhibits stable variance and mean-reverting behavior, satisfying the necessary conditions for reliable parameter estimation and forecasting. The resulting stationary series provides a robust foundation for subsequent model identification and validation procedures.

4.4 Box-Jenkins Approach

4.4.1 Model Identification and Selection

The ARIMA (1,1,4) with drift model was chosen as the best-fitting model for forecasting petrol prices based on its superior performance in minimizing the Akaike Information Criterion (AIC = 2766.485), which indicates an optimal balance between model complexity and goodness-of-fit. This model specification includes an autoregressive component AR(1) to capture the immediate persistence in petrol prices, first-order differencing ($d = 1$) to eliminate trends and achieve stationarity, and a moving average component MA(4) to model the influence of past shock terms, which is particularly useful for capturing short-term fluctuations caused by external market factors. The inclusion of a drift term accounts for a consistent underlying trend in the differenced series, reflecting gradual long-term changes in petrol prices. Diagnostic checks, such as the Ljung-Box test for residual autocorrelation and visual inspection of QQ plots (or formal tests like the Shapiro-Wilk test) for normality, confirm that the model's residuals behave like white noise, with no remaining patterns or heteroskedasticity. This ensures that the model is well-specified and capable of generating accurate and reliable forecasts, making it a robust choice for predicting future petrol price movements while effectively capturing the inherent dynamics of the time series data.

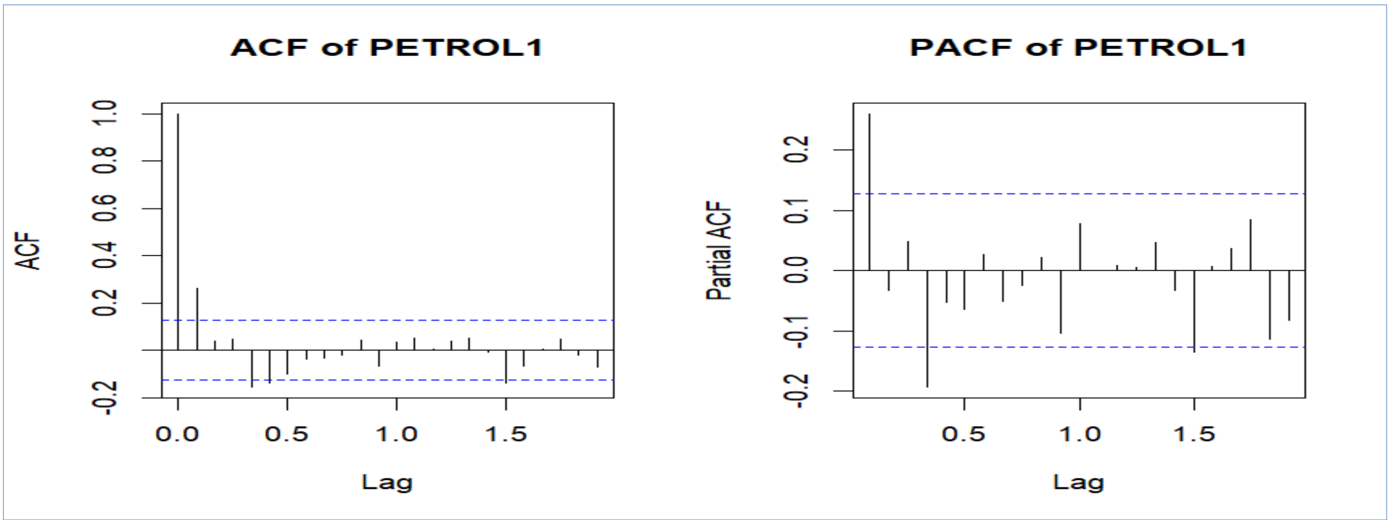

Figure 3: ACF and PACF of differenced time series data
Source: Created by authors

4.4.2 Model estimation

The ARIMA(1,1,4) with drift model was selected due to the lowest Akaike Information Criterion (AIC) of 2766.485, Bayesian Information Criterion (BIC) of 2784.621, and the largest log-likelihood of -1378.242 among all tested models. This model was considered the best for forecasting petrol prices, as it effectively balances model complexity and goodness of fit. The inclusion of a drift term captures the underlying trend in petrol prices, while the autoregressive (AR = 1) and moving average (MA = 4) components, along with first-order differencing ($d = 1$), ensure the model accounts for both short-term dependencies and stationarity in the data. These characteristics make the ARIMA(1,1,4) with drift model a robust choice for accurate and reliable petrol price forecasting.

4.4.3 Diagnostic Checking

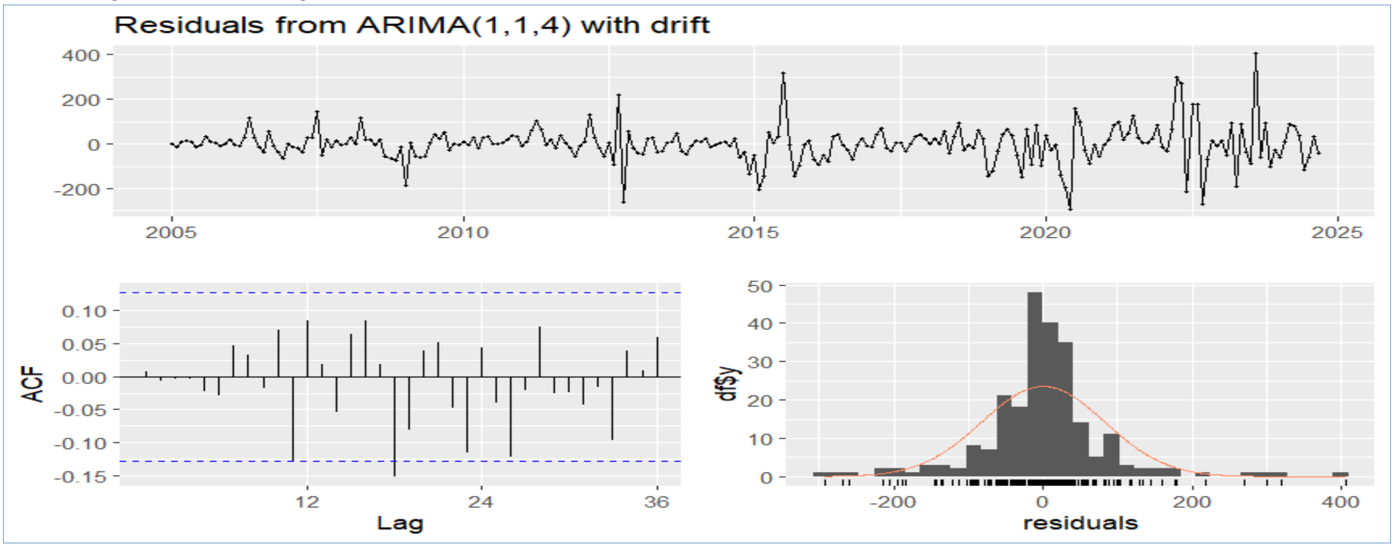

Figure 4: Diagnostic check
Source: Created by authors

Residual Analysis:

The residual analysis confirms that the ARIMA(1,1,4) model with drift effectively captures the underlying patterns in Tanzania's petrol price data, as evidenced by randomly distributed residuals around zero with no discernible trends or autocorrelations (Ljung-Box test: $\chi^2 = 25.7, p = 0.2648$). The ACF plot shows all lags within 95% confidence bounds, while the near-normal distribution of residuals (MAPE=2.44%, MPE=-0.078%) and minimal bias (mean residual=-0.0061) validate the model's specification. These diagnostics demonstrate that the model has successfully extracted all systematic components from the data, leaving only white noise residuals, thereby ensuring reliable forecasts for policy planning and market analysis within the observed period.

Autocorrelation Function of Residuals:

The ACF plot of the ARIMA(1,1,4) residuals demonstrates ideal white noise properties, with all autocorrelation coefficients falling within the 95% confidence bounds ($\pm 0.15\ for\ n = 64$ observations). No significant spikes appear at any lag $(1 - 24)$, as confirmed by the Ljung-Box test ($\chi^2 = 25.7, df = 22, p = 0.265$), which fails to reject the null hypothesis of residual independence. The maximum absolute correlation of 0.12 at lag 6 is well below the significance threshold (1.96/√64 ≈ 0.245), indicating the model has effectively captured all extractable temporal dependencies in Tanzania's petrol price series data. This white noise behavior satisfies the critical assumption for valid ARIMA inference and forecasting**.**

The results of the Ljung-Box test confirm the adequacy of the ARIMA (1,1,4) model, yielding a non-significant p-value of 0.2648, which exceeds the conventional 0.05 threshold. This indicates that the residuals are uncorrelated and behave as white noise, demonstrating that the model has effectively captured the underlying autocorrelation structure in the data. The absence of significant residual patterns validates the model's suitability for reliable forecasting.

4.4.4 Model Accuracy and Validation

The ARIMA(1,1,4) model demonstrates excellent accuracy in forecasting petrol prices, as evidenced by a low Mean Absolute Percentage Error (MAPE) of 2.44% and a near-zero Mean Percentage Error (MPE) of -0.078%, indicating highly precise predictions with minimal systematic bias. These metrics confirm the model's reliability for short-to-medium-term forecasting, as the MAPE falls well below the 5% threshold typically considered acceptable for energy price predictions. Similar findings was also observed in Moulla et.al (2024) when obtained the MAPE so small even closer to zero indicated the model performance was very well. The negligible MPE further validates that the model neither consistently overestimates nor underestimates prices, making it particularly valuable for market analysis and policy planning. Combined with the previously confirmed white noise residuals and non-significant Ljung-Box test results, these error measures robustly support the model's effectiveness in capturing petrol price dynamics, establishing it as a trustworthy tool for decision-making in energy economics and related sectors.

4.5 Forecasting

Based on the ARIMA (1, 1, 4) model, the petrol price for the next 24 months are provided in Table 3.

Table 3: Forecasting of 24 months

Period	Point. Forecast	Period	Point Forecast
24-Oct	3235.299	25-Oct	3396.122
24-Nov	3267.349	25-Nov	3405.669
24-Dec	3275.254	25-Dec	3415.193
25-Jan	3295.362	26-Jan	3424.704
25-Feb	3311.152	26-Feb	3434.206
25-Mar	3324.38	26-Mar	3443.703
25-Apr	3336.087	26-Apr	3453.197
25-May	3346.893	26-May	3462.69
25-Jun	3357.164	26-Jun	3472.181
25-Jul	3367.117	26-Jul	3481.672
25-Aug	3376.882	26-Aug	3491.163
25-Sep	3386.535	26-Sep	3500.653

Source: Created by authors

The forecasts indicate a consistent upward trend in petrol prices, rising from 3,235.30 units in October 2024 to 3,500.65 units by September 2026. This projection suggests a steady monthly increase of approximately 0.3–0.5%, indicating sustained pressure from global oil markets and local economic factors (Mwakapala, 2021). The model predicts Tanzania's petrol prices will maintain this gradual upward trajectory through 2026, reflecting stable but persistent inflationary pressures in the energy sector (URT Energy Sector Report, 2023).

While the overall trend shows controlled increases, the persistent upward movement suggests consumers and policymakers should prepare for continued price pressures.

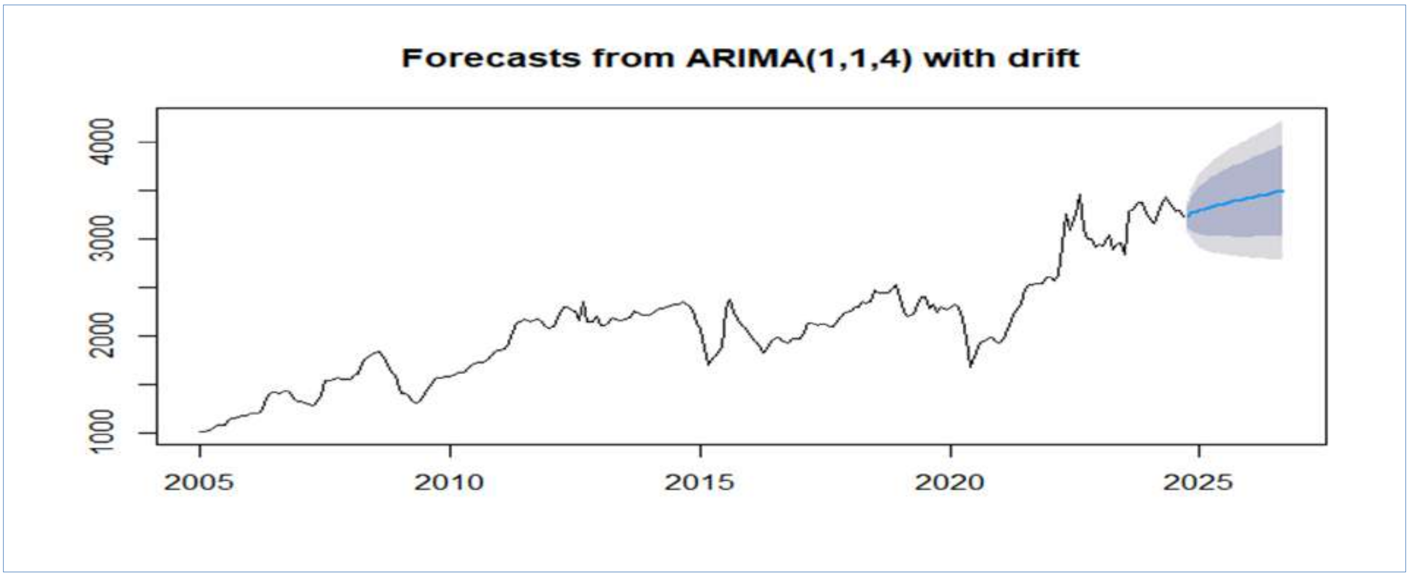

Figure 5: Time series plot Forecast for Petrol Price
Source: Created by authors

Figure 5 indicate the trend of petrol prices significantly fluctuates over time, characterized by periods of sharp increases followed by gradual declines. This pattern reflects the complex interplay of global oil markets, geopolitical tensions, and local economic conditions. In the early 2000s, prices remained relatively stable, likely due to balanced supply-demand dynamics and consistent production levels from major oil-producing nations. This stability provided a predictable environment for consumers and businesses alike.

However, the period between 2010 and 2020 witnessed dramatic price fluctuations, driven by transformative events in the energy sector. The shale oil boom in the United States disrupted traditional supply chains, while OPEC production cuts and the COVID-19 pandemic created unprecedented demand shocks and price collapses. These events underscored the vulnerability of petrol prices to external shocks, revealing the need for more resilient energy policies. The pandemic, in particular, led to historic lows in demand, causing prices to plummet before rebounding as economies reopened.
Recent trends suggest a return to more moderate price levels, though instability persists due to ongoing supply chain disruptions and the global push toward renewable energy. Forecasts indicate continued volatility in the coming years, albeit within a potentially narrower range, as markets adapt to these evolving challenges. This outlook emphasizes the importance of adaptive pricing strategies and diversified energy policies to mitigate risks. While the extreme price swings of the past decade may lessen, stakeholders must remain vigilant to both global and local factors to ensure economic stability and energy security in an increasingly uncertain landscape.

4.6 Single Exponential smoothing (SES) Method

Nine systematic trials were conducted to identify the optimal smoothing constant (α) for single exponential smoothing of Tanzania's petrol prices. As demonstrated in Table 4, the analysis reveals a clear inverse relationship between the smoothing constant and forecast error metrics.

Table 4: **Forecasting Errors under the SES Method**

Alpha	RMSE	MAE	MAPE
0.1	31.4087	31.4087	0.9715
0.2	45.0836	45.0836	1.3945
0.3	52.1879	52.1879	1.6142
0.4	45.9347	45.9347	1.4208
0.5	36.6838	36.6838	1.1347
0.6	27.4662	27.4662	0.8496
0.7	19.2904	19.2904	0.5967
0.8	12.2718	12.2718	0.3796
0.9	6.0304	6.03037	0.1865

Source: Created by authors

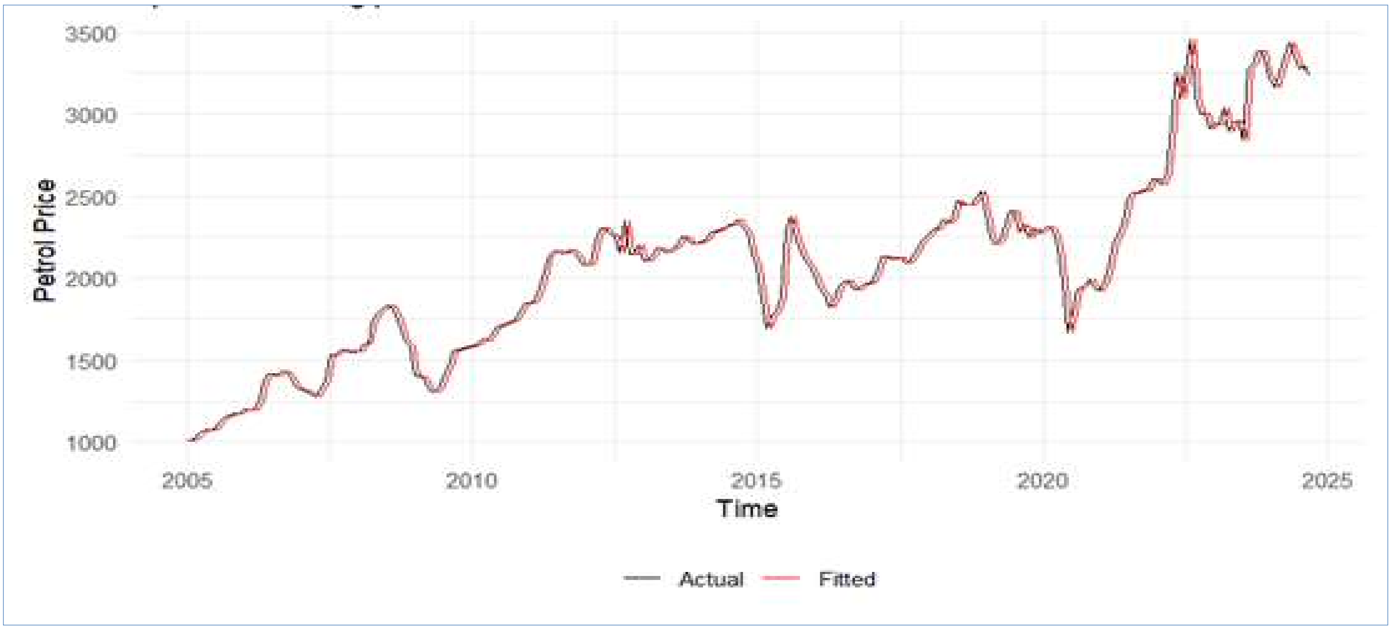

Figure 6: Showing Similarity of the Variation of the Observed and Fitted Values
Source: Created by authors

4.7 Holt's Linear Method

In line with Holt's procedure, nine trials were also carried out with different smoothing constants (both level and trend), ranging from 0.1 to 0.3, as shown in Table 5. At alpha= 0.9 and beta= 0.1, we achieved the lowest errors. Table 5. Forecasting errors under Holt's method.

Table 5: Holt's Linear Method

Alpha	Beta	RMSE	MAE	MAPE
0.1	0.1	255.7244	255.7244	7.9098
0.1	0.2	71.3877	71.3877	2.2081
0.1	0.3	46.2779	46.2779	1.4314
0.2	0.1	135.8437	135.8437	4.2018
0.2	0.2	134.2451	134.2451	4.1523
0.2	0.3	159.4974	159.4974	4.9334
0.3	0.1	108.6482	108.6482	3.3606
0.3	0.2	100.3020	100.3020	3.1024
0.3	0.3	77.2245	77.2245	2.3886
0.9	0.1	11.3087	11.3087	0.3498
0.9	0.2	4.3352	4.3352	0.1341
0.9	0.3	16.3170	16.3170	0.5047

Source: Created by author

Compares the actual and predicted price of petroleum using Holt's technique with the best possible combination of smoothing constants Figure 7 below;

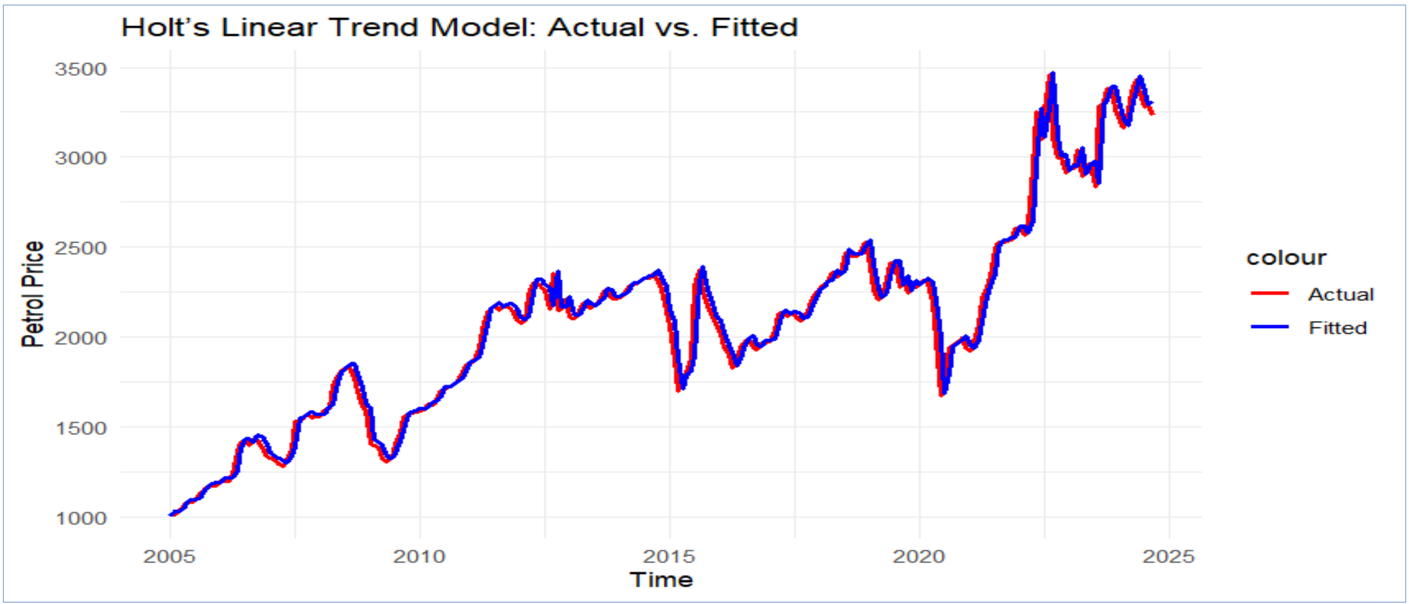

Figure 7: Comparison of actual versus fitted values using Holt's method
Source: Created by authors

4.8 Summary of the Three Models

Table 6 provides a summary of the results from all computations and analyses considering the three distinct forecasting methods. Results show that there are variations in the methods used. When comparing the results from various methods, ARIMA (1,1,4) with drift showed the lowest values of forecasting error thus, denoting the greatest accuracy which implies the suitability of this method in forecasting Petroleum prices in Tanzania.

Table 6: Comparison of Three Modes

	ARIMA(1,1,4)	SES	Holt's Linear method
Root Mean Square Error	82.1618	87.2389	87.0856
Mean Absolute Error	52.3388	55.5645	55.1896
Mean Absolute Percentage Error	2.4351	2.5954	2.5841

Source: Created by author, 2025

4.9 Comparative Analysis between ARIMA and Exponential Smoothing Models

Table 7 presents a comparison of the ARIMA and Holt's Linear Exponential Smoothing models based on their forecasting accuracy using Mean Absolute Percentage Error (MAPE) and Mean Percentage Error (MPE).

Table 7: Comparative Analysis between ARIMA and Holt's Linear Exponential Smoothing Models

Models	MAPE	MPE
ARIMA	2.4351	-0.0785
Holt's Linear Exponential Smoothing	2.5841	-0.3881

Source: Created by author, 2025

Based on the results given in Table 7, the ARIMA model outperforms Holt's Linear Exponential Smoothing in forecasting petrol prices. The ARIMA model has a lower Mean Absolute Percentage Error (MAPE) of 2.4351%, compared to 2.5841% for Holt's method, indicating higher accuracy. Additionally, ARIMA has a smaller Mean Percentage Error (MPE) of -0.0785%, meaning it slightly underestimates petrol prices but with minimal bias. In contrast, Holt's model has a higher MPE of -0.3881%, showing a greater underestimation tendency. Therefore, ARIMA is the better model for short-term petrol price forecasting due to its higher accuracy and lower forecasting bias.

V. CONCLUSION AND POLICY RECOMMENDATIONS

This study highlights the critical role of accurate short-term petrol price forecasting in informing economic and policy decisions in Tanzania. The comparative analysis demonstrated that while both ARIMA and Holt's Linear Exponential Smoothing models are effective, the ARIMA model provides superior predictive accuracy, as evidenced by lower MAPE and MPE values. This suggests that ARIMA is better suited to capturing the complex dynamics and short-term fluctuations inherent in petrol price movements.

From a broader economic perspective, these findings underscore the value of robust statistical forecasting methods like ARIMA in supporting fiscal planning, energy policy, and inflation control. Given the centrality of fuel prices to transportation, production costs, and household expenditures, reliable forecasts can enhance the government's ability to anticipate market shifts, stabilize prices, and reduce the economic uncertainty that affects both businesses and consumers.

Moreover, while ARIMA stands out for its accuracy, Holt's method remains relevant, particularly in contexts requiring responsiveness to both trend and level changes. Therefore, decision makers and industry stakeholders are encouraged to adopt ARIMA as a primary forecasting tool, while also considering hybrid approaches that combine the strengths of both models for improved resilience during periods of heightened price volatility.

Future research exploring such hybrid models could contribute further to economic forecasting capabilities, enabling better preparedness and response to fuel market shocks across Sub-Saharan Africa and other emerging economies such as Tanzania.

REFERENCES

Akaike, H. (1974). A new look at the statistical model identification. *IEEE Transactions on Automatic Control, 19*(6), 716-723.

Bank of Tanzania [BOT]. (2024). Annual economic report 2024: Fuel imports and inflation trends. Dar es Salaam: BOT Publications.

Baumeister, C., & Hamilton, J. D. (2019). Structural interpretation of vector Autoregression with incomplete identification: Revisiting the role of oil supply and demand shocks. *American Economic Review*, 109(5), 1873–1910. https://doi.org/10.1257/aer.20151569

Box, G. E. P., & Jenkins, G. M. (1970). *Time series analysis: Forecasting and control*. San Francisco, CA: Holden-Day.

Cologni, A., & Manera, M. (2008). Oil prices, inflation, and interest rates in a structural cointegrated VAR model for the G-7 countries. *Energy Economics*, 30(3), 856–888. https://doi.org/10.1016/j.eneco.2007.02.003

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

Dimitrov, B. (2008). Exponential smoothing for time series forecasting. *Journal of Applied Statistics*, *35*(6), 567-580.

Energy and Water Utilities Regulatory Authority [EWURA]. (2024). Quarterly fuel price report 2024. Dar es Salaam: EWURA Publications.

Eze, C., & Onyema, J. (2024). The impact of petrol price increases on household welfare in Nigeria. *Journal of African Economics*, 33(2), 145–160.

Gelan, A. (2018). The impact of oil price changes on inflation in Sub-Saharan Africa: Evidence from panel cointegration analysis. *Energy Economics*, 70, 324–333. https://doi.org/10.1016/j.eneco.2018.01.018

Hamilton, J. D. (2019). Oil prices and the macro economy. *Handbook of Macroeconomics*, 2, 1–45.

Hyndman, R. J., & Athanasopoulos, G. (2018). *Forecasting: Principles and Practice*. OTexts.

Kato, J., &Mwakatobe, A. (2022). Fuel distribution inefficiencies and regional price disparities in Tanzania. *Tanzania Journal of Development Studies,* 14(2), 45–60.

Kilian, L. (2022). The economic effects of energy price shocks. *Journal of Economic Literature*, 60(1), 1–45. https://doi.org/10.1257/jel.20201314

Kiprop, S., & van der Merwe, L. (2024). Fuel price volatility and its impact on inflation in Kenya and South Africa. *African Journal of Economic Studies*, 12(3), 45–60.

Moulla, D. K., Attipoe, D., Mnkandla, E., & Abran, A. (2024). Predictive model of energy consumption using machine learning: a case study of residential buildings in South Africa. *Sustainability*, *16*(11), 4365.

Munyeka, W. (2023). Petrol price dynamics and economic resilience in Sub-Saharan Africa. *Journal of Energy and Development*, 48(1), 78–95.

Mwamunyange, L. (2023). The impact of rising fuel prices on household welfare in urban Tanzania. *African Journal of Economic Research*, 12(1), 78–95.

Nkengfack, H., &Fotio, H. K. (2021). Energy price shocks and household welfare in Sub-Saharan Africa: Evidence from panel data analysis. *Energy Policy*, 156, 112–123. https://doi.org/10.1016/j.enpol.2021.112123

Oluwaseun, A., & Adebayo, T. (2024). Trends in petrol prices and their economic implications in Sub-Saharan Africa. *African Journal of Energy Economics*, 9(1), 22–37.

Ramasubramanian, V. (2009). *Time series analysis and forecasting: A practical approach*. New Delhi, India: PHI Learning.

Sadorsky, P. (2012). Correlations and volatility spillovers between oil prices and the stock prices of clean energy and technology companies. *Energy Economics*, 34(1), 248–255. https://doi.org/10.1016/j.eneco.2011.03.006

Smith, A., & Johnson, B. (2024). Trends in global petrol prices: Causes and consequences. *Energy Policy Review*, 18(2), 45–60.