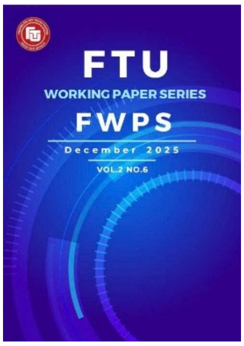

Working Paper 2025.2.6.13

- Vol. 2, No. 6

ỨNG DỤNG MÔ HÌNH ARIMA TRONG DỰ BÁO NGẮN HẠN TỶ GIÁ HỐI ĐOÁI CỦA VIỆT NAM (USD/VND) TRONG NGẮN HẠN

Phạm Ngọc Nữ [1]

Sinh viên K62 CLC Kinh tế Đối ngoại - Viện Kinh tế và Kinh doanh quốc tế

Trường Đại học Ngoại thương, Hà Nội, Việt Nam

Phùng Duy Quang

Giảng viên Khoa Công nghệ và Khoa học dữ liệu

Trường Đại học Ngoại thương, Hà Nội, Việt Nam

Tóm tắt

Nghiên cứu này áp dụng mô hình ARIMA (Autoregressive Integrated Moving Average) để dự báo tỷ giá USD/VND dựa trên dữ liệu quan sát hằng ngày từ tháng 1/2022 đến tháng 4/2025. Phương pháp Box-Jenkins được sử dụng nhằm xây dựng và kiểm định mô hình ARIMA tối ưu thông qua các bước phân tích chẩn đoán dữ liệu, kiểm định tính dừng, nhận dạng mô hình, ước lượng tham số và đánh giá khả năng dự báo. Mô hình ARIMA(3,1,1) được lựa chọn dựa trên ý nghĩa thống kê và các tiêu chí thông tin. Nhiều thước đo hiệu quả dự báo, bao gồm RMSE, MAE, MAPE, MASE và kiểm định tự tương quan phần dư, được sử dụng để đánh giá độ chính xác của mô hình. Kết quả cho thấy mô hình cuối cùng có khả năng dự báo tốt, đặc biệt trong việc nắm bắt các biến động ngắn hạn của tỷ giá với mức tự tương quan phần dư rất thấp và sai số phần trăm nhỏ. Điều gợi ý rằng mô hình ARIMA tiếp tục là công cụ đáng tin cậy cho dự báo tỷ giá hối đoái tại thị trường Việt Nam trong ngắn hạn.

Từ khóa: ARIMA, mô hình chuỗi thời gian, thị trường ngoại hối (FOREX), dự báo tỷ giá, USD/VND

[1] Tác giả liên hệ, Email: k62.2312150185@ftu.edu.vn

APPLICATION OF ARIMA MODEL IN FORECASTING VIETNAM'S FOREIGN EXCHANGE RATE (USD/VND) IN THE SHORT TERM

Abstract

This study applies the ARIMA (Autoregressive Integrated Moving Average) model to forecast the USD/VND exchange rate using daily observations from January 2022 to April 2025. We adopt the Box - Jenkins methodology to build and validate an optimal ARIMA model through data diagnostics, stationarity testing, model identification, parameter estimation, and forecast evaluation. The selected model, ARIMA(3,1,1), was identified based on statistical significance and information criteria. Multiple performance metrics - including RMSE, MAE, MAPE, MASE, and residual autocorrelation - are employed to assess forecast accuracy. The final model demonstrates good predictive performance, particularly in capturing short-term exchange rate dynamics with minimal residual autocorrelation and a very low percentage error. The results suggest that the ARIMA model remains a reliable tool for short-term exchange rate forecasting in the recent Vietnamese context.

Keywords: ARIMA model, time series analysis, foreign exchange market (FOREX), exchange rate forecasting, USD/VND

1. Introduction

Foreign exchange (FOREX) refers to the price of one currency relative to another, and plays a crucial role in international trade and finance. Every cross-border transaction - whether by individuals, corporations, or governments - requires currency conversion, making the foreign exchange market the largest and most liquid financial market globally. Due to its high volatility and significant influence on global capital flows, the FOREX market presents both opportunities and risks. This is particularly relevant for developing economies, where currency fluctuations can greatly affect inflation, trade balances, and investment decisions. For such countries, accurate exchange rate forecasting is vital - not only for investors seeking profit and businesses aiming to hedge exposure, but also for policymakers striving to manage economic stability.

Among the many statistical approaches for time series forecasting, the ARIMA model, developed by Box and Jenkins (1970), has remained a foundational method due to its simplicity and effectiveness in modeling linear trends. This paper applies the ARIMA model to forecast the USD/VND exchange rate using daily data from 2022 to 2025. By focusing on short-term prediction accuracy, the study aims to contribute empirical insights for both market participants and monetary authorities in Vietnam.

Compared with previous ARIMA-based studies in Vietnam, this research utilizes high-frequency daily data from 2022 to 2025, a period marked by post-pandemic recovery, global inflationary pressures, and significant monetary tightening by the U.S. Federal Reserve. These factors have led to unprecedented short-term volatility in the USD/VND exchange rate, capturing the dynamics of an inflation-driven and post-pandemic economic environment. Consequently, this study provides new empirical evidence on the applicability of ARIMA models in the Vietnamese foreign exchange market under rapidly evolving conditions.

The paper is structured as follows: Section 2 presents a review of relevant literature on exchange rate forecasting using ARIMA and alternative models. Section 3 describes the dataset used in the analysis. Section 4 introduces the ARIMA methodology. Section 5 reports and discusses the empirical results. Finally, Section 6 concludes the paper with key findings and several implications.

2. Literature Review

ARIMA (Autoregressive Integrated Moving Average), introduced by Box and Jenkins (1970), has long been recognized as a robust and widely used method for time series forecasting. Its application spans across various fields, particularly in macroeconomic and financial domains. For example, Meyler et al. (1998) successfully employed ARIMA models to forecast inflation in Ireland, while Guha and Bandyopadhyay (2016) applied it to gold price forecasting.

2.1. ARIMA Models in Exchange Rate Forecasting

Specifically in exchange rate forecasting, Rasheed (2019) used five years of data to forecast the USD/PKR exchange rate and concluded that ARIMA(1,1,9) yielded the most accurate forecasts among several candidates. Nwankwo (2014), using annual data from 1982 to 2011, showed that ARIMA(1,0,0) was suitable for forecasting the Nigerian naira to U.S. dollar exchange rate. Adetunde et al. (2011) modeled the Ghanaian cedi/USD exchange rate using ARIMA(1,1,1), achieving excellent statistical fit based on MAPE, RMSE, and R^2 metrics. Quinta et al. (2016) applied ARIMA to model the Indonesian Rupiah against the US dollar and attained a forecast accuracy of over 98.74%. Yildiran et al. (2017), using over 3,000 daily observations of the Turkish Lira, also validated ARIMA's performance for both short- and long-term forecasts. Ahmed and Keya (2019) used the Box-Jenkins methodology and identified ARIMA(2,1,1) as the optimal model over a 44-year dataset.

2.2. Comparison with Alternative Forecasting Models

Several studies have explored the relative accuracy of alternative models in predicting exchange rates. In Pakistan, Khan et al. (2013) applied GARCH(1,1) to model exchange rate volatility and found international currencies (like USD) more stable than local currencies such as PKR. Similarly, Naeem et al. (2020) compared five machine learning methods - logistic regression, random forest, bagging, naïve Bayes, and SVM - finding that logistic regression yielded the highest forecast accuracy (82.14%) for USD/PKR exchange rate. Akhtar et al. (2022) applied ARIMA and GARCH models to forecast the USD/PKR exchange rate using daily data, concluding that while ARIMA captured short-term trends, ARCH modeling outperformed GARCH in volatility prediction. These results highlight that while ARIMA is useful for trend prediction, models like GARCH or ML algorithms may be better suited for capturing volatility or regime shifts.

Hybrid and machine learning models have also emerged as strong alternatives. Wang et al. (2016) and Khashei et al. (2020) demonstrated that combining ARIMA with Artificial Neural Networks (ANN) can significantly enhance forecast accuracy. Matroushi (2011) compared ARIMA, ARIMA-ANN, and ARIMA-MLP models and concluded that the hybrid MLP model outperforms both standalone ARIMA and ANN in capturing the nonlinear dynamics of exchange rates.

Studies by Dunis et al. (2006) and Amat et al. (2018) demonstrated that no single model consistently outperforms across all scenarios. Ensemble and hybrid models combining ARIMA, GARCH, and ML components often yield superior predictive performance. Nevertheless, for short-term forecasting, ARIMA continues to be considered a reliable baseline by many researchers, including Liu and Lv (2011).

2.3. ARIMA Applications in Vietnam

In Vietnam, ARIMA has been applied to different financial and macroeconomic time series. Nguyen et al. (2025) compared ARIMA and deep learning models in predicting stock prices of Vietnam's five largest firms and found that while deep learning slightly outperformed, ARIMA remained a useful baseline for linear trend forecasting. Quang et al. (2024) used a hybrid ARIMA-GARCH model to forecast Bitcoin prices and reported that incorporating GARCH improved predictive accuracy over the standalone ARIMA model. Le and Nguyen (2020) applied ARIMA(3,1,3) to model Vietnam's annual GDP growth from 1985 to 2025, demonstrating its applicability to medium-term macroeconomic forecasting despite sensitivity to structural shocks.

A direct application of ARIMA to Vietnam's foreign exchange market was conducted by Tran (2016), who modeled the USD/VND exchange rate using monthly data from 2013 to 2015. The study followed the Box–Jenkins methodology and demonstrated that while ARIMA performs well in the short term (under 15 days), its accuracy degrades over longer horizons. To address this limitation, Tran proposed a rolling forecast strategy that continuously updates the model's inputs to maintain short-term accuracy. More recently, Tu et al. (2024) expanded upon this approach by combining ARIMA with two machine learning techniques, Random Forest (RF) and Artificial Neural Network (ANN), to forecast the VND/USD rate using data from 2000 to 2023. Their hybrid models, ARIMA-ANN and ARIMA-RF, split the forecasting task into two components: the linear structure handled by ARIMA, and the nonlinear patterns captured by ANN or RF. The results show that while standalone machine learning models outperformed ARIMA in accuracy, the hybrid methods yielded even better predictive performance.

Although recent studies have introduced machine learning and hybrid models, traditional ARIMA models remain valuable for short-term exchange-rate forecasting, especially in emerging economies such as Vietnam. Its advantages include transparent structure and strong performance on high-frequency linear data. In this study, we apply the ARIMA model to high-frequency daily exchange rate data over a longer and more recent time horizon (2022–2025), capturing the dynamics of the post-pandemic and inflation-driven economic environment. Moreover, the study conducts a comprehensive evaluation of forecasting performance using multiple error metrics (MAPE, MAE, RMSE, and R^2). As such, it contributes new empirical evidence to the application of ARIMA in the Vietnamese foreign exchange market under evolving market conditions.

3. Data Description

The time-series data used in this analysis is denoted as $FOREX_t$, representing the daily exchange rate between the U.S. dollar (USD) and the Vietnamese dong (VND), where USD is the base currency and VND is the pricing currency. The rates reflect the adjusted closing price of USD expressed in VND. The dataset spans the period from *January 1, 2022 to April 30, 2025*, comprising a total of 866 daily observations. The exchange rate is treated as the dependent variable, while the corresponding Date serves as the independent time index.

Figure I. Daily USD/VND Exchange Rate from Jan 2022 to Apr 2025

Figure I illustrates the evolution of the USD/VND exchange rate over the sample period. The exchange rate shows a general upward trend, indicating a depreciation of the VND against the USD. The exchange rate increased from 22,825 to 25,980 VND per USD over the sample period. Basic descriptive statistics show a mean of 24,228, a median of 24,268, with a maximum of 26,026 and a minimum of 22,625 VND per USD. There are several sharp increases in late 2022 and early 2024, followed by periods of correction or stabilization.

The observed upward trend in the series suggests a potential random walk of the exchange rate, which means it may require techniques suitable for non-stationary time series, such as ARIMA with differencing, ARCH/GARCH models to account for volatility clustering, or other trend-following approaches.

4. Research Methodology

Box–Jenkins ARIMA is known as the ARIMA(p,d,q) model, where p is the number of autoregressive (AR) terms, d is the number of differences taken, and q is the number of moving average (MA) terms. ARIMA models always assume the variance of data to be constant. Gujarati (2014) introduced four steps to estimate an ARIMA model: 1. Recognizing the model; 2. Estimating variables and choosing the model; 3. Testing the model; and 4. Forecasting (Tran et al., 2016). The steps can be explained as follows:

4.1. Recognizing the Model

The ARIMA model combines three components: autoregressive (AR) terms, integrated (I) differencing, and moving average (MA) terms.

AR(p) – Autoregressive Process of Order p:
A time series Y_t is said to follow an AR(p) process if it depends linearly on its own past values:

$$Y_t = \nu + a_1 Y_{t-1} + a_2 Y_{t-2} + \cdots + a_p Y_{t-p} + \varepsilon_t$$

This can be expressed using the lag operator L as:

$$A(L)Y_t = \nu + \varepsilon_t \text{ where } A(L) = 1 - a_1 L - a_2 L^2 - \cdots - a_p L^p$$

MA(q) – Moving Average Process of Order q
A time series Y_t follows an MA(q) process if it is a linear function of current and past error terms:

$$Y_t = \nu + \varepsilon_t + b_1 \varepsilon_{t-1} + b_2 \varepsilon_{t-2} + \cdots + b_q \varepsilon_{t-q}$$

Which can also be expressed as:

$$Y_t = \nu + B(L)\varepsilon_t \text{ where } B(L) = 1 + b_1 L + b_2 L^2 + \cdots + b_q L^q$$

ARMA(p, q) – Combined Model
When both AR and MA components are present:

$$A(L)Y_t = \nu + B(L)\varepsilon_t$$

ARIMA(p, d, q) – General Model with Differencing
To model non-stationary time series, the series is differenced $d$times to achieve stationarity, and the ARMA model is applied to the differenced data:

$$A(L)(1 - L)^d Y_t = \nu + B(L)\varepsilon_t$$

Here, d represents the number of times the series is differenced to remove trends or seasonality.

4.2. Estimation and Model Selection

The modeling process is conducted using *RStudio*, applying relevant time series packages such as *forecast* and *tseries*. The steps involved in estimating the ARIMA model include:

- *Stationarity Testing*: The original exchange rate series is tested for stationarity using the Augmented Dickey–Fuller (ADF) test. If the series is found to be non-stationary, differencing is applied until stationarity is achieved. The number of differences required determines the integration order d
- *Identification of AR and MA Orders*: Once stationarity is attained, the Autocorrelation Function (ACF) and Partial Autocorrelation Function (PACF) plots are used to preliminarily determine appropriate values for the autoregressive order p and moving average order q
- *Model Selection:* Several ARIMA(p,d,q) models are estimated, and the optimal specification is chosen based on information criteria such as the Akaike Information Criterion (AIC) and Bayesian Information Criterion (BIC). The model with the lowest AIC/BIC values and statistically significant coefficients (with p-values < 0.05) is selected

.

4.3. Diagnostic Checking

After selecting a candidate model, diagnostic tests are performed to evaluate its adequacy

- *Ljung–Box Q Test*: Conducted to assess whether the residuals are independently distributed, i.e., exhibit white noise behavior
- *ACF plots of residuals*: Examined to ensure there is no remaining autocorrelation.
- *ADF test on residuals*: applied to confirm the stationarity of residuals.

Awell-specified ARIMA model should have residuals that are uncorrelated and stationary, indicating that the time series dynamics have been adequately captured.

4.4. Forecast Evaluation

After model estimation and diagnostics, the model's in-sample performance is evaluated using forecast accuracy metrics, including Root Mean Squared Error (RMSE), Mean Absolute Error (MAE), Mean Absolute Percentage Error (MAPE), Coefficient of Determination (R2) These indicators provide a quantitative assessment of the model's goodness-of-fit and help determine the predictive reliability of the selected ARIMA specification.

5. Results & Findings

After understanding the ARIMA method, we use *RStudio* to apply statistical techniques on data from *1 January 2022 to 30 April 2025.*

5.1. Descriptive Statistics

The original series (the USD/VND exchange rate) is denoted as $FOREX_t$, and the first difference of the original series is represented by $dFOREX_t = FOREX_t - FOREX_{t-1}$.

We then use descriptive statistics to provide an overview of both the original series $FOREX_t$ and the differenced series $dFOREX_t$.

	$FOREX_t$	$dFOREX_t$
Mean	24228.092	3.647
Median	24268.000	0.000
Maximum	26026.000	290.000
Minimum	22625.000	-370.000
Standard Deviation	911.383	46.870
Skewness	0.050	-0.250
Kurtosis	1.737	0.000
Jarque–Bera Statistic	57.890	4110.363

	$FOREX_t$	$dFOREX_t$
Jarque–Bera p-value	0.000	0.000
Sum	20981528	3155.000
Sum of Squared Deviations	718485933	1898039.457
Observations	866	865

Table 1. Descriptive Statistics of the Original Series ($FOREX_t$) and Its First Difference ($dFOREX_t$)

- *Skewness*: The original series $FOREX_t$ has positive skewness, nearly zero, suggesting a roughly symmetric distribution. Whereas, $dFOREX_t$ is slightly negatively skewed, which means the left tail is a little more extreme. This implies that negative changes in the exchange rate occur with slightly more frequency than positive changes.
- *Kurtosis*: $FOREX_t$ has a kurtosis of 1.74, which is less than the benchmark value of 3, indicating a platykurtic distribution. In contrast, the $dFOREX_t$ series has a kurtosis of 13.67, significantly exceeding 3, indicating a leptokurtic distribution. This implies the presence of fat tails and a higher likelihood of extreme values or outliers.
- *Jarque–Bera test:* Test statistics for both series are very high, with p-values of 0.0000, representing the rejection of the null hypothesis even at the 1% significance level. This indicates that both series follow a non-normal distribution.

Figure II: Kernel density plots of $FOREX_t$ (left) and $dFOREX_t$ (right)

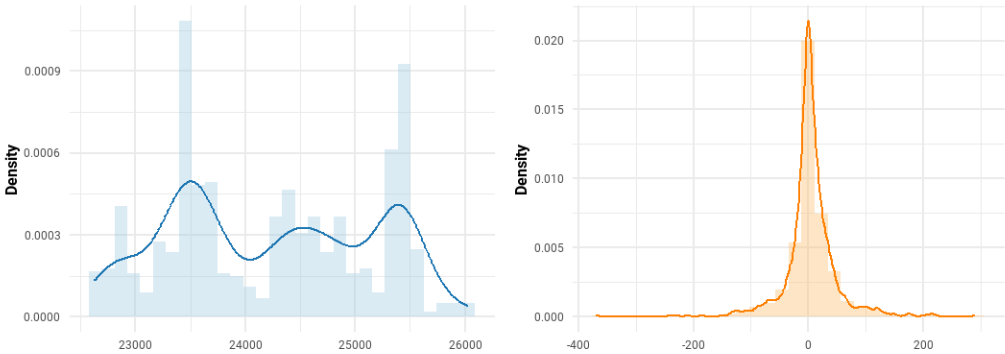

5.2. Testing Stationary

Stationarity is the basic assumption to be fulfilled for time series forecasting. In this paper, we use the augmented Dickey–Fuller (ADF) test introduced by Dickey and Fuller in 1981 to diagnose stationarity in the time series. For $FOREX_t$ data, the hypothesis can be defined as follows:

- Null Hypothesis (H0): $\phi = 0 \rightarrow$ The series $FOREX_t$ is non-stationary.
- Alternative Hypothesis (H1): $\phi \neq 0 \rightarrow$ The series $FOREX_t$ is trend-stationary.

Model 1	ADF statistic	Critical values	Results' interpretation
		-3.96 at 1% level	Nonstationary
$FOREX_t$	-1.736922	-3.41 at 5% level	Nonstationary
		-3.12 at 10% level	Nonstationary

Table 2. ADF test for $FOREX_t$ at levels with trend and intercept (trend-stationary)

Because –1.736922 > –3.41, we do not reject the null hypothesis at the 5% level. Therefore, the series $FOREX_t$ is said to be non-stationary and contains a unit root with drift. In order to fulfill the assumption of the ARIMA model, $FOREX_t$ must have a unit root. The next step is to transform the series to make it stationary. Hence, $dFOREX_t$ is computed by taking the first difference of $FOREX_t$.

For $dFOREX_t$ data series, the hypotheses of the ADF test can be defined as follows:

- Null Hypothesis (H0): $\phi = 0 \rightarrow$ The series is non-stationary.
- Alternative Hypothesis (H1): $\phi < 0 \rightarrow$ $dFOREX_t$ is stationary around a non-zero mean.

Model 2:	ADF statistic	Critical values	Results' interpretation
		-3.43 at 1% level	Stationary
$dFOREX_t$	-24.80931	-2.86 at 5% level	Stationary
		-2.57 at 10% level	Stationary

Table 3. ADF test for $dFOREX_t$ at levels with trend and intercept (trend-stationary)

Because –24.80931 < –3.43, we reject the null hypothesis at the 1% level. Therefore, the series $dFOREX_t$ is said to be stationary around its mean and will be used for further analysis.

5.3. Correlogram

Approximate significance bounds can be constructed, as shown by the red lines in the plots, to help identify large values. The approximate (1–α)×100% significance bounds for the autocorrelation

$$\pm z_{1-\frac{\alpha}{2}} \times \frac{1}{\sqrt{n}}$$

where $z_{1-\frac{\alpha}{2}}$ is the critical value from the standard normal distribution and n is the sample size.

In this paper, we use a 5% significance level, α = 0.05 and $z_{1-\frac{\alpha}{2}} = z_{0.975} = 1.96$

For the $FOREX_t$ series with $n = 866$ observations, the approximate bounds ar

$$\pm 1.96 \times \frac{1}{\sqrt{886}} \approx \pm 0.0666$$

.

For the *dFOREX*$_t$ series with $n = 865$ observations:

$$\pm 1.96 \times \frac{1}{\sqrt{885}} \approx \pm 0.0666$$

Values lying outside these bounds indicate statistically significant correlations. The correlograms for the *FOREX*$_t$ and *dFOREX*$_t$ series are shown in the Figure III below.

Figure III. Correlogram of *FOREX*$_t$ and *dFOREX*$_t$

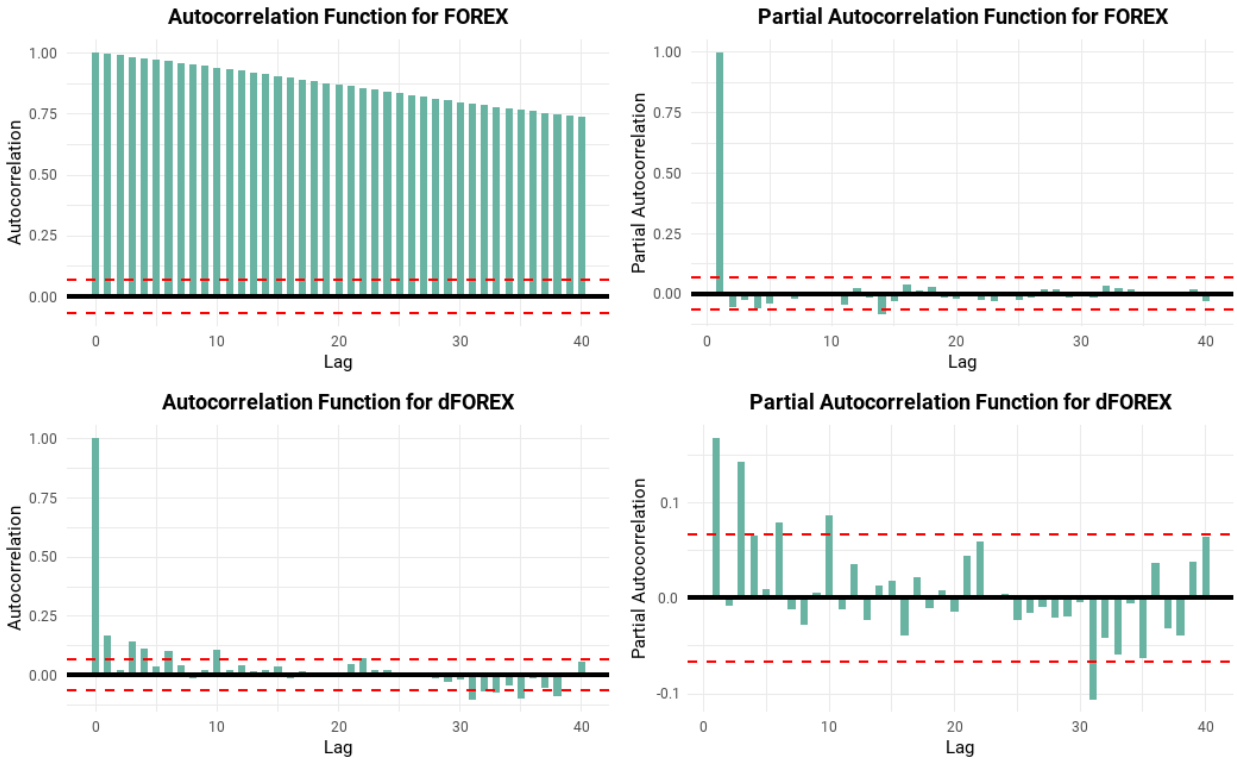

FOREX$_t$ Series:

- ACF: The autocorrelation decreases very slowly and remains significant over a long lag period, showing almost no clear decline within 40 lags. This indicates non-stationarity.
- PACF: There is only one large spike at lag 1, which suggests an AR(1) process. This series is non-stationary, so a differencing transformation is needed.

dFOREX$_t$ Series:

- ACF: The autocorrelation drops sharply after lag 1 and then fluctuates around zero. This pattern is typical of a stationary series.
- PACF: There are several significant spikes at lags 1 and 3, with smaller but still noticeable spikes at lags 6 and 10, which exceed the confidence bounds. This suggests that an autoregressive model of order 3 may be appropriate, while models of order 6 (AR(6)) or 10 (AR(10)) could also be considered.

5.4. Model Selection and Application

A variety of criteria are available for determining the most appropriate ARIMA model order, which may include the following features (simultaneously considered):

- the most significant coefficients
- the highest adjusted R^2
- the lowest sigma square (volatility)
- the lowest information criteria: the Akaike Information Criterion (AIC) and the Schwarz Information Criterion (SIC)

Model	Sig. Coeff.	Sigma Sq.	Adj. R^2	AIC	SIC	DW
ARIMA(1,1,1)	0	2144.976	0.997418	9093.868	9108.156	2.001505
ARIMA(3,1,1)	4	2092.610	0.997481	9074.565	9098.379	2.001000
ARIMA(6,1,1)	2	2089.713	0.997484	9076.374	9114.476	2.003833
ARIMA(10,1,1)	3	2080.367	0.997495	9076.543	9133.696	2.005711

Table 4. ARIMA Model Evaluation Summary

Table 4 represents the outcome of tentative ARIMA models. Among the evaluated models, ARIMA(3,1,1) appears to be the most appropriate specification. It contains the highest number of significant coefficients (4), the highest adjusted R^2 value of 0.997481, and relatively low volatility, with a sigma square of 2092.610. Moreover, it yields the lowest AIC (9074.565) and SIC (9098.379) values, indicating better model fit and parsimony compared to the others. The Durbin–Watson statistic of 2.001000 also supports the absence of serious autocorrelation.

Therefore, based on the diagnostic indicators, **ARIMA(3,1,1)** can be considered the best-fitted model among the proposed alternatives.

In addition to the manual evaluation of potential ARIMA models, the automated selection approach using the *auto.arima()* function in Rstudio also supports the same conclusion. Specifically, the function identifies **ARIMA(3,1,1) with drift** as the most appropriate model.

The estimated coefficients and their standard errors (s.e.) are as follows:

Coefficient	Estimate	s.e.
AR(1)	0.6794	0.1862
AR(2)	-0.1159	0.0492
AR(3)	0.1414	0.0376
MA(1)	-0.5242	0.1883
Drift	3.5928	2.4934

.

Table 5. ARIMA(3,1,1) with Drift

These coefficients suggest that the model captures moderate short-term dependencies in the time series. The AR(1) and MA(1) coefficients are particularly influential, indicating the presence of short memory and shock effects. The drift term (3.5928) implies a positive average change in the differenced series $dFOREX_t$, which may reflect a long-run upward trend in the original series $FOREX_t$

5.5. Diagnostic Tests of the Selected ARIMA Model

Table 6 presents the results of the ACF, PACF, and p-values obtained from the residual diagnostic Q-statistic test performed on the ARIMA(3,1,1) model, which has been selected based on the previous model evaluation.

Most ACF and PACF values of the residuals are close to zero. The Ljung–Box Q-statistic increases steadily, and the p-values from the Ljung–Box Q-test are mostly greater than 0.05 across all lags. This means there is no significant autocorrelation left in the residuals, indicating that the residuals are stationary. The ARIMA(3,1,1) model fits the data well and adequately captures the time series patterns. Overall, the residuals behave like random noise, confirming the model is appropriate for forecasting.

Table 6. Residual Diagnostic Results of ARIMA(3,1,1) Model

Lag	ACF	PACF	Q_stat	Prob
1	-0.0007	-0.0007	0.0004	
2	-0.0035	-0.0035	0.0112	
3	-0.0020	-0.0020	0.0146	
4	-0.0013	-0.0013	0.0160	0.000000
5	-0.0416	-0.0416	1.5286	0.216328
6	0.0462	0.0462	3.3956	0.183082
7	-0.0201	-0.0205	3.7490	0.289871
8	-0.0417	-0.0417	5.2732	0.260394
9	-0.0191	-0.0192	5.5933	0.347821
10	0.0878	0.0864	12.3701	0.054203
14	0.0034	0.0061	13.5120	0.196435
15	0.0224	0.0297	13.9553	0.235479
16	-0.0258	-0.0377	14.5418	0.267455
17	0.0077	0.0136	14.5946	0.333342
18	-0.0154	-0.0127	14.8036	0.391712
19	-0.0054	-0.0005	14.8295	0.463768
20	-0.0113	-0.0161	14.9430	0.528819
21	0.0345	0.0333	15.9999	0.523840
22	0.0686	0.0687	20.1891	0.322276
23	0.0113	0.0122	20.3031	0.376550
24	0.0165	0.0150	20.5466	0.424237
25	0.0048	-0.0006	20.5670	0.485646

Lag	ACF	PACF	Q_stat	Prob
11	-0.0140	-0.0110	12.5431	0.084053
12	0.0324	0.0290	13.4651	0.096817
13	-0.0064	-0.0081	13.5017	0.141189

To strengthen our conclusion, we also apply the ADF test to the residuals.

- Null hypothesis: H0: $\phi = 0$ → Residual has a unit root → Non-stationary residuals
- Alternative hypothesis: H1: $\phi < 0$ → Residual does not have a unit root → Stationary residuals

ADF test statistic	Critical values	Results' interpretation
	-3.43 at 1% level	Stationary
-29.398	-2.86 at 5% level	Stationary
	-2.57 at 10% level	Stationary

Table 7. ADF Test Results on Residuals of ARIMA(3,1,1)

The test yields a test statistic of –29.398, which is less than the 1% critical value of –3.43. Therefore, we reject the null hypothesis at the 1% level, concluding that the residuals are stationary, which means that this model is a suitable model for forecasting.

5.6. Forecast Accuracy Metrics

Coefficient	Estimate	s.e.
AR(1)	0.6794	0.1862
AR(2)	-0.1159	0.0492
AR(3)	0.1414	0.0376
MA(1)	-0.5242	0.1883
Drift	3.5928	2.4934
Model Fit: $\sigma^2 = 2090$ Log Likelihood = - 4531.28 AIC = 9074.57 AICc = 9074.67 BIC = 9103.15		
Training Error Metrics		
ME	0.0544	
RMSE	45.5605	
MAE	27.9664	

.

Training Error Metrics	
MPE	0.00010183
MAPE	0.1145799
MASE	1.0143
ACF1	0.0015

Table 8. Estimation and Forecast Accuracy Metrics for ARIMA(3,1,1)

These metrics show that the model performs well in terms of predictive accuracy on the training data.

- The Mean Error (ME) is 0.0544, which is close to zero, suggesting that the model does not systematically overestimate or underestimate the exchange rate. The Root Mean Squared Error (RMSE) and Mean Absolute Error (MAE) are 45.56 and 27.97 respectively, indicating the average deviation between predicted and actual values in absolute terms. These values, while not extremely low, are acceptable given the natural fluctuations of daily exchange rate data.
- Most notably, the Mean Absolute Percentage Error (MAPE) is only 0.115%, which reflects a very high level of forecast accuracy in percentage terms. This low MAPE suggests that the predicted values from the ARIMA(3,1,1) model deviate very little from the actual observed exchange rates, reflecting that the model effectively captures the short-term dynamics and trends of the USD/VND series.
- The Mean Absolute Scaled Error (MASE) is 1.014, which is approximately equal to 1. This indicates that the model's forecast performance is comparable to that of a simple benchmark (such as the naive forecast), though not substantially better.
- Finally, the autocorrelation of the residuals at lag 1 (ACF1) is very low (0.0015), implying that there is minimal autocorrelation remaining in the forecast errors. This is a desirable property and indicates that the residuals resemble white noise, which confirms that the ARIMA model has captured most of the autocorrelated structure in the data.

6. Conclusion

This paper demonstrates the effectiveness of the ARIMA model in forecasting the USD/VND exchange rate using daily data from 2022 to 2025. Among the candidate models evaluated, ARIMA(3,1,1) emerges as the most appropriate specification. The results indicate strong short-term predictive power, with low forecast errors across multiple metrics (e.g., RMSE: 45.56; MAE: 27.97; MAPE: 0.115%). This contrasts with some previous studies in other emerging markets, where lower-order ARIMA models such as ARIMA(1,1,1), ARIMA(2,1,1) or ARIMA(1,1,9) were found optimal (e.g., Ahmed & Keya, 2019; Rasheed, 2019).

The findings reinforce the applicability of ARIMA as a baseline short-term forecasting tool in Vietnam's foreign exchange market, especially for short-term projections. The USD/VND exchange rate exhibits relatively strong short-term fluctuations and is influenced by multiple factors, including monetary policy, inflation, trade balance, capital flows, and other macroeconomic conditions. The model can assist businesses and market participants by providing reliable short-term exchange rate forecasts, supporting better decision-making regarding foreign currency transactions, hedging, and risk management, thereby optimizing profits and minimizing exposure. For policymakers, short-term exchange rate forecasts help the central bank make timely decisions on interest rate adjustments, intervene in the foreign exchange market, or manage foreign reserves to stabilize the economy, while also identifying sudden or unexpected fluctuations to support rapid contingency planning and mitigate immediate macroeconomic risks.

However, as highlighted by numerous previous studies, ARIMA models exhibit certain limitations when applied to long-term forecasting due to their inability to capture structural changes, regime shifts, or nonlinear patterns inherent in financial time series. Future research may consider hybridizing ARIMA with other techniques, for example ARCH/GARCH models for volatility dynamics, or machine learning approaches to better capture nonlinearity and further improve forecast performance, as suggested by Tu et al. (2024) and Khashei et al. (2020).

REFERENCES

Adetunde, I.A. & Appiah, S.T. (2011), "Forecasting exchange rate between the Ghana Cedi and the US Dollar using time series analysis", *African Journal of Basic Applied Sciences*, Vol. 3, No. 6, pp. 255–264.

Ahmed, F. & Keya, J.A. (2019), "The time series analysis for predicting the exchange rate of USD to BDT", *International Journal of Academic Research in Business, Arts and Science*, Vol. 1, No. 2, pp. 282–294. DOI:10.5281/zenodo.3369751

Box, G.E.P. & Jenkins, G.M. (1970), *Time Series Analysis: Forecasting and Control*, Holden-Day, San Francisco.

Dunis, C. & Chen, Y. (2006), "Alternative volatility models for risk management and trading: Application to the EUR/USD and USD/JPY rates", *Derivatives Use, Trading & Regulation*, Vol. 11, No. 2, pp. 126–156. https://doi.org/10.1057/palgrave.dutr.1840013

Guha, B. & Bandyopadhyay, G. (2016), "Gold price forecasting using ARIMA model", *Journal of Advanced Management Science*, Vol. 4, No. 2, pp. 117–121. doi: 10.12720/joams.4.2.117-121

Khan, A.J. & Azim, P. (2013), "One-step-ahead forecastability of GARCH(1,1): A comparative analysis of USD- and PKR-based exchange rate volatilities", *The Lahore Journal of Economics*, Vol. 18, No. 1, pp. 1–38.

Khashei, M. & Mahdavi Sharif, B. (2020), "A Kalman filter-based hybridization model of statistical and intelligent approaches for exchange rate forecasting", *Journal of Modelling in Management*. https://doi.org/10.1108/JM2-12-2019-0277

Liu, Z. & Lv, Y. (2011), "A study of the USDX prediction based on ARIMA and GARCH models", in *Proceedings of the 2011 Fourth International Conference on Business Intelligence and Financial Engineering*, pp. 82–85. doi: 10.1109/BIFE.2011.9.

Matroushi, S. (2011), *Hybrid computational intelligence systems based on statistical and neural networks methods for time series forecasting: The case of the gold price*, Master's thesis, Lincoln University.

Meyler, A., Kenny, G. & Quinn, T. (1998), "Forecasting Irish inflation using ARIMA models", *Central Bank and Financial Services Authority of Ireland Working Paper*.

Naeem, S., Mashwani, W.K., Ali, A., et al. (2020), "Machine learning-based USD/PKR exchange rate forecasting using sentiment analysis of Twitter data", *Computers, Materials & Continua*, Vol. 67, No. 3, pp. 3451–3461. DOI:10.32604/cmc.2021.015872

Nguyen, T.L., Dang, D.H. & Vu, T.V. (2025), "Application of ARIMA model and deep learning in forecasting stock price in Vietnam", *Salud, Ciencia y Tecnología – Serie de Conferencias*, Vol. 4.

Nwankwo, S.C. (2014), "Autoregressive Integrated Moving Average (ARIMA) model for exchange rate (Naira to Dollar)", *Academic Journal of Interdisciplinary Studies*, Vol. 3, No. 4, pp. 429–434. DOI: 10.5901/ajis.2014.v3n4p429

Phung Duy, Q., Nguyen Thi, O., Le Thi, P.H., Pham Hoang, H.D., Luong, K.L. & Nguyen Thi, K.N. (2024), "Estimating and forecasting bitcoin daily prices using ARIMA-GARCH models", *Business Analyst Journal*, Vol. 45, No. 1, pp. 11–23. https://doi.org/10.1108/BAJ-05-2024-0027

Quinta, A., Pertiwi, A.G. & Widiyaningtyas, T. (2016), "Prediction of rupiah against US dollar by using ARIMA", in *Proceedings of the International Conference on Electrical Engineering, Computer Science and Informatics (EECSI)*. doi: 10.1109/EECSI.2017.8239205.

Rasheed, A., Ullah, M.A. & Uddin, I. (2019), "PKR exchange rate forecasting through univariate and multivariate time series techniques", *NICE Research Journal*, Vol. 13, No. 4, pp. 49–67. https://doi.org/10.51239/nrjss.v13i4.226

Tran, M.U. (2016), "Forecasting Foreign Exchange Rate by using ARIMA Model: A Case of VND/USD Exchange Rate*", Research Journal of Finance and Accounting*, Vol. 7, No. 12.

Tu, N.A., Khoa, D.D., Nhan, N.C.T. & Le Kieu Oanh, D. (2024), "Hybrid ARIMA and Machine Learning Approach for the VND/USD Exchange Rate Prediction in Vietnam: A Comparative Analysis", in *Partial Identification in Econometrics and Related Topics*, pp. 587–604, Cham: Springer Nature Switzerland. https://doi.org/10.1007/978-3-031-59110-5_39

Wang, S., Tang, Z. & Chai, B. (2016), "Exchange rate prediction model analysis based on improved artificial neural network algorithm", in *Proceedings of the 2016 International Conference on Communication and Electronics Systems (ICCES)*, pp. 1–5. doi: 10.1109/CESYS.2016.7889912

Yildiran, C.U. & Fettahoglu, A. (2017), "Forecasting USDTRY rate by ARIMA method", *Cogent Economics & Finance*, Vol. 5, No. 1, Article 1335968. https://doi.org/10.1080/23322039.2017.1335968

** The dataset used in this study was obtained from Yahoo Finance (https://finance.yahoo.com/), a publicly accessible financial data platform (Accessed: 30 July 2025).*

APPENDIX A: ARIMA FORMULAS

AR(p) - Autoregressive Process of Order p:

$$Y_t = \nu + a_1 Y_{t-1} + a_2 Y_{t-2} + \cdots + a_p Y_{t-p} + \varepsilon_t \tag{1}$$

Use the lag operator L:

$$A(L)Y_t = \nu + \varepsilon_t \text{ where } A(L) = 1 - a_1 L - a_2 L^2 - \cdots - a_p L^p \tag{2}$$

MA(q) - Moving Average Process of Order q:

$$Y_t = \nu + \varepsilon_t + b_1 \varepsilon_{t-1} + b_2 \varepsilon_{t-2} + \cdots + b_q \varepsilon_{t-q} \tag{3}$$

Use the lag operator L:

$$Y_t = \nu + B(L)\varepsilon_t \text{ where } B(L) = 1 + b_1 L + b_2 L^2 + \cdots + b_q L^q \tag{4}$$

ARMA(p, q) - Autoregressive Moving Average Process:

$$A(L)Y_t = \nu + B(L)\varepsilon_t \tag{5}$$

ARIMA(p, d, q) - Autoregressive Integrated Moving Average Process:

$$A(L)(1 - L)^d Y_t = \nu + B(L)\varepsilon_t \tag{6}$$