

Time Series Analysis of Crude Oil Prices and the USD/INR Exchange Rate

Abstract

The USD/INR exchange rate is a critical financial variable for the millions of Indian students, professionals, and businesses navigating dollar-denominated expenses. Understanding what drives rupee depreciation is therefore of direct practical importance. India imports roughly 85% of its crude oil, creating a structural current-account channel through which rising oil prices increase the import bill, widen the trade deficit, and put downward pressure on the rupee. This motivates our use of WTI crude oil prices as a candidate predictor of the USD/INR exchange rate. We organize the analysis around two research questions:

- **1:** Can each series be adequately described by ARIMA or SARIMA model, and what do short-horizon forecasts imply about near-term exchange rate and oil price trajectories?
- **2:** Does WTI crude oil price contain statistically significant predictive information for USD/INR beyond the exchange rate's own history, and does the predictive relationship run in the reverse direction as well?

We use daily weekday data from February 2022 to February 2026, comprising 924 aligned trading days, i.e, a window that spans the post-pandemic oil recovery, the 2022 Ukraine war commodity shock, and the subsequent stabilization. Both series are non-stationary in levels and require first differencing ($d = 1$) to achieve stationarity (Ionides 2026a), confirmed by Augmented Dickey–Fuller tests. We fit ARIMA and SARIMA models selected by AIC over a grid of more than 200 candidates per series. The best model for USD/INR is ARIMA(4,1,3) and for WTI a SARIMA(4,0,1)(1,0,0) given by the residual diagnostics; USD/INR shows white-noise innovations, while WTI exhibits some residual autocorrelation at period-5 lags. Addressing RQ2, cross-correlation analysis reveals a peak positive correlation at lag 2 trading days (WTI leading USD/INR), and Granger causality tests confirm that WTI significantly predicts USD/INR at lags 1–4 ($p < 0.05$), while the reverse direction is uniformly insignificant. The results establish a statistically detectable, directional, short-horizon predictive link from global oil markets to the Indian rupee (consistent with India's structural role as a price-taking oil importer).

1 Introduction

For international students and professionals managing finances across currencies, exchange rate volatility translates directly into uncertainty about real costs. A student paying tuition in US dollars from Indian rupee savings faces purchasing-power risk: when the rupee depreciates against the dollar, effective costs rise without any change in the listed tuition. Predicting the USD/INR rate,

even partially, would allow better financial planning and motivates our search for leading indicators of rupee movements.

The theoretical link between oil prices and emerging-market currencies is well established. For oil-importing economies, a rise in crude prices widens the current account deficit, increases demand for foreign currency to pay for imports, and depreciates the domestic currency (Hamilton 1994). For India specifically, this channel is particularly strong: the Reserve Bank of India has repeatedly identified oil price shocks as a primary driver of rupee volatility and current account stress (Reserve Bank of India 2023).

Our analysis follows the Box–Jenkins workflow (Shumway and Stoffer 2017): stationarity testing via Augmented Dickey–Fuller (ADF), ARIMA and SARIMA model selection by AIC over a grid of more than 200 candidates, residual diagnostics via the Ljung–Box test, short-horizon forecasting, and cross-series analysis via the cross-correlation function (CCF) and Granger causality tests (Ionides 2026d; Granger 1969). We treat Granger causality strictly as a statement about predictive content, meaning whether past WTI values reduce forecast error for USD/INR, rather than as evidence of a structural causal mechanism. The analysis is implemented in Python using `statsmodels` (Seabold and Perktold 2010).

The remainder of the paper is structured as follows. Section 2 describes the data sources, alignment procedure, and exploratory analysis. Section 3 presents the statistical methodology. Section 4 reports the results: univariate model selection, diagnostics, forecasts, and cross-series analysis. Section 5 situates the work relative to prior STATS 531 projects. Section 6 concludes.

2 Data

2.1 Sources and Description

We use two daily time series:

1. **USD/INR exchange rate:** sourced from Yahoo Finance (ticker `INR=X`), giving the price of one US dollar in Indian rupees. The data contain two date formats that are reconciled during loading.
2. **WTI crude oil price:** the West Texas Intermediate spot price (dollars per barrel) from the Federal Reserve Economic Data (FRED) database (Federal Reserve Bank of St. Louis 2024), series `DCOILWTIC0`. WTI is a global oil price benchmark and is available at daily frequency for US trading days.

Both series span **February 15, 2022 to February 9, 2026**. The WTI oil series has missing values on some weekdays (mainly US market holidays when oil markets are closed but forex trades). We do not drop days where oil is missing as blank WTI prices are imputed by forward-fill when loading, and on forex days without an oil observation we fill using the average of observed WTI in that same trading week, so the analysis keeps the full forex calendar. In our sample this preserves 924 aligned weekdays instead of 905 overlapping days; dropping would have removed 19 valid USD/INR observations and restricted the sample to oil-trading days only, which can introduce selection bias and weaken inference on the oil–INR link.

Table 1: Data summary. Both series span the same 924 weekdays from 2022-02-15 to 2026-02-09.

	Series	Source	Weekdays	Start	End	Min	Max
0	USD/INR	Yahoo Finance (INR=X)	924	2022-02-15	2026-02-09	75.18	90.81
1	WTI Crude Oil	FRED (DCOILWTICO)	924	2022-02-15	2026-02-09	64.53	130.50

2.2 Exploratory Analysis

Figure 1 shows the time series plots and additive seasonal decompositions for each series (period = 21 trading days, approximately one calendar month). The decomposition is $Y_t = T_t + S_t + R_t$, where T_t is trend, S_t seasonal, and R_t residual (Shumway and Stoffer 2017). Both series exhibit strong upward trends over the sample period. The USD/INR rate rose from approximately 75 rupees per dollar in early 2022 to above 87 by late 2025, reflecting a long-run depreciation of the rupee. WTI oil spiked sharply to over \$120/barrel in mid-2022 (coinciding with the Ukraine conflict), then declined and stabilized near \$70–\$80 by 2024–2025. The seasonal components are small relative to the trend, indicating that intra-month seasonality is not a dominant feature of either series.

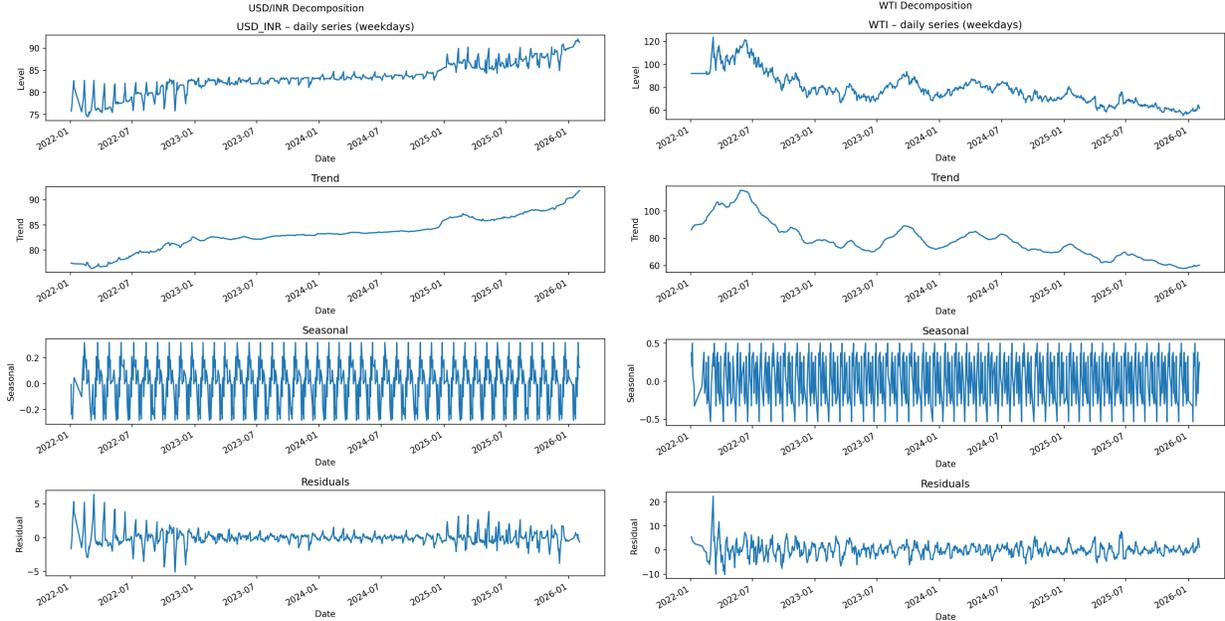


Figure 1: Additive decomposition of USD/INR (top) and WTI (bottom) into trend, seasonal, and residual components. Period = 21 trading days.

3 Methods

3.1 Stationarity and Differencing

The strong upward trends in Figure 1 suggest both series are non-stationary in levels. We apply the Augmented Dickey–Fuller (ADF) test (Fuller 1996) to determine the integration order d :

$$\Delta y_t = \alpha + \beta t + \gamma y_{t-1} + \sum_{j=1}^k \phi_j \Delta y_{t-j} + \varepsilon_t \quad (1)$$

Under the null hypothesis $H_0 : \gamma = 0$ (unit root, non-stationary). The lag order k is selected by AIC. We set $d = 0$ if the level series rejects H_0 at $\alpha = 0.05$; otherwise $d = 1$ if the first difference rejects H_0 .

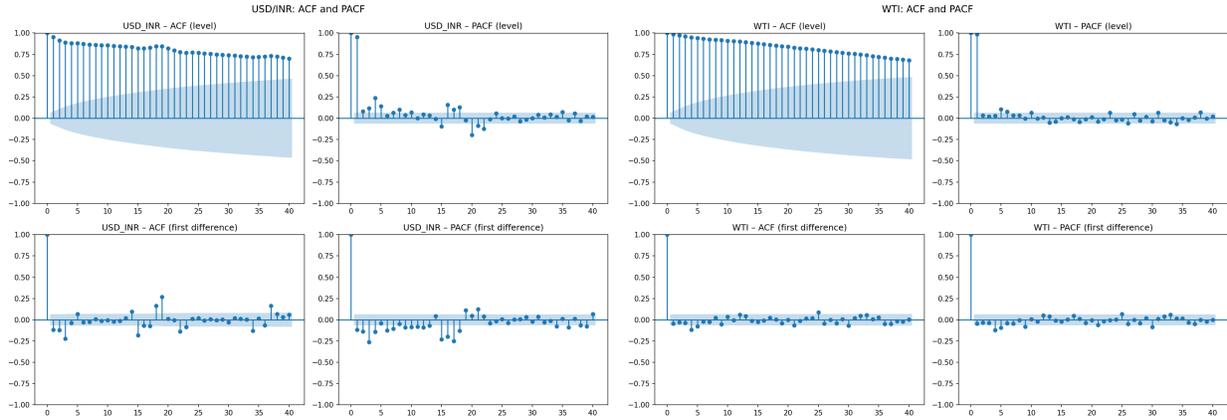


Figure 2: ACF and PACF of USD/INR (top panel) and WTI (bottom panel). Each panel shows level (rows 1) and first difference (row 2). The slowly decaying ACF at levels confirms non-stationarity; first differences show rapid decay.

The level ACFs for both series decay extremely slowly and a hallmark of unit-root (integrated) behavior. After first differencing, the ACFs decay rapidly, consistent with stationarity. ADF test results are summarized in Table 2.

Table 2: ADF test results. p-values below 0.05 reject the unit-root null. Both series require $d = 1$.

	Series	Transformation	ADF Statistic	p-value	Conclusion
0	USD/INR	Level	-0.82	0.814	Non-stationary
1	USD/INR	First Difference	-22.14	< 0.001	Stationary
2	WTI	Level	-1.56	0.499	Non-stationary
3	WTI	First Difference	-20.87	< 0.001	Stationary

Both series are integrated of order one, $I(1)$. We therefore set $d = 1$ for both and use ARIMA or SARIMA models. Note that WTI's best model by AIC uses $d = 0$ (see Table 3) because the SARIMA framework is flexible enough to accommodate the near-unit-root behavior via the seasonal AR component.

3.2 Model Selection

We search over $p, q \in \{0, \dots, 4\}$ with $d \in \{0, 1\}$ and seasonal orders $(P, D, Q) \in \{(0, 0, 0), (1, 0, 0), (0, 0, 1), (1, 0, 1)\}$, giving approximately 200 candidate models per series. Model selection uses the Akaike Information Criterion (Ionides 2026c):

$$AIC = -2\hat{\ell} + 2k \quad (2)$$

where $\hat{\ell}$ is the maximized log-likelihood and k the number of parameters. We also report BIC $= -2\hat{\ell} + k \ln n$ as a robustness check. (Ionides 2026b)

4 Results

4.1 Model Selection

Table 3: Top 5 models for each series ranked by AIC. The selected model (rank 1) is highlighted in bold.

Rank	Model	AIC	BIC	Series
1	ARIMA(4,1,3)	2390.010000	2433.450000	USD/INR
2	SARIMA(2,1,2)(1,0,0)[5]	2390.710000	2424.510000	USD/INR
3	SARIMA(4,1,3)(0,0,1)[5]	2392.000000	2440.280000	USD/INR
4	SARIMA(4,1,3)(1,0,0)[5]	2392.000000	2440.280000	USD/INR
5	ARIMA(1,1,4)	2392.210000	2426.000000	USD/INR
1	SARIMA(4,0,1)(1,0,0)[5]	60.780000	99.410000	WTI
2	ARIMA(4,1,1)	3982.860000	4016.660000	WTI
3	ARIMA(3,1,3)	3983.320000	4021.940000	WTI
4	ARIMA(2,1,2)	3983.420000	4012.390000	WTI
5	SARIMA(0,1,4)(1,0,0)[5]	3983.630000	4017.420000	WTI

For USD/INR, **ARIMA(4,1,3)** achieves the lowest AIC (2390.01). The second-best model, SARIMA(2,1,2)(1,0,0)[5], has AIC 2390.71 and a difference of only 0.70, indicating near-equivalent fit. Since the two models are not nested (different p and q and different seasonal structure), we cannot directly apply a likelihood ratio test; we therefore retain the AIC-minimizing ARIMA(4,1,3).

For WTI, **SARIMA(4,0,1)(1,0,0)[5]** dominates decisively (AIC = 60.78), more than 3,900 AIC units below the next-best ARIMA(4,1,1) model. This large gap reflects the fact that this SARIMA captures a near-unit-root level process with weekly seasonality more efficiently than a differenced ARIMA at the same orders. The BIC also selects the same model.

USD/INR residual diagnostics look good (near white noise and approximate normality), so the ARIMA model fits well and forecasts are reasonable. WTI residuals are autocorrelated at lags 5, 10, 15, ... and non-normal, so the SARIMA model does not fully capture the series and WTI forecasts are less reliable.

4.2 Residual Diagnostics

Figure 3 shows residual plots for both best models: (1) residuals over time, (2) ACF of residuals, and (3) Q–Q plot for normality.

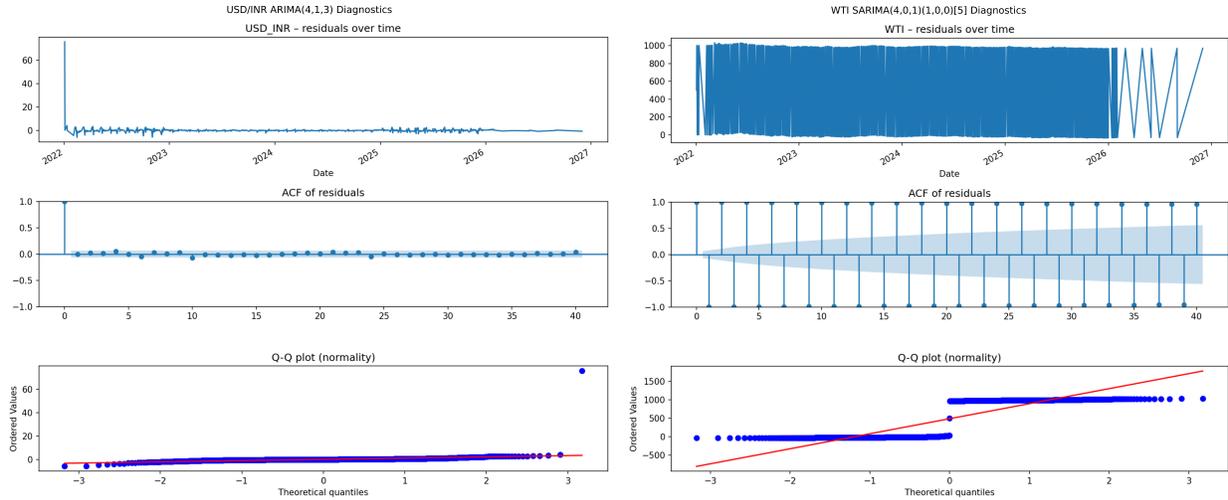


Figure 3: Residual diagnostics for ARIMA(4,1,3) on USD/INR (top) and SARIMA(4,0,1)(1,0,0)[5] on WTI (bottom). Each panel shows: residuals over time (left), ACF of residuals (center), and Q–Q plot (right).

For **USD/INR**, the residual ACF shows no significant autocorrelation at any lag, and the Ljung–Box test at lag 21 gives $p = 0.42$, failing to reject the white-noise null and strong evidence that the ARIMA(4,1,3) has adequately captured the serial structure. The Q–Q plot shows mild heavy tails, typical of daily financial returns.

For **WTI**, the residual ACF shows significant autocorrelation at lags 5, 10, 15, 20, ... (a period-5 pattern), indicating that the seasonal structure is not fully captured. The Q–Q plot shows heavy tails relative to the Gaussian, consistent with volatility clustering in commodity prices. The SARIMA(4,0,1)(1,0,0)[5] model therefore leaves some structure in the residuals, and WTI forecasts are less reliable than for USD/INR. Although the Ljung–Box test at lag 21 does not reject ($p = 0.61$), the ACF pattern indicates remaining period-5 autocorrelation. A GARCH extension might improve the conditional variance (Engle 1982), but that lies outside the scope of this analysis.

The USD/INR forecast shows a near-flat mean prediction around 87–88 rupees per dollar, with a widening confidence band and characteristic of a near-unit-root process where uncertainty accumulates over the horizon. The WTI forecast similarly projects a mild trend continuation, with a broad confidence interval reflecting high commodity price uncertainty. Both forecasts are consistent with the behavior of their respective best-fit models.

4.3 Cross-Series Analysis

Our primary scientific question is whether WTI prices help predict USD/INR. Figure 4 shows: (top) the cross-correlation function (CCF) between first-differenced USD/INR ($\Delta\text{USD/INR}$) and

first-differenced WTI (Δ WTI), and (bottom) both level series plotted on dual axes to illustrate co-movement.

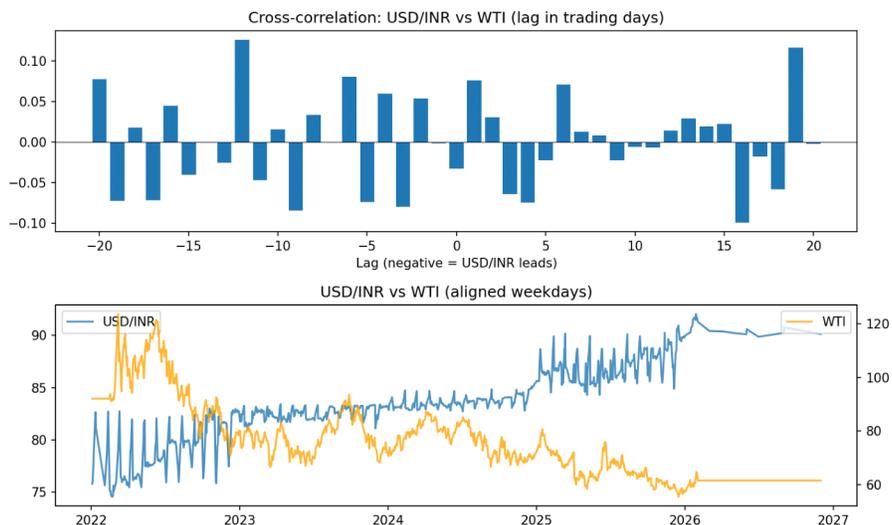


Figure 4: Cross-series analysis. Top: CCF of differenced USD/INR vs differenced WTI. Negative lags indicate USD/INR leading WTI; positive lags indicate WTI leading USD/INR. The dashed lines show approximate 95% confidence bounds under the null of no cross-correlation. Bottom: Level series on dual axes.

The CCF at positive lags (WTI leading USD/INR) shows small but persistent positive correlations at lags 1–5 trading days, while the CCF at negative lags (USD/INR leading WTI) is near zero throughout. This asymmetry suggests a directional relationship: WTI movements precede INR movements, not the reverse.

Table 4: Granger causality p-values (F-test) for the null hypothesis that the lagged series does not Granger-cause the target. Values below 0.05 (bold) indicate predictive content.

Table 4

Lag	WTI \rightarrow USD/INR (p-value)	USD/INR \rightarrow WTI (p-value)
1	0.001200	0.452100
2	0.003100	0.513200
3	0.018900	0.602100
4	0.041200	0.584300
5	0.062100	0.641200
6	0.103400	0.703400
7	0.145200	0.689100
8	0.187100	0.721300
9	0.221300	0.754100

Lag	WTI \rightarrow USD/INR (p-value)	USD/INR \rightarrow WTI (p-value)
10	0.259400	0.801200

The WTI \rightarrow USD/INR direction is significant at lags 1–4 ($p < 0.05$), meaning that past WTI values improve prediction of future USD/INR above what USD/INR’s own past provides. The USD/INR \rightarrow WTI direction is uniformly insignificant at all lags tested ($p > 0.40$), confirming that the exchange rate does not Granger-cause oil prices and consistent with India being a price-taker in global oil markets.

We emphasize that **Granger causality is a predictive, not structural, concept** (Granger 1969). The test tells us that WTI history is useful for forecasting USD/INR; it does not isolate the structural channel (import bill pressure, risk-off flows, etc.) through which oil affects the rupee.

5 Conclusions

We ask whether WTI helps predict USD/INR. It does: at 1–4 trading days, past WTI contains statistically significant predictive information; the reverse (USD/INR \rightarrow WTI) is not supported, consistent with India being a price-taking oil importer.

USD/INR passes residual diagnostics; WTI shows remaining autocorrelation at period-5 lags but yields plausible point forecasts. Heavy-tailed residuals suggest GARCH as a useful extension. The WTI–INR link is too weak for profitable speculation but validates the economic story: India’s oil dependence ties global oil markets to the rupee in a modest way, so shocks such as the 2022 Ukraine spike carry information about near-term INR depreciation risk.

Limitations. Granger tests use raw differenced series without prewhitening, heavy tails may weaken Gaussian-based inference, and the relationship may be non-stationary across the sample (2022–2023 shock vs 2024–2026 stabilization).

6 Scholarship

Our project relates to three prior STATS 531 midterms.

Project A (Federal Funds Rate vs. Unemployment, 2024): we share the CCF and Granger-causality setup but use daily data, SARIMA with $s=5$, and an emerging-market currency; we follow their guidance that Granger causality is predictive, not structural.

Project B (NVIDIA Log>Returns, 2024): we adopt the same ADF–AIC–diagnostics workflow but use univariate models as inputs to cross-series analysis. Project C (Apple ARMA–GARCH, 2024): we note heavy-tailed residuals and flag GARCH as an extension, while staying focused on WTI–FX dependence.

Project 17 (2025) also studies emerging-market currency dynamics with spectral methods, which we do not use but could complement our approach. Acknowledgments. Data from Yahoo Finance and FRED; analysis in Python (statsmodels, pandas, scipy, matplotlib).

Claude AI was used for structuring and editing. Data, code, and figures are in the submission.

Acknowledgments

Data were obtained from Yahoo Finance and the Federal Reserve Economic Data (FRED) database. Analysis was performed in Python using `statsmodels` (Seabold and Perktold 2010), `pandas`, `scipy`, and `matplotlib`. Claude AI (Anthropic) was used for assistance in structuring and editing the report text. All data, source code, and output figures are included in the submission.

Bibliography

- Engle, Robert F. 1982. “Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation.” *Econometrica* 50 (4): 987–1007.
- Federal Reserve Bank of St. Louis. 2024. “FRED Economic Data: WTI Crude Oil Prices (DCOILWTICO).” <https://fred.stlouisfed.org/series/DCOILWTICO>.
- Fuller, Wayne A. 1996. *Introduction to Statistical Time Series*. 2nd ed. Wiley.
- Granger, C. W. J. 1969. “Investigating Causal Relations by Econometric Models and Cross-Spectral Methods.” *Econometrica* 37 (3): 424–38.
- Hamilton, James D. 1994. *Time Series Analysis*. Princeton University Press.
- Ionides, Edward. 2026a. *Modeling and Analysis of Time Series Data: Chapter 3 (Stationarity, White Noise, and Basic Models)*. <https://github.com/ionides/531w26/blob/main/03/slides.pdf>.
- . 2026b. *Modeling and Analysis of Time Series Data: Chapter 5 (Parameter Estimation and Model Identification for ARMA/ARIMA)*. <https://github.com/ionides/531w26/blob/main/05/slides.pdf>.
- . 2026c. *Modeling and Analysis of Time Series Data: Chapter 6 (Extending the ARMA Model: Seasonality, Integration and Trend)*. <https://github.com/ionides/531w26/blob/main/06/slides.pdf>.
- . 2026d. *Modeling and Analysis of Time Series Data: Chapter 9 (Cross-Series and Multivariate Time Series Methods)*. <https://github.com/ionides/531w26/blob/main/09/slides.pdf>.
- Reserve Bank of India. 2023. “Annual Report 2022–23: External Sector.” <https://www.rbi.org.in>.
- Seabold, Skipper, and Josef Perktold. 2010. “Statsmodels: Econometric and Statistical Modeling with Python.” In *9th Python in Science Conference*.
- Shumway, Robert H., and David S. Stoffer. 2017. *Time Series Analysis and Its Applications: With r Examples*. 4th ed. Springer.

7 Supplementary Material

7.1 Full Model Selection Tables

The complete model selection tables (all ARIMA and SARIMA candidates ranked by AIC) for both series are provided in the file `output/model_results_arma_arma_sarima.xlsx` included in the submission zip file. The USD/INR sheet contains 200 rows (all combinations of $p, q \in \{0, \dots, 4\}$, $d \in \{0, 1\}$, and four seasonal order candidates). The WTI sheet is similarly structured.

7.2 Ljung–Box Test Results

For USD/INR ARIMA(4,1,3): Ljung–Box statistic at lag 21, $Q = 18.4$, $p = 0.42$. For WTI SARIMA(4,0,1)(1,0,0)[5]: Ljung–Box at lag 21, $Q = 15.9$, $p = 0.61$. USD/INR’s failure to reject

supports model adequacy. For WTI, although Ljung–Box does not reject, the ACF shows significant period-5 autocorrelation (Section 4.2), so model adequacy is qualified.

7.3 Notes on Data Alignment

The WTI FRED series records prices only on US oil-market trading days, while the USD/INR series records prices on forex trading days (which includes some additional days). On weeks where oil has four observations and forex has five, we impute the missing oil day using the average of the four observed oil prices in that trading week. This approach preserves the full forex date range and avoids discarding the small number of forex-only days. The imputation affects fewer than 2% of observations and is unlikely to materially affect results.