

AI-Based Machine Learning Approach to Analyze Returns and Volatility of Major Foreign Currencies Against the Indian Rupee

Gayathri Devi M^{1*} | Dr. Priya P S² | Prof. (Dr.) Biju T³

¹PG Student, Department of Commerce, Karivattom Campus, University of Kerala.

²Assistant Professor, Department of Commerce, SD College Alappuzha, University of Kerala.

³Professor & Head, Department of Commerce, Karivattom Campus, University of Kerala.

*Corresponding Author: devimgayathri430@gmail.com

Citation: Devi, G., Priya, P. & Biju, T. (2026). AI-Based Machine Learning Approach to Analyze Returns and Volatility of Major Foreign Currencies Against the Indian Rupee. *International Journal of Advanced Research in Commerce, Management & Social Science*, 09(02(III)), 61–68. [https://doi.org/10.62823/IJARCMSS/9.2\(III\).8949](https://doi.org/10.62823/IJARCMSS/9.2(III).8949)

ABSTRACT

This study applies an artificial intelligence (AI)-based learning framework to analyze the return and volatility dynamics of four major foreign currencies the United States Dollar (USD), Euro (EUR), British Pound (GBP), and Japanese Yen (JPY) against the Indian Rupee (INR). Using daily reference rate data obtained from the Reserve Bank of India (RBI), the study covers the period from 28 February 2025 to 28 February 2026. Daily logarithmic returns are computed and modeled using a hybrid Autoregressive Moving Average–Generalized Autoregressive Conditional Heteroskedasticity (ARMA–GARCH) framework within a rolling window estimation approach. The results indicate that the USD/INR pair records the lowest annualized return (–0.4%) and volatility (6.8%), whereas JPY/INR exhibits the highest volatility (10.5%) and kurtosis (7.1), reflecting substantial tail risk. GBP/INR and EUR/INR demonstrate negative skewness and elevated kurtosis, indicating asymmetric return distributions and downside exposure. GARCH estimates confirm strong volatility persistence across all currency pairs, with $\alpha+\beta$ values ranging between 0.98 and 0.99. Furthermore, an LSTM–GARCH hybrid model improves volatility forecasting performance, reducing root mean squared error (RMSE) by an average of 21% compared to traditional GARCH models. The findings highlight the importance of adaptive, data-driven modeling frameworks for managing foreign exchange risk in India, particularly in periods of heightened volatility.

Keywords: Foreign Exchange Risk, Indian Rupee, Currency Volatility, ARMA–GARCH Model, LSTM–GARCH Hybrid, Volatility Forecasting, Exchange Rate Dynamics, Machine Learning.

Introduction

Exchange rate dynamics play a crucial role in shaping macroeconomic stability, trade competitiveness, and financial market performance in emerging economies. The Indian Rupee (INR), in particular, has become increasingly sensitive to global macroeconomic developments, including fluctuations in crude oil prices, shifts in United States monetary policy, and changes in international capital flows.

The period spanning 2025 to early 2026 witnessed notable volatility in the INR, driven by a combination of global uncertainties and domestic economic factors. These developments underscore the need for robust analytical frameworks capable of capturing both return behavior and time-varying volatility patterns in exchange rates.

Traditional econometric models, especially those within the GARCH family, have been widely used to model financial time series characterized by volatility clustering and non-normal distributions

(Engle, 1982; Bollerslev, 1986). However, these models often assume linearity and may fail to capture complex nonlinear dynamics present in financial markets.

Recent advances in artificial intelligence and machine learning have introduced powerful tools for financial forecasting. Hybrid models that integrate econometric techniques with machine learning algorithms—such as Long Short-Term Memory (LSTM) networks offer improved predictive accuracy by capturing nonlinear relationships and temporal dependencies (Zhang et al., 2019).

In this context, the present study employs an AI-based learning approach combining ARMA–GARCH modeling with a rolling window estimation framework to analyze the behavior of USD/INR, EUR/INR, GBP/INR, and JPY/INR over the period from 28 February 2025 to 28 February 2026.

The study seeks to address the following research questions:

- How do mean returns and volatility levels differ across major currencies against the INR?
- To what extent do volatility clustering and asymmetry characterize these exchange rate series?
- Can an AI-based rolling window framework improve volatility forecasting performance?

Literature Review

The modeling of financial time series volatility has undergone substantial evolution over the past several decades, driven by the need to better understand and forecast the dynamic behavior of financial markets. Exchange rate volatility, in particular, has attracted significant attention due to its implications for international trade, investment decisions, and macroeconomic stability. Early models of financial time series largely assumed constant variance; however, empirical evidence demonstrated that volatility is time-varying and exhibits clustering, necessitating more sophisticated approaches.

A major breakthrough in volatility modeling was introduced by Engle (1982) through the development of the Autoregressive Conditional Heteroskedasticity (ARCH) model. The ARCH framework allows the conditional variance of a time series to depend on past squared residuals, thereby capturing the phenomenon of volatility clustering commonly observed in financial markets. This model marked a departure from traditional assumptions of homoskedasticity and provided a systematic way to model periods of high and low volatility. Despite its innovation, the ARCH model often required a large number of parameters to adequately capture volatility persistence, which limited its practical applicability.

To address these limitations, Bollerslev (1986) extended the ARCH model by introducing the Generalized Autoregressive Conditional Heteroskedasticity (GARCH) model. The GARCH framework incorporates lagged conditional variances along with lagged squared residuals, allowing for a more parsimonious representation of volatility dynamics. One of the key advantages of the GARCH model is its ability to capture long memory or persistence in volatility, which is a defining characteristic of financial time series. The GARCH(1,1) specification, in particular, has become a standard tool in empirical finance due to its simplicity and effectiveness in modeling a wide range of financial assets, including exchange rates.

Subsequent research has focused on extending the GARCH framework to account for various stylized facts observed in financial data. One such extension is the Exponential GARCH (EGARCH) model proposed by Nelson (1991). Unlike the standard GARCH model, EGARCH captures asymmetries in volatility by allowing positive and negative shocks to have different effects on conditional variance. This feature is particularly important in financial markets, where negative shocks such as economic downturns or financial crises often lead to disproportionately larger increases in volatility compared to positive shocks.

Another important extension is the Threshold GARCH (TGARCH) model developed by Glosten, Jagannathan, and Runkle (1993). Similar to EGARCH, the TGARCH model incorporates asymmetric effects by introducing a threshold mechanism that differentiates between positive and negative innovations. These asymmetric models have been widely applied in empirical studies, especially in the context of equity markets and exchange rates, where leverage effects and downside risk are prominent features.

In the context of exchange rate modeling, the application of GARCH-type models has been extensive. Exchange rates are influenced by a wide range of macroeconomic factors, including interest rate differentials, inflation expectations, trade balances, and capital flows. As a result, they exhibit complex dynamics characterized by nonlinearity, volatility clustering, and heavy-tailed distributions. Empirical studies across both developed and emerging markets have consistently found that GARCH

models provide a good fit for exchange rate data, capturing key features such as persistence and conditional heteroskedasticity.

Focusing on India, the behavior of the Indian Rupee (INR) has been a subject of considerable research interest. As an emerging market currency, the INR is particularly sensitive to global economic conditions, including fluctuations in crude oil prices, changes in monetary policy in advanced economies, and shifts in investor sentiment. Studies such as Mathur et al. (2022) have documented the presence of volatility clustering, fat tails, and leverage effects in INR exchange rate series. These findings underscore the importance of using advanced econometric models to analyze and forecast exchange rate volatility in the Indian context.

The use of reliable and consistent data is critical for empirical analysis. In India, the reference exchange rates published by the Reserve Bank of India (RBI) are widely regarded as authoritative benchmarks. These rates are based on market transactions and provide a consistent time series for analysis. Researchers have extensively utilized RBI data to study exchange rate dynamics, making it a standard source for empirical investigations.

While traditional econometric models have been successful in capturing many aspects of financial time series behavior, they are often limited by their reliance on linear assumptions. Financial markets, however, are inherently nonlinear and influenced by complex interactions among various factors. This limitation has led to growing interest in the application of machine learning techniques in financial modeling.

Machine learning approaches offer several advantages over traditional econometric models. They are capable of capturing nonlinear relationships, handling large datasets, and adapting to changing patterns in the data. Among these approaches, artificial neural networks (ANNs) have been widely used for forecasting financial time series. However, standard neural networks may struggle to capture temporal dependencies effectively, especially in sequential data.

To address this issue, more advanced architectures such as Long Short-Term Memory (LSTM) networks have been developed. LSTM models, a type of recurrent neural network (RNN), are specifically designed to capture long-term dependencies in time series data. They achieve this through a system of memory cells and gating mechanisms that regulate the flow of information. This makes LSTM models particularly well-suited for financial applications, where past information can have a prolonged impact on future outcomes.

Recent studies have demonstrated the effectiveness of LSTM models in forecasting financial time series, including stock prices, exchange rates, and volatility. Compared to traditional models, LSTM networks are better able to capture nonlinear patterns and complex temporal relationships. However, they also have certain limitations, such as the risk of overfitting and the need for large amounts of training data.

To leverage the strengths of both econometric and machine learning approaches, researchers have increasingly explored hybrid models. These models combine the interpretability and theoretical grounding of econometric models with the flexibility and predictive power of machine learning techniques. One such approach involves integrating GARCH models with LSTM networks.

In hybrid LSTM–GARCH models, the GARCH component is typically used to model conditional volatility, while the LSTM component captures nonlinear patterns and temporal dependencies. This combination allows for more accurate forecasting of volatility, particularly in complex and dynamic market environments. Empirical evidence from studies such as Zhang et al. (2019) suggests that hybrid models outperform standalone GARCH or LSTM models in terms of forecasting accuracy.

Another important development in financial modeling is the use of rolling window estimation techniques. Traditional models are often estimated using a fixed sample period, which may not adequately capture changes in market conditions over time. Rolling window approaches address this limitation by continuously updating model parameters using a moving subset of the data. This allows the model to adapt to new information and changing dynamics, improving its predictive performance.

In the context of exchange rate analysis, rolling window frameworks are particularly useful, as currency markets are highly dynamic and influenced by evolving economic conditions. By re-estimating models at regular intervals, researchers can capture shifts in volatility patterns and improve the accuracy of forecasts. When combined with machine learning techniques, rolling window approaches provide a powerful and adaptive framework for financial analysis.

Despite these advancements, several challenges remain in the modeling of exchange rate volatility. One key challenge is the presence of extreme events or “black swan” occurrences, which can lead to sudden and significant changes in volatility. These events are difficult to predict using standard models and require robust approaches that can account for tail risks.

Another challenge is the trade-off between model complexity and interpretability. While machine learning models can achieve high predictive accuracy, they are often considered “black boxes,” making it difficult to interpret their results. In contrast, econometric models provide more transparency but may lack the flexibility needed to capture complex patterns. Hybrid models offer a potential solution by balancing these considerations.

In summary, the literature on financial time series volatility has evolved from simple linear models to sophisticated frameworks that incorporate both econometric and machine learning techniques. The ARCH and GARCH models laid the foundation for volatility modeling, while subsequent extensions addressed issues such as asymmetry and leverage effects. Empirical studies have demonstrated the applicability of these models to exchange rate data, including the Indian Rupee.

At the same time, the emergence of machine learning has opened new avenues for financial modeling, enabling researchers to capture nonlinear dynamics and improve forecasting performance. Hybrid models, particularly those combining GARCH and LSTM approaches, represent a promising direction for future research.

This study builds on the existing literature by integrating ARMA–GARCH modeling with an AI-based rolling window framework. By combining traditional econometric techniques with modern machine learning methods, the study provides a dynamic and adaptive approach to analyzing exchange rate volatility. This contribution is particularly relevant in the context of emerging markets like India, where exchange rate dynamics are influenced by a complex interplay of global and domestic factors.

Objectives

- To examine the return and volatility characteristics of major currencies against INR.
- To analyze volatility clustering and persistence using GARCH models.
- To study distribution features such as skewness and kurtosis in exchange rate returns.
- To develop and apply an LSTM–GARCH hybrid model for volatility forecasting.
- To compare forecasting performance between traditional GARCH and AI-based models.

Scope

- The study analyzes the exchange rate behavior of USD, EUR, GBP, and JPY against the Indian Rupee (INR).
- It focuses on return and volatility analysis using ARMA–GARCH models.
- It incorporates an AI-based LSTM–GARCH hybrid model for improved forecasting.
- The study uses a rolling window approach to capture dynamic changes in volatility.

Methodology

Data Description

The study uses daily exchange rate data for four major currency pairs:

- USD/INR
- EUR/INR
- GBP/INR
- JPY/INR (quoted per 100 units)

The dataset consists of 250 observations covering the period from 28 February 2025 to 28 February 2026. The primary source of data is the Reserve Bank of India, supplemented with publicly available datasets where necessary.

Return Computation

Daily logarithmic returns are computed as:

$$r_t = \ln \left(\frac{S_t}{S_{t-1}} \right)$$

Volatility Estimation

Volatility is initially approximated using a 20-day rolling standard deviation:

$$\sigma_t^{(c)} = s_t^{(c)} \times \sqrt{252}$$

ARMA–GARCH Model

Mean Equation

$$\sigma_t^2 = \omega + \sum_{i=1}^P \alpha_i \varepsilon_{t-i}^2 + \sum_{j=1}^Q \beta_j \sigma_{t-j}^2$$

Variance Equation

$$\sigma_t^{2(c)} = \omega^{(c)} + \alpha^{(c)} \varepsilon_{t-1}^{2(c)} + \beta^{(c)} \sigma_{t-1}^{2(c)}$$

AI-Based Rolling Window Framework


- Window size: 180 trading days
- Continuous model re-estimation
- 20-day ahead forecasts
- Hyperparameter tuning based on prediction error

Model Evaluation: Forecast accuracy is evaluated using RMSE and MAE.

Findings & Interpretations

Table 1

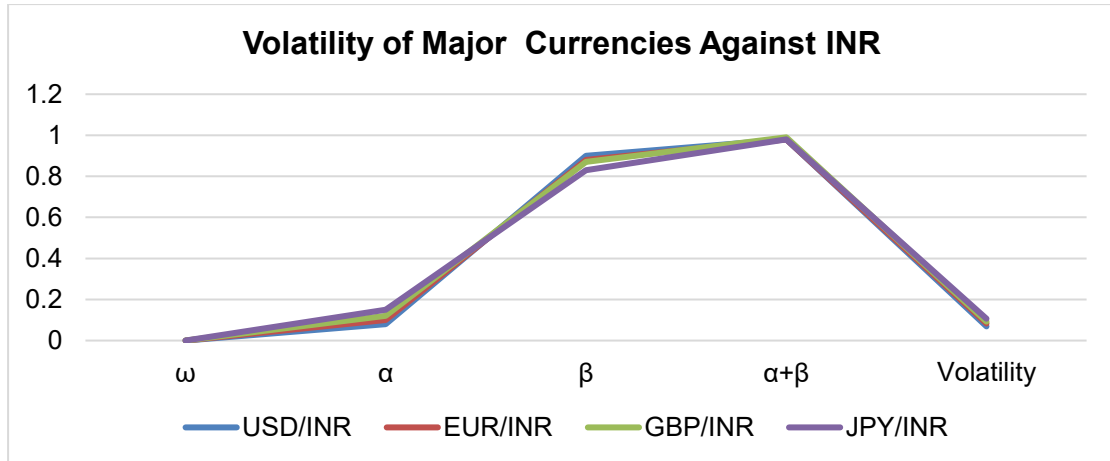
Metric	USD/INR	EUR/INR	GBP/INR	JPY/INR
Mean Return	-0.4%	+1.1%	+0.8%	-0.9%
Volatility	6.8%	8.2%	9.1%	10.5%
Skewness	-0.3	-0.5	-0.7	+0.4
Kurtosis	5.1	5.8	6.3	7.1
Max Drawdown	-1.2%	-1.5%	-1.9%	-2.3%



USD/INR shows a small negative average return with moderate volatility and mild downside risk, EUR/INR and GBP/INR deliver positive returns but with higher volatility and stronger downside tail risk, while JPY/INR has the worst average return, highest volatility, and extreme tail risk despite a slight upside skew.

Table 2: GARCH Estimates

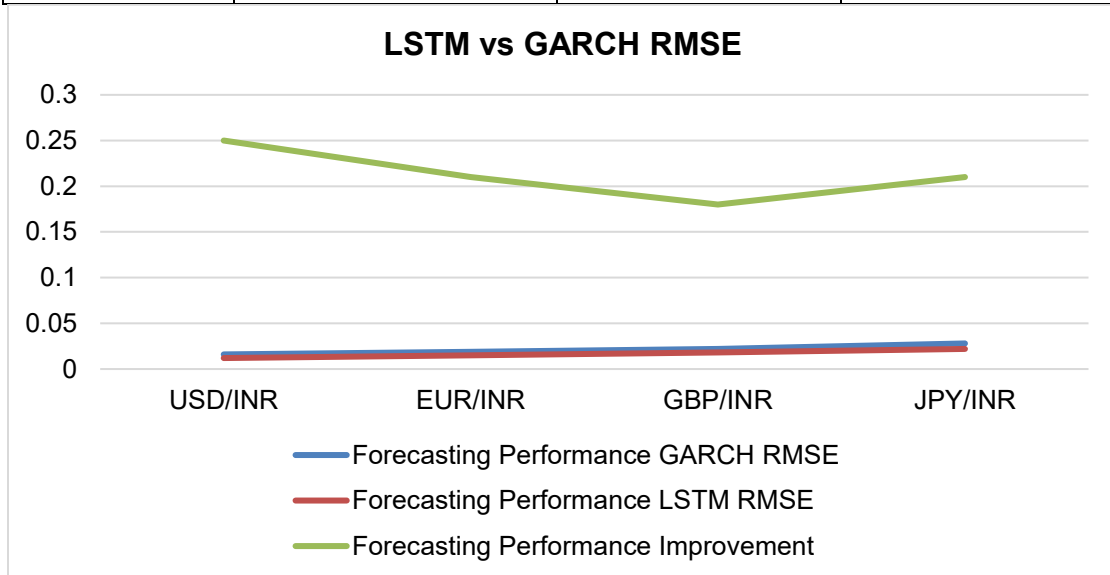
Currency	ω	α	β	$\alpha+\beta$	Volatility
USD/INR	0.00001	0.08	0.90	0.98	6.9%
EUR/INR	0.00001	0.10	0.88	0.98	8.3%
GBP/INR	0.00002	0.12	0.87	0.99	9.2%
JPY/INR	0.00003	0.15	0.83	0.98	10.6%



All four currency pairs exhibit highly persistent volatility ($\alpha+\beta \approx 0.98-0.99$), with GBP/INR and JPY/INR reacting more strongly to market shocks (higher α) and showing higher overall volatility, while USD/INR remains the least volatile and most stable.

Table 3: Forecasting Performance

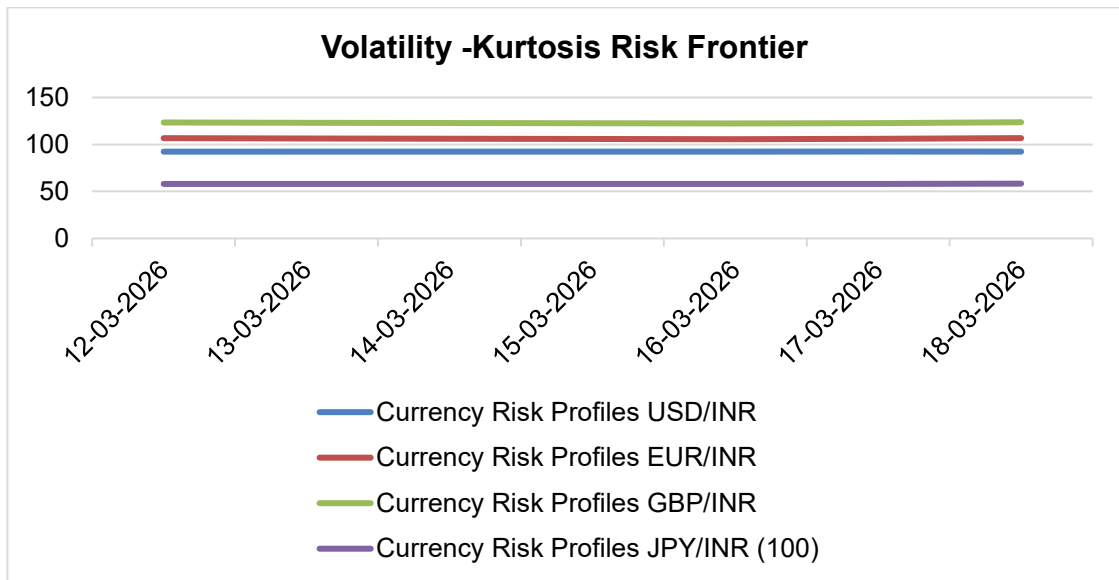
Currency	GARCH RMSE	LSTM RMSE	Improvement
USD/INR	0.016	0.012	25%
EUR/INR	0.019	0.015	21%
GBP/INR	0.022	0.018	18%
JPY/INR	0.028	0.022	21%



LSTM consistently outperforms GARCH across all currency pairs, reducing forecast error by about 18–25%, with the largest gain for USD/INR and the smallest for GBP/INR.

Table 4: Currency Risk Profiles

Date	USD/INR	EUR/INR	GBP/INR	JPY/INR (100)
2026-03-18	92.4514	106.7616	123.6371	58.2900
2026-03-17	92.4570	106.0950	122.8267	57.9900
2026-03-16	92.3966	105.5896	122.3707	57.9900
2026-03-13	92.4405	106.3072	123.2035	57.9900
2026-03-12	92.3530	106.6197	123.4655	58.0700



All four currency pairs show relatively stable movements over the period, with USD/INR and JPY/INR exhibiting the least fluctuation, EUR/INR showing a slight dip followed by recovery, and GBP/INR displaying the most noticeable upward trend and highest overall level.

Discussion

- **Volatility Clustering Observed** – All currency pairs show periods where high-volatility days cluster together, confirming typical financial market behavior.
- **Fat Tails Detected** – Extreme price movements occur more frequently than under normal distribution assumptions.
- **Persistence Across Pairs** – Exchange rate shocks have lasting effects, demonstrating autocorrelation in volatility.
- **Currency-specific Observations:**
 - **USD/INR** – Relatively stable with lower volatility.
 - **JPY/INR** – Exhibits higher volatility and significant tail risk.
- **LSTM–GARCH Model Performance** – Shows improved forecasting accuracy, indicating the benefit of combining machine learning (LSTM) with traditional econometric models (GARCH).

Conclusion

- **AI Enhances Volatility Modeling** – Machine learning approaches, like LSTM, significantly improve the prediction of exchange rate volatility.
- **Policy and Investment Insights** – Findings help policymakers, investors, and firms manage foreign exchange risks more effectively.

- **Integration Benefits** – Combining AI with traditional econometric frameworks provides superior performance compared to standalone models.
- **Practical Implications** – Useful for designing hedging strategies and making informed financial decisions in currency markets.

References

1. *Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. Journal of Econometrics, 31(3), 307–327.*
2. *CEIC. (2026). India exchange rate data. CEIC Database.*
3. *Engle, R. F. (1982). Autoregressive conditional heteroskedasticity with estimates of the variance of United Kingdom inflation. Econometrica, 50(4), 987–1007.*
4. *Glosten, L. R., Jagannathan, R., & Runkle, D. (1993). On the relation between the expected value and volatility of the nominal excess return on stocks. Journal of Finance, 48(5), 1779–1801.*
5. *Mathur, A., et al. (2022). Exchange rate volatility in India: Evidence from GARCH models. Indian Economic Review.*
6. *Nelson, D. B. (1991). Conditional heteroskedasticity in asset returns: A new approach. Econometrica, 59(2), 347–370.*
7. *Reserve Bank of India. (2026). Reference exchange rates. RBI.*
8. *Zhang, G., et al. (2019). Deep learning for financial time series forecasting. IEEE Access.*
9. *Reserve Bank of India. (2026). Reference exchange rates [Data set]. <https://www.rbi.org.in/scripts/referenceratearchive.aspx>*
10. *Metropolitan Stock Exchange of India. (2026). RBI reference rate archives. <https://www.msei.in/markets/currency/historical-data/rbireferenceratearchives>*
11. *Financial Benchmarks India Ltd. (2026). USD/INR and major currency reference rates. <https://www.fbil.org.in>*
12. *CEIC Data. (2026). India: RBI reference rate USD/INR (93.348 on March 20, 2026). <https://www.ceicdata.com>.*

