IRJEMS International Research Journal of Economics and Management Studies
ISSN: 2583 – 5238 / Volume 5 Issue 8 August 2026 / Pg. No: 37-53
Paper Id: IRJEMS-V5I8P104, Doi: 10.56472/25835238/IRJEMS-V5I8P104

Original Article

Comparative Analysis of the Volatility of Indian Rupee Exchange Rates against Major Foreign Currencies: Evidence from ARCH and GARCH Models

[1]Hari V. G.
[1]*Department of Statistics, Christian College, Kattakada, Thiruvananthapuram, Kerala, India.*

Received Date: 02 July 2026 Revised Date: 22 July 2026 Accepted Date: 28 July 2026 Published Date: 04 August 2026

Abstract: *This study examines and compares the conditional volatility of the Indian rupee against the US dollar, pound sterling, euro and Japanese yen using daily exchange rate data covering April 2010 to March 2026. After examining the time-series properties and stylised features of the return series, a range of ARCH and GARCH-family models is estimated under alternative innovation distributions. The analysis identifies the most appropriate volatility specification for each currency and compares their relative volatility and persistence. The findings indicate substantial differences in the volatility characteristics of the four exchange rates, with the Japanese yen exhibiting the highest unconditional volatility and the US dollar the lowest. Volatility is highly persistent across the currencies, although persistence varies across regimes and declines when the sample is divided into pre-pandemic, pandemic and post-pandemic periods, suggesting that structural changes can influence estimates of volatility persistence. Out-of-sample results further indicate that GARCH-based models provide useful improvements over a random-walk volatility benchmark, although their forecasting advantage over simple rolling measures is limited. The study provides comparative evidence on the dynamics and persistence of major foreign exchange risks faced by the Indian rupee.*

Keywords: *Emerging Market Currencies, Exchange Rate Volatility, GARCH, Indian Rupee, Volatility Persistence.*

I. INTRODUCTION

The exchange rate occupies a central position among macro-financial variables. It transmits shocks between domestic and international markets, shapes the competitiveness of tradable sectors, conditions the pass-through of foreign price movements into domestic inflation, and determines the domestic-currency value of external assets and liabilities. For an economy such as India, where merchandise and services trade together approach two-fifths of gross domestic product, where external commercial borrowing and portfolio flows are substantial. Where a large share of the import bill is denominated in dollars and concentrated in energy, the second moment of the exchange rate is of direct consequence for firms, investors and the monetary authority alike.

The period examined here, April 2010 to March 2026, spans an unusually dense sequence of disturbances. It opens in the aftermath of the global financial crisis. It encompasses the Euro-area sovereign debt episode, the 2013 taper tantrum that placed acute pressure on the rupee, the 2016 referendum on United Kingdom membership of the European Union, the demonetisation of high-denomination notes later that year, the COVID-19 pandemic and the associated collapse and rebound in global activity, the synchronised monetary tightening of 2022, and the sustained accumulation of foreign exchange reserves by the Reserve Bank of India in the years that followed. Each of these events had the capacity to alter not merely the level of the rupee but the intensity and persistence of its fluctuations.

Financial returns are well known to depart from the assumptions of classical linear time-series models. They exhibit volatility clustering, in which large movements of either sign follow large movements; leptokurtosis, so that extreme outcomes occur far more often than a Gaussian law would allow; and persistence, whereby the effect of a shock on conditional variance decays only slowly. The autoregressive conditional heteroscedasticity model of Engle (1982) and its generalisation by Bollerslev (1986) were designed precisely to accommodate these features, and the subsequent asymmetric extensions of Nelson (1991) and Glosten, Jagannathan and Runkle (1993) allow the sign of a shock, and not only its magnitude, to influence subsequent variance.

This paper applies that apparatus comparatively. Rather than modelling a single bilateral rate, it estimates a common set of specifications across four rupee cross-rates against the dollar, sterling, the euro and the yen over an identical sample, so that differences in estimated volatility, asymmetry and persistence can be attributed to the currency pair rather than to differences in sample period, data frequency or estimation choices. Four questions organise the analysis. First, what are the stylised statistical properties of daily rupee returns against each of the four currencies? Second, are ARCH effects present, and which member of the ARCH family best describes each series? Third, how do the currencies rank on unconditional volatility, on the persistence of variance shocks, and on the half-life of those shocks? Fourth, does the pandemic mark a genuine change in volatility dynamics, or does apparent persistence over the full sample instead reflect unmodelled shifts in the variance regime?

The contribution is threefold. The comparative design, holding sample and method fixed across four currencies, permits a cleaner ranking of exchange-rate risk than the single-pair studies that dominate the applied literature. The sample extends through March 2026 and therefore covers the post-pandemic normalisation, a period during which rupee and dollar volatility has fallen to levels not previously observed in the series. And the treatment of persistence is explicit rather than incidental: the paper reports that the rupee–dollar conditional variance is empirically indistinguishable from an integrated process, examines whether this reflects genuine long memory or the neglected-regime-shift mechanism identified by Lamoureux and Lastrapes (1990), and declines to report a half-life where the parameter estimates do not support one.

The remainder of the paper is organised as follows. Section 2 reviews the relevant literature. Section 3 describes the data and sets out the econometric methodology. Section 4 presents the empirical results. Section 5 discusses their economic interpretation and policy implications. Section 6 concludes and identifies limitations.

II. LITERATURE REVIEW

A) The Statistical Character of Exchange Rate Returns

That speculative price changes are not normally distributed has been understood since Mandelbrot (1963) and Fama (1965) documented tails far heavier than the Gaussian law permits. Mandelbrot also observed the clustering of large changes that later motivated conditional variance modelling. For foreign exchange specifically, Baillie and Bollerslev (1989) established that daily returns display both marked leptokurtosis and strong dependence in squared returns. Conditional variance models fitted with Student-t innovations describe the data substantially better than Gaussian alternatives a finding replicated across a wide range of currency pairs and sample periods.

A related literature concerns the level of the exchange rate rather than its variance. Following Meese and Rogoff (1983), the difficulty of outperforming a driftless random walk in out-of-sample exchange-rate prediction has proved remarkably durable. This result underpins the standard practice, adopted here, of treating the log exchange rate as a unit-root process and directing modelling effort at the conditional second moment rather than the conditional mean.

B) Conditional Variance Models

Engle (1982) introduced the ARCH model, in which conditional variance is a linear function of past squared innovations. In practice, the specification requires long lag structures to capture observed persistence, and Bollerslev (1986) resolved this by admitting lagged conditional variances, yielding the GARCH model. The GARCH(1,1) parameterisation has since become the workhorse of applied volatility analysis; Hansen and Lunde (2005), comparing 330 specifications across equity and currency data, found no evidence that any of them outperformed GARCH(1,1) for exchange rates.

Symmetric models restrict the conditional variance to depend on the magnitude of the innovation alone. Nelson (1991) proposed the exponential GARCH model, which specifies the logarithm of conditional variance and thereby dispenses with non-negativity constraints while admitting an asymmetric response through a signed innovation term. Glosten, Jagannathan and Runkle (1993) achieved a similar end within the linear variance framework by adding an indicator for negative innovations. In equity markets, these extensions capture the leverage effect, whereby negative returns raise volatility more than positive returns of equal size. In foreign exchange, the direction of asymmetry is not theoretically determined. It is an empirical matter, since the sign of a return depends on the convention by which the pair is quoted.

Bollerslev (1987) introduced Student-t innovations to accommodate conditional leptokurtosis beyond that generated by the variance dynamics themselves, and Bollerslev and Wooldridge (1992) developed the quasi-maximum-likelihood standard errors that render inference robust to misspecification of the innovation density. Both are standard in current practice and are employed here.

C) Persistence and Structural Change

Estimated GARCH persistence in long financial samples is frequently close to unity, a configuration formalised as integrated GARCH by Engle and Bollerslev (1986), in which variance shocks never fully dissipate. Lamoureux and Lastrapes (1990) offered an influential alternative reading: if the unconditional variance shifts across regimes and the model does not accommodate those shifts, the estimated persistence is biased towards unity. Their diagnostic is straightforward – partition the sample and re-estimate, and if persistence falls materially, near-integration was at least partly an artefact of neglected regime change.

Detecting such shifts requires a formal procedure. Inclán and Tiao (1994) proposed the iterated cumulative sums of squares algorithm, which identifies multiple variance change points by recursively applying a centred cumulative-sum-of-squares statistic to subsamples. A parallel literature addresses breaks in the level: Perron (1989) showed that unaccounted structural breaks bias unit-root tests towards non-rejection, and Zivot and Andrews (1992) developed a test in which the break date is chosen endogenously to be least favourable to the unit-root null.

D) Forecast evaluation

Evaluating volatility forecasts is complicated by the latency of the target. Andersen and Bollerslev (1998) demonstrated that the squared daily return, while an unbiased proxy for conditional variance, is extremely noisy, and that the low coefficients of determination reported in earlier work reflected the proxy rather than genuine forecast failure. Patton (2011) established that only certain loss functions squared error and the logarithmic QLIKE criterion among them preserve the ranking of competing forecasts when a noisy proxy is substituted for the latent variance. Loss functions expressed in relative terms, such as mean absolute percentage error, are not robust in this sense and become unstable when the proxy approaches zero. Diebold and Mariano (1995) supplied the standard test of equal predictive accuracy, and the present paper follows all of these conventions.

E) Evidence on the Indian rupee

India-specific empirical research on rupee exchange-rate volatility has extensively documented the presence of volatility clustering, conditional heteroscedasticity, persistence and, in several cases, asymmetric responses to exchange-rate shocks. Early studies using ARCH-family models established the relevance of time-varying volatility in the Indian foreign-exchange market, with evidence that exchange-rate volatility is influenced by monetary conditions, exchange-market intervention, capital flows and international financial developments (Karmakar, 2005; Patnaik and Shah, 2010; Bhattacharya and Mukherjee, 2013; Goyal et al., 2013). Studies employing GARCH and related models have generally found substantial persistence in INR volatility, although the magnitude and dynamics vary across exchange-rate regimes and currency pairs (Choudhry et al., 2005; Kumar, 2013; Murari, 2015; Bhowmik, 2016; Sahoo and Kamaiah, 2017). Comparative evidence for the rupee against major international currencies suggests that asymmetric specifications may outperform conventional symmetric GARCH models in capturing the effects of positive and negative exchange-rate innovations, particularly where leverage-type effects are present (Murari, 2015). Subsequent research has examined USD/INR volatility using GARCH, EGARCH, TGARCH and APARCH models, confirming the continuing importance of conditional heteroscedasticity and indicating that model performance can differ substantially across alternative volatility specifications (Mahalwala, 2022). Recent studies have further extended the analysis to the volatility and forecasting behaviour of the rupee against the US dollar, euro, Japanese yen and pound sterling, reporting persistent volatility shocks in the Indian foreign-exchange market (Raja et al., 2024). Related research has also shown that INR exchange-rate returns exhibit non-normality, fat tails, long-range dependence and multifractal behaviour, reinforcing the limitations of simple constant-variance models (Datta, 2023).

The broader literature additionally documents significant volatility spillovers and interconnectedness between the Indian rupee and other Asian and global currencies, suggesting that domestic exchange-rate volatility is shaped not only by country-specific fundamentals but also by international financial transmission mechanisms (Kumar, 2013; Bouri et al., 2020; Anwer et al., 2022; Das and Roy, 2023; He et al., 2024). Despite this substantial body of work, the literature remains fragmented across individual currency pairs, different sample periods and alternative modelling frameworks. In particular, an updated, systematic comparison of the volatility magnitude, persistence and regime sensitivity of the rupee against the US dollar, pound sterling, euro and Japanese yen over a common long daily sample remains relatively limited. The present study contributes to this literature by providing such a comparative analysis using a unified set of ARCH and GARCH-family models and by explicitly examining how volatility persistence changes across major market regimes. Against whatever prior evidence is established there, the present study is distinguished by its comparative design across four currencies over a common sample, by an updated sample extending to March 2026, and by explicit treatment of the distinction between genuine and spurious volatility persistence.

III. DATA AND METHODOLOGY

A) Data

The analysis uses daily Reserve Bank of India reference rates for the Indian rupee against the US dollar, pound sterling, euro and Japanese yen from April 5, 2010 to March 30, 2026. The yen rate is quoted per 100 yen, in accordance with Reserve Bank convention; because the analysis is conducted throughout in logarithmic differences, this scaling affects neither the return series nor any subsequent result.

Continuously compounded returns are computed as:

$$r_t = 100 \times \left[ln\left(S_t\right) - ln\left(S_{t-1}\right) \right], \quad (1)$$

Where S_t is the rupee price of one unit of foreign currency, yielding 3,861 return observations per series. Under this convention, a positive return denotes a depreciation of the rupee. This point is material to the interpretation of the asymmetry parameters reported below, and is stated here to avoid ambiguity.

B) Preliminary testing

Three complementary procedures assess stationarity. The augmented Dickey–Fuller test (Dickey and Fuller, 1979; Said and Dickey, 1984) estimates

$$\Delta r_t = \alpha + \delta t + \gamma r_{t-1} + \sum_{i=1}^{n} \varphi_i \Delta r_{t-i} + \varepsilon_t, \quad (2)$$

and tests $H_0 : \gamma = 0$ against the stationary alternative, with lag order selected by the Akaike criterion. Because the ADF test has low power against persistent stationary alternatives, it is supplemented by the Phillips–Perron (1988) test, which corrects the first-order statistic non-parametrically for serial correlation using a Newey–West long-run variance, and by the KPSS test of Kwiatkowski et al. (1992), which reverses the hypotheses by taking stationarity as the null. Agreement across all three provides considerably firmer ground for an order-of-integration conclusion than any one test alone.

Because unaccounted level breaks bias unit root tests towards non-rejection (Perron, 1989), the Zivot–Andrews (1992) test is also applied, with the break date selected endogenously as the one minimising the t-statistic on the autoregressive parameter and the search trimmed by 15 per cent at each end of the sample.

Serial dependence is examined using the Ljung–Box (1978) Q statistic applied to returns and to squared returns; normality using the Jarque–Bera (1980) statistic. Conditional heteroscedasticity is tested by the Lagrange multiplier procedure of Engle (1982), in which squared residuals are regressed on q of their own lags and the statistic TR^2 is referred to a chi-square distribution with q degrees of freedom.

C) Volatility models

All models share a constant conditional mean, $r_t = \mu + \varepsilon_t$, with $\varepsilon_t = \sigma_t z_t$ and z_t independently and identically distributed with zero mean and unit variance. Five variance specifications are estimated.

The ARCH(q) model of Engle (1982):

$$\sigma_t^2 = \omega + \sum_{i=1}^{q} \alpha_t \varepsilon_{t-i}^2, \quad (3)$$

estimated for $q = 1$ and $q = 5$.

The GARCH(1,1) model of Bollerslev (1986):

$$\sigma_t^2 = \omega + \alpha \varepsilon_{t-1}^2 + \beta \sigma_{t-1}^2, \quad (4)$$

With $\omega > 0$, $\alpha \geq 0$, $\beta \geq 0$ for positive variance and $\alpha + \beta < 1$ for covariance stationarity. The quantity $\alpha + \beta$ measures persistence; the implied half-life of a variance shock is $ln(0.5) / ln(\alpha + \beta)$ trading days and the unconditional variance is $\omega / (1 - \alpha - \beta)$, both defined only when persistence lies strictly below unity.

The GJR-GARCH(1,1,1) model of Glosten, Jagannathan and Runkle (1993):

$$\sigma_t^2 = \omega + \alpha \varepsilon_{t-1}^2 + \gamma \varepsilon_{t-1}^2 I[\varepsilon_{t-1} < 0] + \beta \sigma_{t-1}^2, \quad (5)$$

in which $\gamma \neq 0$ indicates asymmetry and persistence is $\alpha + \gamma / 2 + \beta$.

Finally, the EGARCH(1,1) model of Nelson (1991):

$$ln(\sigma_t^2) = \omega + \alpha \left[|z_{t-1}| - E|z_{t-1}| \right] + \gamma z_{t-1} + \beta \, ln(\sigma_{t-1}^2), \quad (6)$$

Where modelling the logarithm removes the need for non-negativity restrictions and β alone governs persistence. Given the quotation convention adopted here, a positive and significant γ implies that rupee depreciations raise subsequent volatility by more than appreciations of equal magnitude.

Each specification is estimated under both Gaussian and standardised Student-*t* innovations, the latter introducing a degrees-of-freedom parameter $\nu > 2$ estimated jointly with the remaining parameters (Bollerslev, 1987). Model selection is by the Schwarz information criterion, which penalises parameters more heavily than Akaike and is therefore the more conservative choice; the Akaike and Hannan–Quinn criteria are also reported.

D) Sub-period analysis and forecast evaluation

To assess whether pandemic-era disruption altered volatility dynamics, and to implement the Lamoureux–Lastrapes diagnostic, the sample is partitioned using World Health Organization milestones as exogenous and verifiable boundaries: a pre-pandemic period from the start of the sample to January 29, 2020; a pandemic period from January 30, 2020, when a public health emergency of international concern was declared, to May 05, 2023, when that declaration was terminated; and a post-pandemic period running to the end of the sample. GARCH(1,1)-*t* is re-estimated within each. In parallel, the iterated cumulative sums of squares algorithm of Inclán and Tiao (1994) is applied to each full return series to identify variance change points without imposing dates a priori.

Forecast performance is assessed on the final 500 observations, beginning February 29 2024, which are withheld from estimation entirely. Parameters are estimated once on the initial 3,361 observations and the variance recursion is then filtered forward, so that each one-step-ahead forecast uses only information available at the time. Because conditional variance is unobservable, the squared return serves as its proxy, with the well-understood consequence that all reported losses are inflated by proxy noise (Andersen and Bollerslev, 1998). Performance is measured by root mean squared error, mean absolute error, Theil's U relative to a random-walk variance benchmark, and the QLIKE criterion; mean absolute percentage error is reported for completeness but, as discussed in Section 4.6, is not informative in this setting. Diebold tests the equality of predictive accuracy–Mariano (1995) statistic with a Newey West variance estimator. Three benchmarks are used: a random walk in variance, a 22-day rolling sample variance, and the in-sample unconditional variance.

IV. EMPIRICAL RESULTS

A) Descriptive statistics and stylised facts

Figure 1 plots the four exchange rates in levels and Figure 2 the corresponding returns. The level plots show the sustained depreciation of the rupee against the dollar, from roughly 45 to 95 over sixteen years, together with the sharp adjustments of 2013 and 2018 and the more gradual drift thereafter. The return plots display the clustering characteristic of financial data: quiet intervals are interrupted by concentrated bursts of activity around 2013, early 2020 and 2022.

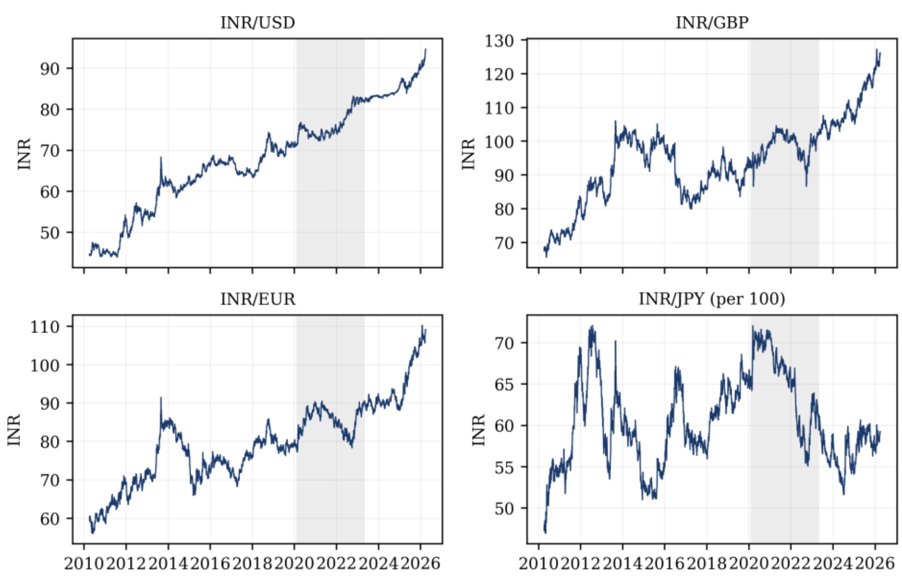

Figure 1. Daily Exchange Rates in Levels, April 2010 – March 2026

Note: *Shaded region denotes the WHO public health emergency of international concern (January 30 2020 – May 5 2023). The yen rate is quoted per 100 yen.*

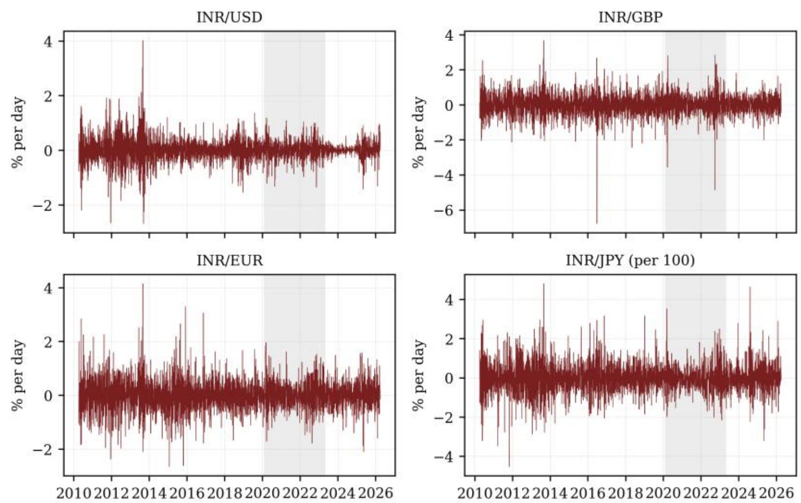

Figure 2. Daily Continuously Compounded Returns (Per Cent)

Note: *Volatility clustering is visible in all four series, most pronounced around mid-2013 and the first half of 2020.*

Table 1 reports the descriptive statistics. Mean daily returns are positive for all four pairs, consistent with trend depreciation of the rupee, but are economically small and except the dollar statistically indistinguishable from zero. The ordering of standard deviations is unambiguous: the yen is the most volatile counterpart at 11.29 per cent annualised, followed by sterling at 9.30 per cent and the euro at 8.76 per cent, with the dollar substantially the least volatile at 6.46 per cent. That the rupee should be most stable against the currency in which the bulk of Indian trade is invoiced is not accidental, and Section 5 returns to the point.

Table 1: Descriptive Statistics of Daily Log Returns, April 2010 – March 2026

Statistics	INR/USD	INR/GBP	INR/EUR	INR/JPY
Observations	3,861	3,861	3,861	3,861
Mean (%)	0.0194	0.0158	0.0153	0.0058
Median (%)	0.0122	0.0224	0.0103	−0.0177
Maximum (%)	4.0200	3.6789	4.1545	4.8114
Minimum (%)	−2.6793	−6.7758	−2.6520	−4.5293
Std. deviation (%)	0.4069	0.5861	0.5519	0.7109
Annualised volatility (%)	6.459	9.303	8.761	11.285
Skewness	0.299	−0.446	0.303	0.287
Excess kurtosis	7.181	7.683	3.048	3.351
Jarque–Bera	8353.1***	9622.9***	1553.1***	1859.6***
Ljung–Box Q(10)	31.69***	16.89*	26.07***	4.47
Ljung–Box $Q^2(10)$	1992.2***	243.8***	492.4***	520.0***

Notes: *Returns are 100 × the first difference of log exchange rates. Annualised volatility is the daily standard deviation scaled by √252. Q(10) and $Q^2(10)$ are Ljung–Box statistics at ten lags on returns and squared returns respectively, and denote significance at the 1, 5 and 10 per cent levels.*

Every series is leptokurtic, with excess kurtosis ranging from 3.05 for the euro to 7.68 for sterling, and Jarque–Bera statistics reject normality decisively in all four cases. Skewness is positive for the dollar, euro and yen – indicating that sharp rupee depreciations outnumber comparable appreciations – but negative for sterling, reflecting the extreme single-day appreciation of the rupee against sterling of 6.78 per cent that followed the 2016 referendum, the largest single observation anywhere in the dataset. The two Ljung–Box statistics carry different messages. Linear dependence in returns is weak: the statistic is significant at the one per cent level for the dollar and euro, marginal for sterling and absent for the yen, and the implied autocorrelations are small. Dependence in squared returns is by contrast overwhelming, with statistics between 244 and 1,992 against a critical value of 23.2. This is the signature of volatility clustering and the empirical justification for the conditional variance models that follow.

B) Unit root and structural break tests

Table 2 reports the stationarity tests. In log levels, the ADF and PhillipsPerron statistics fail to reject the unit root null for all four series at the five per cent level, while the KPSS statistics between 1.55 and 4.09 against a critical value of 0.146 – reject stationarity emphatically. In first differences the position reverses completely: ADF statistics between −11.10 and −12.79 and Phillips–Perron statistics around −60 reject the unit root at any conventional level, and KPSS statistics between 0.06 and 0.14 lie far below the critical value. The three procedures agree, and all four exchange rates are integrated of order one with stationary returns.

Table 2: Unit Root Tests

	ADF	PP	KPSS	ADF	PP	KPSS
	Log levels			**Returns**		
INR/USD	−2.685	−2.747	4.090***	−11.579***	−62.53***	0.072
INR/GBP	−2.342	−2.423	2.802***	−12.792***	−59.79***	0.088
INR/EUR	−2.683	−2.619	1.549***	−11.962***	−59.05***	0.060
INR/JPY	−3.182*	−3.179*	2.514***	−11.101***	−62.50***	0.141
5% critical value	−3.422	−3.422	0.146	−2.866	−2.866	0.463

Notes: *Level tests include a constant and linear trend; return tests include a constant only. ADF lag order selected by AIC; Phillips Perron and KPSS use a Newey–West bandwidth of* $\lfloor 4(T/100)^{1/4} \rfloor$*. ADF and PP critical values obtained by Monte Carlo simulation of the Dickey–Fuller distribution at the relevant sample size. For ADF and PP, the null is a unit root; for KPSS the null is stationarity, and denote significance at the 1, 5 and 10 per cent levels.*

Table 3 addresses the possibility that non-rejection in levels reflects an unaccounted structural break. The Zivot–Andrews statistics range from −4.02 to −4.63, and none approaches the five per cent critical value of −5.08, so the unit-root null survives even when a break in intercept and trend is permitted at the least favourable date. The endogenously selected break dates are nonetheless informative in their own right. For sterling, the algorithm selects June 24 2016 the day the referendum result was announced. For the dollar, it selects May 3 2013, immediately preceding the taper tantrum. These dates were not imposed; that the procedure recovers them independently is a useful indication that the series behave as the historical narrative would suggest.

Table 3: Zivot–Andrews test with endogenous break, and ICSS variance change points

Series	ZA statistic	Break date	Interpretation	ICSS breaks
INR/USD	−4.187	May 03, 2013	Pre-taper tantrum	19
INR/GBP	−4.486	June 24, 2016	EU referendum result	16
INR/EUR	−4.632	August 11, 2014	Euro-area stagnation	13
INR/JPY	−4.016	January 9 2018	BoJ policy adjustment	19
5% critical value	−5.080			

Notes: *Zivot–Andrews Model C allows a break in both intercept and trend, with the break date chosen to minimise the t-statistic on the autoregressive parameter and 15 per cent trimming. No statistic rejects the unit-root null. ICSS breaks are variance change points in the return series detected by the Inclán–Tiao algorithm with a minimum segment length of 60 observations. Interpretations in column four are the authors' reading of the associated events and are indicative rather than tested.*

The final column of Table 3 reports a different and, for present purposes, more consequential result. The ICSS algorithm identifies between 13 and 19 distinct variance regimes in each return series over the sixteen-year sample. The unconditional variance of these series is therefore very far from constant, which bears directly on the interpretation of the persistence estimates in Section 4.4.

C) ARCH effects and model selection

Table 4 reports Engle's LM test. At five lags, the statistics range from 157.9 for sterling to 646.2 for the dollar, against a one per cent critical value of 15.1; the null of no ARCH effect is rejected at every lag length and for every currency by an enormous margin. Conditional heteroscedasticity is present, and modelling it is not optional.

Table 4: Engle ARCH-LM Test on Mean-Adjusted Returns

Series	LM(1)	LM(5)	LM(10)
INR/USD	203.52***	646.22***	760.60***
INR/GBP	43.53***	157.92***	173.37***
INR/EUR	157.66***	245.47***	300.88***
INR/JPY	147.80***	218.01***	290.41***
1% critical value	6.63	15.09	23.21

Notes: *The statistic is* TR^2 *from a regression of squared residuals on q of their own lags, distributed* $\chi^2(q)$ *under the null of no ARCH effect, and denotes significance at the 1 per cent level.*

Table 5 compares the ten estimated specifications for each currency by the Schwarz criterion. Two regularities are immediate. First, the Student-t distribution improves fit relative to the Gaussian for every currency and every variance specification, by margins of 150 to 290 Schwarz units differences far too large to attribute to sampling variation, and consistent with the excess kurtosis documented in Table 1. Second, the ARCH specifications are decisively rejected in favour of the GARCH family. ARCH(1) is the worst model in every case, and even ARCH(5) remains 150 to 290 units behind the best GARCH specification, confirming that a parsimonious lagged-variance term captures dependence that a finite distributed lag of squared innovations cannot.

Table 5: Model Selection by Schwarz Information Criterion

Specifications	INR/USD		INR/GBP		INR/EUR		INR/JPY	
	Normal	*t*	Normal	*t*	Normal	*t*	Normal	*t*
ARCH(1)	3457.3	2849.2	6763.7	6400.6	6200.4	5975.3	8147.7	7791.5
ARCH(5)	2704.2	2414.6	6574.2	6354.9	6114.8	5924.1	7929.7	7663.1
GARCH(1,1)	2362.5	2139.6	6484.1	6289.0	5983.9	5829.6	7771.2	7548.0
GJR-GARCH(1,1,1)	2374.3	2149.2	6493.5	6297.2	5992.2	5839.1	7779.5	7564.6
EGARCH(1,1)	2334.9	2123.8	6496.7	6292.4	5994.7	5832.5	7742.3	7532.9

Notes: *Lower values indicate better fit. The preferred specification for each currency is EGARCH(1,1)-t for INR/USD (2123.8) and INR/JPY (7532.9), and GARCH(1,1)-t for INR/GBP (6289.0) and INR/EUR (5829.6). The Akaike and Hannan–Quinn criteria select the same models in all four cases.*

Within the GARCH family, the choice is closer. For sterling and the euro, symmetric GARCH(1,1)-t is preferred, though EGARCH trails by only 3.4 and 2.9 Schwarz units respectively margins too narrow to treat as decisive. For the dollar and the yen, the asymmetric EGARCH specification is preferred by clearer margins of 15.8 and 15.1 units. GJR-GARCH is never selected; its indicator parameterisation of asymmetry does not suit these data, and for the dollar the estimated asymmetry coefficient is effectively zero.

D) Estimated volatility dynamics

Table 6 reports GARCH(1,1)-*t* estimates for all four currencies, providing a common basis for comparison, and Table 7 the preferred specifications from Table 5.

Table 6: GARCH(1,1)-

Parameters	INR/USD	INR/GBP	INR/EUR	INR/JPY
μ (mean)	0.0101***	0.0151*	0.0065	−0.0196**
	(0.0032)	(0.0081)	(0.0075)	(0.0090)
ω (constant)	0.0002**	0.0150***	0.0049**	0.0120***
	(0.0001)	(0.0052)	(0.0024)	(0.0036)
α (ARCH)	0.0806***	0.0656***	0.0558***	0.0942***
	(0.0148)	(0.0137)	(0.0141)	(0.0161)
β (GARCH)	0.9193***	0.8884***	0.9285***	0.8869***
	(0.0148)	(0.0270)	(0.0206)	(0.0194)
ν (d.o.f.)	6.171	7.050***	7.222***	5.498***
	(n.a.)	(0.798)	(0.741)	(0.477)
$\alpha+\beta$	1.0000	0.9540	0.9842	0.9811
Half-life (days)	undefined	14.7	43.7	36.4
Log-likelihood	−1049.14	−3123.84	−2894.16	−3753.35

Notes: *Bollerslev Wooldridge robust standard errors in parentheses, and denote significance at the 1, 5 and 10 per cent levels. For INR/USD, the persistence estimate reaches the boundary of the stationarity region, so the standard error on v and the implied half-life and unconditional variance are not reported; see the discussion in the text.*

All ARCH and GARCH coefficients are positive and significant at the one per cent level, satisfying the non-negativity conditions. The estimated degrees-of-freedom parameters lie between 5.50 for the yen and 7.22 for the euro, all far below the value at which the Student-t would approximate a normal distribution, confirming that conditional leptokurtosis persists after the variance dynamics have been modelled. The yen's low value of 5.50 indicates the heaviest conditional tails of the four.

The ARCH coefficient α measures the immediate response of variance to a shock, and is largest for the yen at 0.094 and the dollar at 0.081, smallest for the euro at 0.056. The GARCH coefficient β measures the carry-over of past variance and is correspondingly large everywhere, between 0.887 and 0.929. Their sum is the persistence measure of principal interest.

For sterling, persistence of 0.954 implies a half-life of 14.7 trading days: a shock to sterling volatility is half dissipated within three weeks. The euro and yen are considerably more persistent, at 0.984 and 0.981, implying half-lives of 43.7 and 36.4

days respectively – around two months and seven weeks. The rupee–dollar rate is a different case entirely. Its persistence estimate reaches the boundary of the stationarity region, and the maximum of the likelihood is not interior. To confirm that this reflects the data rather than a numerical artefact, the likelihood was profiled across a sequence of persistence ceilings from 0.95 to 0.9999; it increases monotonically throughout, and the constraint binds at every point, establishing that the unconstrained maximum-likelihood estimate satisfies $\alpha + \beta \geq 1$. The rupee–dollar conditional variance is therefore best characterised as integrated in the sense of Engle and Bollerslev (1986). Shocks to it do not decay at any finite rate, and neither a half-life nor an unconditional variance can be reported. Doing so would be arithmetically possible but substantively empty.

Table 7 presents the preferred specifications, and the EGARCH estimates for the dollar and yen add an important dimension.

Table 7: Preferred Specifications by Schwarz Criterion

Parameters	**INR/USD**	**INR/GBP**	**INR/EUR**	**INR/JPY**
Specification	EGARCH-t	GARCH-t	GARCH-t	EGARCH-t
μ	0.0112***	0.0151*	0.0065	−0.0117
	(0.0037)	(0.0081)	(0.0075)	(0.0092)
ω	−0.0200**	0.0150***	0.0049**	−0.0252***
	(0.0084)	(0.0052)	(0.0024)	(0.0069)
α	0.2123***	0.0656***	0.0558***	0.1808***
	(0.0432)	(0.0137)	(0.0141)	(0.0242)
γ (asymmetry)	0.0298***	-	-	0.0511***
	(0.0105)			(0.0125)
β	0.9910***	0.8884***	0.9285***	0.9709***
	(0.0038)	(0.0270)	(0.0206)	(0.0076)
β	5.406***	7.050***	7.222***	5.699***
	(0.439)	(0.798)	(0.741)	(0.515)
Persistence	0.9910	0.9540	0.9842	0.9709
Half-life (days)	76.5	14.7	43.7	23.5

Notes: *Bollerslev–Wooldridge robust standard errors in parentheses. For EGARCH, the persistence measure is β, and the variance equation is specified in logarithms, so ω is not restricted to be positive. Given the quotation convention, a positive γ implies that rupee depreciations raise subsequent volatility by more than appreciations of equal magnitude, and denote significance at the 1, 5 and 10 per cent levels.*

The asymmetry coefficient is positive and significant for both the dollar (0.0298, $p = 0.005$) and the yen (0.0511, $p < 0.001$). Because a positive return denotes rupee depreciation under the convention adopted here, this means that depreciations raise subsequent volatility by more than appreciations of the same magnitude. The direction is the opposite of the leverage effect familiar from equity markets. However, it is economically coherent for an emerging-market currency: depreciation episodes are associated with capital outflow, reserve loss and pressure on external balance sheets, whereas appreciations tend to be gradual and are frequently damped by official purchases. Figure 3 displays the implied news impact curves, which are visibly tilted towards positive shocks for the dollar and yen and effectively symmetric for sterling and the euro.

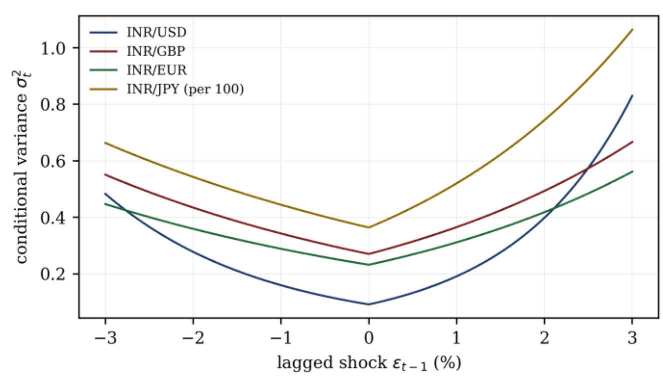

Figure 3. News Impact Curves Implied by EGARCH(1,1)-

Note: *Conditional variance is plotted as a function of the lagged shock, holding the lagged variance at its unconditional level. Under the quotation convention adopted here, positive shocks are rupee depreciations. The asymmetry is significant for INR/USD and INR/JPY only.*

It is worth noting that under EGARCH the rupee dollar persistence estimate is 0.991 high, but interior, and implying a half-life of 76.5 days. That the same series yields an integrated estimate under GARCH and a stationary one under EGARCH indicates that the conclusion of infinite persistence is specification-dependent, and this is stated explicitly rather than resolved in favour of whichever result is more convenient. What both specifications agree on is that rupee and dollar variance is the most persistent of the four by a wide margin.

Table 8 collects the comparative measures that Section 1 posed as the study's central question.

Table 8: Comparative Volatility Measures, GARCH(1,1)-

Measures	INR/USD	INR/GBP	INR/EUR	INR/JPY	Rank (highest risk)
Sample volatility (%, ann.)	6.46	9.30	8.76	11.29	JPY > GBP > EUR > USD
Unconditional volatility (%, ann.)	undefined	9.07	8.85	12.63	JPY > GBP > EUR
Mean conditional volatility (%, ann.)	5.70	8.98	8.51	10.99	JPY > GBP > EUR > USD
Peak conditional volatility (%, ann.)	26.92	30.23	23.34	32.60	JPY > GBP > USD > EUR
Volatility of volatility	3.14	2.06	2.16	3.49	JPY > USD > EUR > GBP
Persistence ($\alpha + \beta$)	≥ 1.000	0.954	0.984	0.981	USD > EUR > JPY > GBP
Half-life (trading days)	undefined	14.7	43.7	36.4	EUR > JPY > GBP
Tail thickness (ν)	6.17	7.05	7.22	5.50	JPY > USD > GBP > EUR

Notes: *Volatility of volatility is the standard deviation of the annualised conditional volatility series. Lower v denotes heavier tails and hence greater tail risk. Rankings for measures involving INR/USD omit that currency where the quantity is undefined.*

The rankings are consistent across the volatility measures: the yen is the riskiest counterpart on every one of them, the dollar the least risky on all but the volatility of volatility. The yen combines the highest average volatility, the highest peak volatility, the most variable volatility and the heaviest tails. It is unambiguously the currency against which the rupee carries the greatest exchange-rate risk. Sterling and the euro are close to one another and intermediate.

Persistence, however, ranks in almost the reverse order. The dollar, the least volatile pair, has the most persistent variance; sterling, considerably more volatile, has the least persistent. Volatility magnitude and volatility persistence are distinct properties, and a risk-management framework that attends to one while ignoring the other will misprice the horizon over which exposure must be hedged.

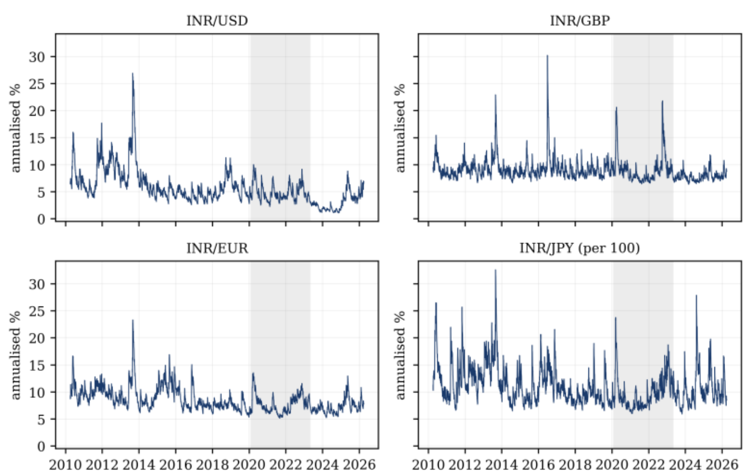

Figure 4. Estimated Annualised Conditional Volatility, GARCH(1,1)-T

Note: *Shaded region denotes the WHO public health emergency period. Common peaks occur in mid 2013 and early 2020; the INR/GBP spike in June 2016 is specific to the referendum.*

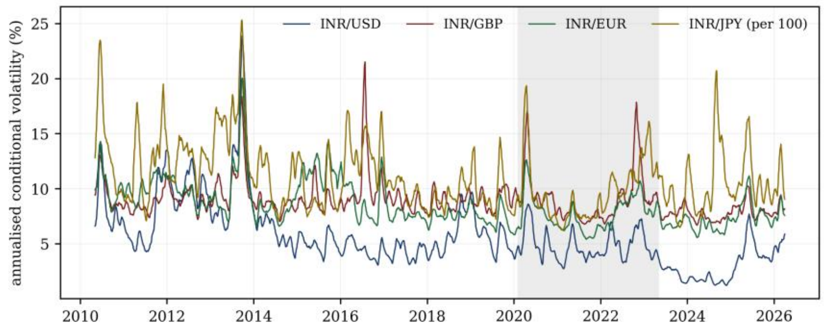

Figure 5. Conditional Volatility of the Four Rupee Cross-Rates Compared (21-Day Moving Average)
Note: *The INR/USD series lies persistently below the other three throughout the sample, and the gap widens after 2023.*

Figures 4 and 5 make the comparative picture visible. All four series peak simultaneously in mid-2013 during the taper tantrum, when conditional volatility reached 25 per cent annualised, and again in the first quarter of 2020. Beyond these common episodes, the series diverge: the rupee and sterling rate shows an isolated spike to 21.5 per cent in June 2016, and the rupee–yen rate exhibits recurrent independent peaks, most recently in mid-2024. The overlay in Figure 5 shows the rupee–dollar series lying below the others throughout, with the separation becoming pronounced after 2023.

E) Sub-period analysis

Table 9 reports GARCH(1,1)-*t* estimates within each of the three sub-periods.

Table 9: Sub-Period Estimates, GARCH(1,1)-

Series	Period	*N*	Ann. vol (%)	α	β	$\alpha+\beta$	Half-life
INR/USD	Pre-COVID	2,376	7.49	0.103	0.885	0.988	57.1
	COVID	784	4.98	0.173	0.735	0.908	7.2
	Post-COVID	701	3.46	0.130	0.870	≥ 1.000	undefined
INR/GBP	Pre-COVID	2,376	9.77	0.062	0.851	0.913	7.6
	COVID	784	9.66	0.100	0.857	0.957	15.8
	Post-COVID	701	6.98	0.059	0.746	0.805	3.2
INR/EUR	Pre-COVID	2,376	9.56	0.058	0.919	0.977	30.3
	COVID	784	7.69	0.053	0.935	0.988	55.0
	Post-COVID	701	6.87	0.079	0.782	0.861	4.6
INR/JPY	Pre-COVID	2,376	12.00	0.092	0.886	0.978	31.6
	COVID	784	9.63	0.070	0.921	0.991	73.0
	Post-COVID	701	10.47	0.133	0.812	0.945	12.2

Notes: *Pre-COVID runs to January 29 2020; COVID from January 30 2020 to May 5 2023 (WHO public health emergency of international concern); post-COVID from May 8 2023 to March 30 2026. Half-life in trading days.*

Three findings emerge. The first concerns the rupee and dollar rate, whose volatility declines monotonically and substantially across the three periods, from 7.49 per cent annualised before the pandemic to 4.98 per cent during it and 3.46 per cent after. Rupee–dollar volatility in the post-pandemic period is less than half its pre-pandemic level, and lower than in any earlier phase of the sample. This is the single most striking result in the study, and it runs contrary to the presumption that the pandemic and subsequent monetary tightening would have raised emerging-market currency volatility.

The second finding is that this pattern is specific to the dollar. Sterling volatility was essentially unchanged between the first two periods, at 9.77 and 9.66 per cent, before falling to 6.98 per cent. Euro volatility declined steadily but modestly. Yen volatility fell during the pandemic and then rose again to 10.47 per cent, remaining the highest of the four throughout. No general compression of rupee volatility is evident; what has changed is the rupee's behaviour against the dollar specifically.

The third finding bears on the interpretation of persistence. In every currency, persistence estimated within sub-periods is lower on average than persistence estimated over the full sample by 0.035 for the dollar, 0.063 for sterling, 0.042 for the euro and 0.010 for the yen. This is precisely the pattern Lamoureux and Lastrapes (1990) predicted, where the unconditional variance shifts across regimes and the model does not accommodate those shifts. Read alongside the 13 to 19 ICSS variance change points reported in Table 3, the evidence suggests that a material component of the near-integrated persistence estimated over the full sample is an artefact of neglected regime change rather than genuine long memory in volatility. The post-pandemic collapse in persistence for sterling (0.805) and the euro (0.861) is particularly marked.

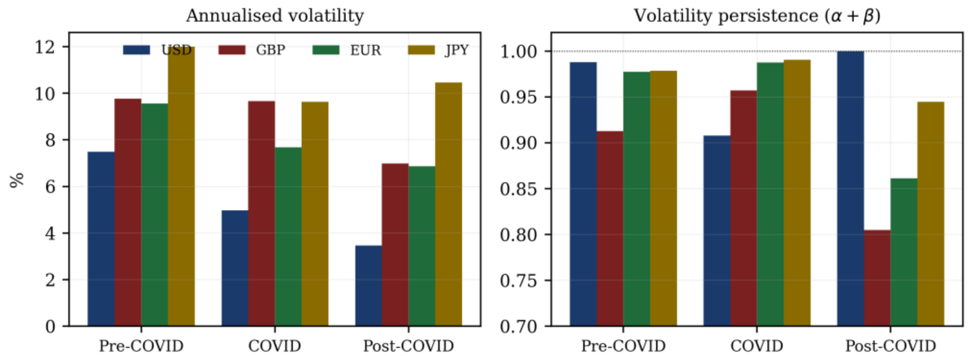

Figure 6. Annualised Volatility and Persistence by Sub-Period

Note: *Left panel, annualised volatility; right panel, GARCH persistence α + β with the unit boundary marked. The INR/USD post-COVID persistence estimate lies at the boundary of the stationarity region.*

F) Residual diagnostics

Table 10 reports diagnostics on the standardised residuals of the preferred models. The Ljung Box statistic on the levels of standardised residuals is insignificant at the five per cent level for all four currencies, indicating that the constant-mean specification is adequate and that no autoregressive term is required.

Table 10: Residual diagnostics for preferred specifications

Diagnostic	INR/USD	INR/GBP	INR/EUR	INR/JPY
Q(10) on *z*	13.12	13.85	16.41	7.85
p-value	(0.217)	(0.180)	(0.088)	(0.643)
Q(10) on z^2	41.85***	27.95***	21.14**	20.03**
p-value	(0.000)	(0.002)	(0.020)	(0.029)
ARCH-LM(5)	29.39***	24.29***	11.63**	14.83**
p-value	(0.000)	(0.000)	(0.040)	(0.011)
Skewness of *z*	−0.049	−0.356	0.221	0.334
Excess kurtosis of *z*	2.083	3.480	1.870	2.172

Notes: *z denotes standardised residuals from the preferred specification (EGARCH-t for INR/USD and INR/JPY; GARCH-t for INR/GBP and INR/EUR), and denote significance at the 1, 5 and 10 per cent levels.*

The diagnostics on squared standardised residuals are less satisfactory and are reported plainly rather than minimised. Residual ARCH remains statistically detectable in all four series, most strongly for the dollar and sterling. The models capture the great bulk of conditional heteroscedasticity the ARCH-LM(5) statistic for the dollar falls from 646.2 in the raw returns to 29.4 in the standardised residuals, a reduction of more than ninety-five per cent – but not all of it. Given a sample of 3,861 observations, statistical significance is attainable for economically trivial departures, and the residual excess kurtosis of 1.87 to 3.48 indicates that the Student-*t*, while a substantial improvement on the Gaussian, does not fully accommodate the tails. Figure 7 shows the corresponding quantile–quantile plots, in which the fit is close through the central mass and deteriorates in the extreme tails.

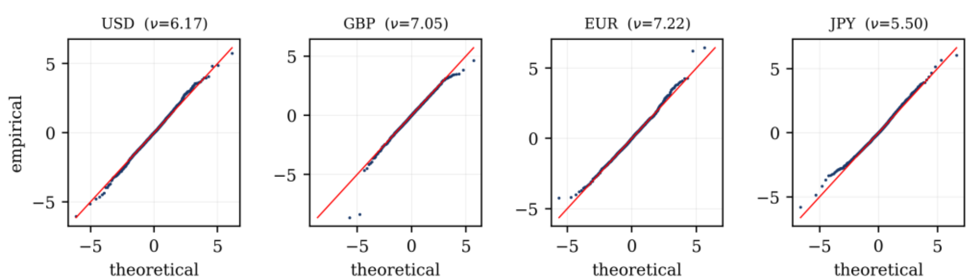

Figure 7. Quantile–Quantile Plots of Standardised Residuals Against the Fitted Student-

Note: *Estimated degrees of freedom shown in each panel. Departures from the 45-degree line are confined largely to the extreme tails.*

A plausible explanation, consistent with the ICSS evidence, is that a single set of parameters applied to sixteen years spanning many variance regimes cannot fully describe the data. A Markov-switching GARCH specification would be the natural response, and is identified in Section 6 as a direction for further work.

G) Out-of-sample forecast performance

Table 11 reports forecast accuracy over the 500 observations beginning February 29, 2024, which were withheld from estimation entirely.

Table 11: Out-of-Sample Volatility Forecast Accuracy, Final 500 Observations

Series	Forecast	RMSE	MAE	QLIKE	Theil U	DM vs RW
INR/USD	GARCH(1,1)-*t*	0.1670	0.0752	−2.246	0.750	−2.40**
	EGARCH(1,1)-*t*	0.1660	0.0725	−2.270	0.746	−2.41**
	Rolling variance (22d)	0.1663	0.0720	−2.114	0.747	−2.39**
	Random walk	0.2227	0.0861	4.336	1.000	—
INR/GBP	GARCH(1,1)-*t*	0.3513	0.2289	−0.599	0.740	−2.81***
	EGARCH(1,1)-*t*	0.3511	0.2269	−0.606	0.740	−2.76***
	Rolling variance (22d)	0.3532	0.2044	−0.551	0.744	−2.78***
	Random walk	0.4747	0.2527	6.840	1.000	—
INR/EUR	GARCH(1,1)-*t*	0.4218	0.2388	−0.609	0.761	−3.76***
	EGARCH(1,1)-*t*	0.4213	0.2368	−0.610	0.760	−3.81***
	Rolling variance (22d)	0.4274	0.2352	−0.486	0.771	−3.55***
	Random walk	0.5543	0.2788	8.661	1.000	—
INR/JPY	GARCH(1,1)-*t*	1.3394	0.5806	0.218	0.766	−1.95*
	EGARCH(1,1)-*t*	1.3211	0.5650	0.198	0.756	−1.96**
	Rolling variance (22d)	1.3665	0.5908	0.392	0.781	−1.88*
	Random walk	1.7487	0.6738	8.888	1.000	—

Notes: *The forecast target is the squared daily return. Theil's U is the ratio of a forecast's RMSE to that of the random-walk benchmark; values below unity indicate improvement. QLIKE is the loss of Patton (2011); lower is better. DM is the Diebold–Mariano statistic against the random-walk benchmark, with negative values favouring the model. All forecasts are floored at one per cent of the in-sample variance to prevent QLIKE from diverging when the proxy approaches zero, and denote significance at the 1, 5 and 10 per cent levels. Mean absolute percentage error is omitted; see text.*

Every GARCH variant improves on the random-walk benchmark, with Theil's U between 0.740 and 0.766, corresponding to reductions in root mean squared error of roughly one quarter. The Diebold Mariano tests confirm that these improvements are statistically significant at the five per cent level for the dollar, sterling and the euro, and at the ten per cent level for the yen. The QLIKE criterion gives the same ordering. The conditional variance models therefore possess genuine predictive content relative to a naive benchmark.

A more demanding comparison is less flattering. The 22-day rolling sample variance a calculation requiring no estimation whatever performs almost identically to the fitted GARCH models. For the dollar it achieves a lower mean absolute error than either; for sterling its mean absolute error is the lowest of all four forecasts. Only on the QLIKE criterion, and only modestly, do the GARCH models establish a consistent advantage. This finding is neither novel nor embarrassing: it is consistent with the broader forecasting literature, in which simple variance estimators are difficult to beat at short horizons, and it reflects the limited information content of a one-step-ahead forecasting exercise. The value of the GARCH framework in this paper lies chiefly in the structural parameters persistence, asymmetry, tail thickness that the rolling variance cannot deliver, rather than in a decisive forecasting advantage.

Mean absolute percentage error, though specified in the original research design, is not reported. Against a squared-return proxy that frequently approaches zero, the measure produced values in the thousands and occasionally hundreds of thousands of per cent, and is dominated entirely by observations for which the denominator is near zero. This is a well-understood property of relative loss functions applied to latent volatility (Patton, 2011) and not a feature of the models. The robust alternatives, squared error and QLIKE, are reported in their place.

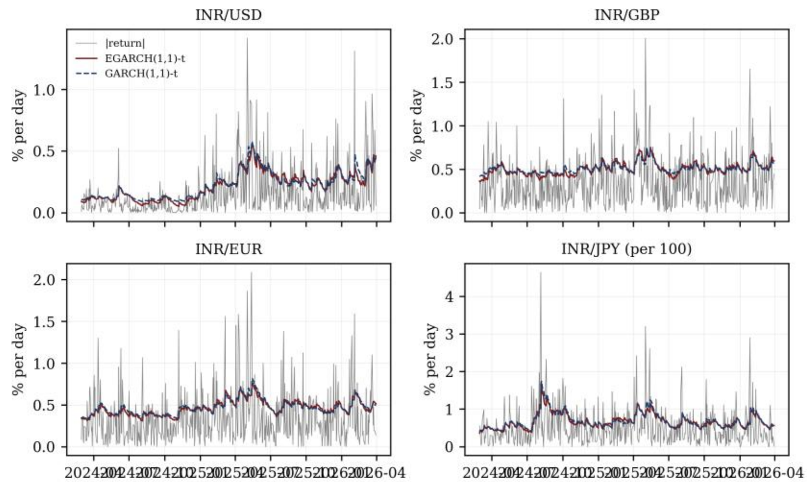

Figure 8. Out-of-Sample Conditional Volatility Forecasts against Absolute Returns

Note: *Forecasts are one-step-ahead, generated from parameters estimated on the first 3,361 observations only. The grey series is the absolute return.*

V. DISCUSSION

A) The exceptional stability of the rupee against the dollar

The most economically consequential finding is that the rupee is far more stable against the dollar than against sterling, the euro or the yen, and that the gap has widened markedly since 2023. Rupee and dollar volatility of 3.46 per cent annualised in the post-pandemic period is low by the standards of any floating currency and is remarkable for an emerging-market economy during a period of substantial global monetary adjustment.

Two mechanisms are consistent with this pattern, and the data presented here cannot separate them. The first is official intervention. The Reserve Bank of India operates a managed float and has accumulated substantial foreign exchange reserves; a policy of leaning against short-horizon fluctuations would compress realised volatility while leaving the trend path of the currency largely unaffected which is what the data show, since the rupee depreciated by roughly 112 per cent against the dollar over the sample while exhibiting the lowest volatility of the four pairs. The second is that dollar invoicing of Indian trade generates a natural two-way order flow in the rupee and dollar market that is absent in the thinner cross-rates, so that greater market depth alone would produce lower volatility.

The near-integrated persistence of rupee and dollar variance is consistent with the first mechanism in a specific way. If the authority damps volatility but permits shocks that survive intervention to pass through, the resulting variance process would exhibit exactly the combination observed here: a low unconditional level with slow decay of whatever variance is realised. This is an interpretation rather than a test, and establishing it would require intervention data that the present study does not employ.

B) Implications for the monetary authority

If the low volatility of the rupee and dollar rate is substantially policy-induced, its cost warrants consideration. Suppressed volatility discourages the development of hedging markets, since firms perceive limited need to hedge an exposure that appears stable, and may therefore concentrate risk on balance sheets that are unprepared for a discrete adjustment. The very high persistence of rupee–dollar variance means that when volatility does rise, it remains elevated for a long time. The combination of ordinarily low but highly persistent volatility is a configuration in which under-hedging is both likely and costly.

The finding that persistence falls materially when the sample is partitioned has a direct operational implication. Full-sample estimates of persistence, of the kind routinely reported in the applied literature, overstate the duration for which a variance shock will elevate risk. Institutions calibrating value-at-risk models or capital buffers on such estimates will systematically misjudge the horizon of exposure. Recursive or rolling estimation, or explicit regime-switching, is preferable to a single set of parameters imposed on a long sample.

C) Implications for exporters, importers and investors

The comparative rankings translate into a straightforward hedging prescription. Yen exposure carries the greatest risk on every measure the highest average and peak volatility, the most variable volatility and the heaviest tails, with an estimated ν of 5.50 implying that extreme outcomes are materially more probable than a normal approximation would suggest. Firms with yen liabilities should hedge more aggressively than the nominal size of the exposure alone would indicate, and value-at-risk calculations for yen positions should not assume normality.

The persistence estimates inform hedge tenor rather than hedge ratio. A shock to sterling volatility is half dissipated within fifteen trading days, so short-dated hedges may be rolled with limited risk of finding the environment unchanged at renewal. Euro and yen volatility shocks require six to nine weeks to decay by half, favouring longer tenors. Rupee and dollar volatility, once elevated, shows no reliable tendency to revert within any horizon the data can identify.

The asymmetry results carry an additional implication for firms with dollar or yen liabilities. Because depreciation shocks raise volatility more than appreciations, the periods in which the rupee value of foreign-currency obligations rises are also periods of elevated uncertainty about how much further they may rise. Risk to such firms is compounded rather than merely directional, and options-based strategies that pay off in exactly those states may be more appropriate than linear forward cover.

D) Interpretation of persistence

A methodological point deserves emphasis. Near-unit persistence estimates are frequently reported in the applied volatility literature and interpreted as evidence of long memory. The evidence assembled here counsels caution. The ICSS algorithm identifies 13 to 19 variance regimes per series; persistence falls in every currency when the sample is partitioned; and for the rupee–dollar rate the persistence estimate itself is not robust to the choice between GARCH and EGARCH. Taken together, these results suggest that a substantial part of the apparent persistence is attributable to neglected variance regime change, exactly as Lamoureux and Lastrapes (1990) argued. Reporting a full-sample persistence estimate of 0.98 without examining its stability across sub-periods risks conflating two distinct phenomena.

VI. CONCLUDING REMARKS

This study models and compares the conditional volatility of the Indian rupee against the US dollar, pound sterling, euro and Japanese yen using sixteen years of daily data. The results confirm non-normal and highly leptokurtic return distributions, significant ARCH effects and substantial volatility persistence across all four exchange rates. Student-*t* innovations consistently outperform Gaussian specifications, while GARCH-family models provide a better representation of volatility dynamics than pure ARCH models. EGARCH(1,1)-*t* is preferred for the dollar and yen, whereas GARCH(1,1)-*t* is preferred for sterling and the euro. The Japanese yen emerges as the most volatile counterpart currency, followed by sterling and the euro, while the dollar is substantially less volatile. Volatility persistence is also pronounced, although it declines across all currencies when the sample is divided into pre-pandemic, pandemic and post-pandemic periods, indicating that unmodelled regime changes can inflate full-sample persistence estimates. Out-of-sample results show that GARCH models improve upon a random-walk volatility benchmark, but do not consistently outperform a simple rolling variance measure. These findings should nevertheless be interpreted with caution because the estimation routines require further independent replication, residual ARCH effects remain, squared returns provide only a noisy measure of daily volatility, and the analysis does not model cross-currency spillovers or directly test the role of official foreign-exchange intervention. Future research could therefore employ regime-switching and multivariate GARCH models, incorporate Reserve Bank intervention and foreign-exchange reserve data, and use intraday-based realised volatility to improve the modelling and evaluation of exchange-rate volatility.

Interest Conflicts

The author(s) declare(s) that there is no conflict of interest regarding the publication of this paper.

Funding Statement

This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

VII. REFERENCES

[1] Andersen, T. G., & Bollerslev, T. (1998). Answering the sceptics: Yes, standard volatility models do provide accurate forecasts. *International Economic Review*, 39(4), 885–905.

[2] Anwer, Z., Naeem, M. A., Hassan, M. K., & Karim, S. (2022). Asymmetric connectedness across Asia-Pacific currencies: Evidence from time-frequency domain analysis. *Finance Research Letters*, 47, 102782. https://doi.org/10.1016/j.frl.2021.102782.

[3] Baillie, R. T., & Bollerslev, T. (1989). The message in daily exchange rates: A conditional-variance tale. *Journal of Business & Economic Statistics*, 7(3), 297–305.

[4] Bollerslev, T. (1986). Generalised autoregressive conditional heteroscedasticity. *Journal of Econometrics*, 31(3), 307–327.

[5] Bollerslev, T. (1987). A conditionally heteroskedastic time series model for speculative prices and rates of return. *Review of Economics and Statistics*, 69(3), 542–547.

[6] Bollerslev, T., & Wooldridge, J. M. (1992). Quasi-maximum likelihood estimation and inference in dynamic models with time-varying covariances. *Econometric Reviews*, 11(2), 143–172.

[7] Bouri, E., Lucey, B., Saeed, T., & Vo, X. V. (2020). Extreme spillovers across Asian-Pacific currencies: A quantile-based analysis. *International Review of Financial Analysis*, 72, 101605. https://doi.org/10.1016/j.irfa.2020.101605.

[8] Das, S., & Roy, S. S. (2023). Following the leaders? A study of co-movement and volatility spillover in BRICS currencies. *Economic Systems*, 47(2), 100980. https://doi.org/10.1016/j.ecosys.2022.100980.

[9] 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.

[10] Diebold, F. X., & Mariano, R. S. (1995). Comparing predictive accuracy. *Journal of Business & Economic Statistics*, 13(3), 253–263.

[11] Engle, R. F. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica*, 50(4), 987–1007.

[12] Engle, R. F., & Bollerslev, T. (1986). Modelling the persistence of conditional variances. *Econometric Reviews*, 5(1), 1–50.

[13] Fama, E. F. (1965). The behaviour of stock-market prices. *The Journal of Business*, 38(1), 34–105.

[14] Glosten, L. R., Jagannathan, R., & Runkle, D. E. (1993). On the relation between the expected value and the volatility of the nominal excess return on stocks. *The Journal of Finance*, 48(5), 1779–1801.

[15] Goyal, A., & Arora, S. (2012). The Indian exchange rate and Central Bank action: An EGARCH analysis. *Journal of Asian Economics*, 23(1), 60–72. https://doi.org/10.1016/j.asieco.2011.09.001.

[16] Hansen, P. R., & Lunde, A. (2005). A forecast comparison of volatility models: Does anything beat a GARCH(1,1)? *Journal of Applied Econometrics*, 20(7), 873–889.

[17] He, S., Cheng, Z., Wang, W., & Luo, Z. (2024). What drives currency connectedness? Evidence from the BRICS currencies. *Applied Economics*, 57(1), 67–85. https://doi.org/10.1080/00036846.2024.2333712.

[18] Inclán, C., & Tiao, G. C. (1994). Use of cumulative sums of squares for retrospective detection of changes of variance. *Journal of the American Statistical Association*, 89(427), 913–923.

[19] Jarque, C. M., & Bera, A. K. (1980). Efficient tests for normality, homoscedasticity and serial independence of regression residuals. *Economics Letters*, 6(3), 255–259.

[20] Kumar, M. (2013). Returns and volatility spillover between stock prices and exchange rates: Empirical evidence from IBSA countries. *International Journal of Emerging Markets*, 8(2), 108–128. https://doi.org/10.1108/17468801311306984.

[21] Kwiatkowski, D., Phillips, P. C. B., Schmidt, P., & Shin, Y. (1992). Testing the null hypothesis of stationarity against the alternative of a unit root. *Journal of Econometrics*, 54(1–3), 159–178.

[22] Lamoureux, C. G., & Lastrapes, W. D. (1990). Persistence in variance, structural change, and the GARCH model. *Journal of Business & Economic Statistics*, 8(2), 225–234.

[23] Ljung, G. M., & Box, G. E. P. (1978). On a measure of lack of fit in time series models. *Biometrika*, 65(2), 297–303.

[24] MacKinnon, J. G. (1996). Numerical distribution functions for unit root and cointegration tests. *Journal of Applied Econometrics*, 11(6), 601–618.

[25] Mahalwala, R. (2022). Analysing exchange rate volatility in India using GARCH family models. *SN Business & Economics*, 2, Article 134. https://doi.org/10.1007/s43546-022-00317-z.

[26] Mandelbrot, B. (1963). The variation of certain speculative prices. *The Journal of Business*, 36(4), 394–419.

[27] Meese, R. A., & Rogoff, K. (1983). Empirical exchange rate models of the seventies: Do they fit out of sample? *Journal of International Economics*, 14(1–2), 3–24.

[28] Nelson, D. B. (1991). Conditional heteroskedasticity in asset returns: A new approach. *Econometrica*, 59(2), 347–370.

[29] Patnaik, I., & Shah, A. (2001). Volatility in the foreign exchange market: A study of selected currencies. *Economic and Political Weekly*, 36(35), 3307–3311.

[30] Patnaik, I., & Shah, A. (2010). Does the currency regime shape unhedged currency exposure? *Journal of International Money and Finance*, 29(5), 760–769. https://doi.org/10.1016/j.jimonfin.2009.12.007.

[31] Patton, A. J. (2011). Volatility forecast comparison using imperfect volatility proxies. *Journal of Econometrics*, 160(1), 246–256.

[32] Perron, P. (1989). The great crash, the oil price shock, and the unit root hypothesis. *Econometrica*, 57(6), 1361–1401.

[33] Phillips, P. C. B., & Perron, P. (1988). Testing for a unit root in time series regression. *Biometrika*, 75(2), 335–346.

[34] Said, S. E., & Dickey, D. A. (1984). Testing for unit roots in autoregressive-moving average models of unknown order. *Biometrika*, 71(3), 599–607.

[35] Zivot, E., & Andrews, D. W. K. (1992). Further evidence on the great crash, the oil-price shock, and the unit-root hypothesis. *Journal of Business & Economic Statistics*, 10(3), 251–270.

Appendix A. Estimation Details and Validation

Critical values for the Dickey Fuller and Phillips Perron statistics were obtained by simulating 20,000 random walks at the relevant sample size and deterministic specification. The resulting five per cent critical value for the constant-only case at $T = 1{,}000$ is −2.853, against MacKinnon's published −2.86; the constant-and-trend value is −3.422 against a published −3.41. *P*-values are the empirical quantiles of the simulated distribution.

The GARCH estimator was validated on 4,000 observations simulated from a GARCH(1,1) process with $\omega = 0.020$, $\alpha = 0.090$, $\beta = 0.880$ and $\mu = 0.010$. Estimated values were $\hat{\omega} = 0.0195$, $\hat{\alpha} = 0.0744$, $\hat{\beta} = 0.8931$ and $\hat{\mu} = -0.0035$, each within 1.4 estimated standard errors of its true value.

The likelihood profile underlying the conclusion that rupee–dollar persistence is not bounded away from unity is reported below. The log-likelihood was maximised subject to a sequence of ceilings on $\alpha+\beta$; the constraint binds at every ceiling, and the likelihood rises monotonically, establishing that the unconstrained optimum lies at or above unity.

Table A1: Profile likelihood for INR/USD GARCH(1,1)-

Ceiling on $\alpha+\beta$	Log-likelihood	ω	α	β	$\alpha+\beta$ at optimum
0.950	−1126.02	0.00430	0.1570	0.7930	0.95000
0.970	−1098.80	0.00213	0.1215	0.8483	0.96973
0.980	−1082.29	0.00121	0.1027	0.8773	0.98000
0.990	−1065.78	0.00067	0.0997	0.8903	0.99000
0.995	−1057.27	0.00045	0.0951	0.8999	0.99500
0.999	−1053.84	0.00033	0.0975	0.9015	0.99900
0.9999	−1050.96	0.00031	0.0959	0.9040	0.99990
0.99999	−1050.43	0.00029	0.0944	0.9056	0.99999

Notes: *The constraint binds at every ceiling, and the log-likelihood increases monotonically, indicating that the unconstrained maximum-likelihood estimate of persistence is not interior to the stationarity region.*