

Financial Inclusion and Investment Uncertainty in the Indian Stock Market: Evidence from Implied Volatility, Regime Shifts and Retail Participation, 2011–2026

Dr B Amarnath Reddy^{1*}, Dr CH V L L Kusuma Kumari²

¹Department of Finance & Marketing, Vignana Jyothi Institute of Management, Hyderabad

²Department of Humanities & Sciences (Mgt), VNR Vignana Jyothi Institute of Engineering and Technology, Hyderabad

***Corresponding Author**

Dr B Amarnath Reddy

Department of Finance & Marketing, Vignana Jyothi Institute of Management, Hyderabad

ABSTRACT

Purpose – India has added more than 18 crore demat accounts since 2019 and its composite Financial Inclusion Index (FI-Index) has climbed from 53.9 to 70.0 in five years. This study asks whether this unprecedented widening of market access has changed the level, persistence, asymmetry and transmission of investment uncertainty in the Indian equity market.

Design/methodology/approach – The study assembles 819 weekly observations (January 2011–September 2026) on the NIFTY 50, the India VIX and a realised-volatility measure built from daily returns, and divides the sample into three inclusion phases anchored on the launch of the Pradhan Mantri Jan Dhan Yojana (August 2014) and the post-pandemic retail surge (April 2020). The empirical design combines GJR-GARCH models, heterogeneous autoregressive (HAR) regressions of realised volatility with out-of-sample tests, asymmetric and quantile regressions of the return–implied-volatility relation, a vector autoregression with Granger-causality and variance-decomposition analysis, multiple structural-break tests and a two-state Markov-switching model. A fiscal-year panel links these measures to the RBI FI-Index and depository statistics.

Findings – Uncertainty in Indian equities is strongly asymmetric: negative shocks raise conditional variance ($\gamma = 0.240$, $p < 0.001$) while positive shocks of equal size do not. Implied volatility carries information about future realised volatility beyond its own history, raising adjusted R^2 from 0.245 to 0.339 and lowering out-of-sample forecast error (Diebold–Mariano = 2.72, $p = 0.006$). Although average volatility and the incidence of turbulent regimes were lower in the two inclusion phases than in the pre-mission phase, a likelihood-ratio test cannot reject stability of the volatility process across phases (LR = 11.18, $p = 0.514$), and phase effects vanish once implied volatility is controlled for. Structural breaks in volatility coincide with global shocks rather than with inclusion milestones. Annual demat growth is highest in years of strong returns and after high-uncertainty years.

Originality/value – The paper is among the first to connect India’s official inclusion metrics with a multi-model, fifteen-year assessment of market uncertainty that covers the entire post-2020 participation boom and the 2026 geopolitical sell-off. It shows that inclusion has broadened the investor base without destabilising volatility dynamics, but that the new cohort is exposed to asymmetric downside risk whose pricing remains dominated by global conditions.

Keywords: Financial inclusion; investment uncertainty; India VIX; GJR-GARCH; realised volatility; retail investors; Markov switching; NIFTY 50

INTRODUCTION

Financial inclusion was, for most of its policy history, a story about bank accounts, remittances and small loans. In India the story has changed. After the Pradhan Mantri Jan Dhan Yojana (PMJDY) opened basic accounts for hundreds of millions of households from August 2014, the combination of biometric identity, electronic know-your-customer checks, the Unified Payments Interface and low-cost digital brokerage lowered the cost of owning listed securities to almost nothing. The number of demat accounts, which stood at about 3.6 crore in March 2019,

reached 22.5 crore by March 2026 (Press Information Bureau [PIB], 2021; Securities and Exchange Board of India [SEBI], 2026). Over the same period the Reserve Bank of India’s composite FI-Index, which aggregates access, usage and quality across banking, investment, insurance, postal and pension services, rose from 53.9 in March 2021 to 70.0 in March 2026 (Reserve Bank of India [RBI], 2021; Business Standard, 2026). Investment, in other words, has become one of the channels through which Indian households are being included in the formal financial system.

This shift raises a question that the inclusion literature has seldom addressed directly: what happens to market uncertainty when millions of first-time investors enter an equity market in a few years? Two opposing arguments can be made. A broader and more diverse investor base may add liquidity, dilute the influence of any single class of trader and anchor domestic demand for equities, which could dampen volatility and shorten episodes of stress (Mehrotra & Yetman, 2015; Han & Melecky, 2013). On the other hand, new investors tend to have less experience, trade on attention and sentiment, and respond more strongly to losses (Barber & Odean, 2000, 2008; Baker & Wurgler, 2006). If their participation is concentrated in rising markets and their exit concentrated in falling ones, inclusion could make uncertainty more asymmetric and more persistent. The evidence from the United States on app-based retail trading during the pandemic suggests that both effects are possible (Ozik et al., 2021; Welch, 2022; Barber et al., 2022).

India offers an unusually clean setting in which to study this tension. The expansion of market access followed identifiable policy milestones, it was large relative to the existing investor base, and it happened in a market that has a liquid, exchange-computed implied-volatility index—the India VIX—covering the whole period. This allows investment uncertainty to be measured in two complementary ways: as the forward-looking expectation of volatility priced in NIFTY options, and as the backward-looking volatility actually realised in returns.

The study makes four contributions. First, it constructs a fifteen-year weekly data set (2011–2026) that spans the pre-inclusion era, the account-opening phase and the entire post-2020 retail boom, including the market stress of March 2026, and links it to official FI-Index and depository data. Second, rather than relying on a single volatility model, it triangulates evidence from conditional-volatility, realised-volatility, quantile, VAR, structural-break and regime-switching frameworks, so that conclusions do not depend on one specification. Third, it tests formally whether the parameters that govern uncertainty—its level, persistence, asymmetry and link to implied volatility—changed across inclusion phases. Fourth, it documents the co-movement between annual investor-participation growth and market conditions, which speaks to the behaviour of the newly included cohort. The main findings can be stated briefly. Uncertainty in the Indian market is dominated by negative news, is moderately persistent with a half-life of about six weeks, and is well anticipated by the India VIX. Average volatility and the frequency of turbulent weeks were lower after 2014 than before, but the underlying dynamics are statistically stable across phases, and the dates at which volatility shifted correspond to global shocks rather than to domestic inclusion events. Participation growth, finally, is procyclical with returns and rises after episodes of high fear, which is consistent with “buy-the-dip” entry by new investors. The remainder of the paper reviews the literature and develops hypotheses (Section 2), describes the data and methods (Section 3), reports results (Section 4), discusses implications (Section 5) and concludes (Section 6).

2. LITERATURE REVIEW AND HYPOTHESES

2.1 Financial inclusion: from access to investment

Early work defined financial inclusion as the ease of access to, availability of and use of formal financial services, and proposed composite indices to measure it across countries (Sarma, 2008; Sarma & Pais, 2011). The developmental case rested on evidence that deeper finance reduces poverty and inequality (Beck et al., 2007). The World Bank’s Global Findex surveys later shifted attention from account ownership to account use and to digital payments, noting that ownership alone overstates effective inclusion (Demirgüç-Kunt et al., 2022). India’s FI-Index, first published in 2021, follows this broader view: it weights access at 35 per cent, usage at 45 per cent and quality at 20 per cent, and it explicitly includes the investment sector (RBI, 2021).

A separate literature in household finance explains why ownership of risky assets remains limited even when access is available. Participation is associated with financial literacy (van Rooij et al., 2011; Lusardi & Mitchell, 2014), trust in markets (Guiso et al., 2008) and the fixed costs of learning and transacting (Campbell, 2006). In India, surveys of young urban workers show substantial gaps in basic financial knowledge (Agarwalla et al., 2015). Digital onboarding and discount brokerage reduce transaction and information costs sharply, which is one reason why the recent jump in Indian participation has been so fast (Ozili, 2018).

2.2 Inclusion, stability and investor behaviour

The link between inclusion and stability is debated. Broad-based deposit funding tends to be more stable in crises (Han & Melecky, 2013), and central banks have argued that inclusion can strengthen monetary transmission and financial resilience if it is accompanied by consumer protection (Mehrotra & Yetman, 2015). For equity markets, however, the relevant mechanism is investor behaviour. Individual investors trade excessively, are drawn to attention-grabbing stocks and tend to under-diversify (Barber & Odean, 2000, 2008). During the COVID-19 lockdown, retail trading rose around the world (Chiah & Zhong, 2020) and in the United States it supplied liquidity at a time of stress (Ozik et al., 2021). Yet herding by app-based investors into attention-grabbing stocks was followed by negative returns (Barber et al., 2022), even though the aggregate portfolio of such investors was not clearly inferior (Welch, 2022). Investor sentiment more broadly moves prices of hard-to-value assets (Baker & Wurgler, 2006). These findings imply that the effect of inclusion on uncertainty is an empirical question whose answer may differ by market and period.

2.3 Measuring investment uncertainty

Uncertainty shocks depress investment and hiring (Bloom, 2009), and news-based measures of policy uncertainty affect asset prices (Baker et al., 2016). In equity markets, uncertainty is typically captured by the conditional variance of returns and by option-implied volatility. The ARCH and GARCH families model volatility clustering (Engle, 1982; Bollerslev, 1986), and asymmetric variants capture the tendency of bad news to raise volatility more than good news (Black, 1976;

Nelson, 1991; Glosten et al., 1993; Engle & Ng, 1993). Despite their simplicity, these models are hard to beat in forecasting comparisons (Hansen & Lunde, 2005; Andersen & Bollerslev, 1998). Realised volatility computed from higher-frequency returns provides a model-free benchmark, and the HAR model of Corsi (2009) captures its long-memory behaviour with a parsimonious cascade of daily, weekly and monthly components. Implied-volatility indices such as the VIX summarise the market's expectation of volatility over the next month and have been described as an "investor fear gauge" (Whaley, 2000). They contain information about future volatility beyond historical measures (Fleming et al., 1995; Christensen & Prabhala, 1998), although part of the index reflects a variance risk premium rather than expected volatility (Bekaert & Hoerova, 2014). Their relation with returns is negative and asymmetric: implied volatility rises more after market falls than it declines after rallies (Giot, 2005). Studies of the India VIX report similar patterns, a strong negative contemporaneous relation with NIFTY returns and a forward-looking component (Chandra & Thenmozhi, 2015; Shaikh & Padhi, 2015). What these studies do not examine is whether the relation has changed as India's investor base has broadened.

2.4 Research gap and hypotheses

Three gaps emerge. First, work on financial inclusion rarely uses market-uncertainty outcomes; its focus is on poverty, banking stability or credit. Second, studies of Indian volatility mostly end before the post-2020 retail surge or treat it as a COVID-19 episode rather than as a structural change in who owns equities. Third, few studies test for parameter stability across inclusion phases while controlling for the global shocks that dominate emerging-market volatility. The study addresses these gaps with the following hypotheses.

- **H1** (asymmetric clustering): Investment uncertainty in the Indian equity market exhibits volatility clustering and responds more strongly to negative than to positive return shocks.
- **H2** (informational role of implied volatility): The India VIX contains information about future realised volatility beyond that in past realised volatility, both in-sample and out-of-sample.
- **H3** (inclusion and volatility dynamics): The level, persistence and asymmetry of volatility differ between the pre-mission phase and the inclusion phases.
- **H4** (inclusion and the fear response): The sensitivity of implied volatility to market declines changed as retail participation deepened.
- **H5** (participation and market conditions): Growth in investor participation and in the FI-Index co-moves with returns and uncertainty at the annual frequency.

3. DATA AND METHODOLOGY

3.1 Data and sources

Daily closing values of the NIFTY 50 index and the India VIX, both computed by NSE Indices Ltd, were retrieved

from the Yahoo Finance historical-data service for 3 January 2011 to 21 September 2026. The India VIX is calculated from the order book of near- and next-month NIFTY options and expresses expected annualised volatility over the following 30 calendar days. Daily data were aggregated to weeks by taking the last trading day of each calendar week, which reduces the microstructure noise and day-of-week effects present in daily data while preserving enough observations for time-series inference. The final partial week (21 September 2026) was dropped, leaving 819 weekly returns from 14 January 2011 to 18 September 2026. The underlying daily return series (3,855 observations) has a mean of 0.035 per cent, a standard deviation of 1.04 per cent, skewness of -0.92 and kurtosis of 17.1; its extreme values are -13.9 per cent on 23 March 2020 and +8.4 per cent on 7 April 2020. Annual financial-inclusion data were collected from official releases. The FI-Index values for March 2021 (53.9), March 2022 (56.4), March 2023 (60.1), March 2024 (64.2), March 2025 (67.0) and March 2026 (70.0) are taken from RBI press releases and their official reporting (RBI, 2021, 2022; PIB, 2025; Business Standard, 2026). The index was not published for 2018–2020. Total demat accounts at the end of each fiscal year come from the Ministry of Finance for FY2019–FY2021 (PIB, 2021), from NSDL and CDSL statistics as reported for FY2022 (EquityPandit, 2022) and FY2023–FY2024 (Deccan Chronicle, 2024), and from the SEBI Annual Report 2025–26 for FY2025 and FY2026 (SEBI, 2026). All values are reported in Table 13 so that the data can be reproduced.

3.2 Variable construction

The weekly log return is $r_t = 100 \times \ln(P_t / P_{t-1})$, where P_t is the NIFTY 50 close. Weekly realised volatility is constructed from the daily log returns $r_{t,d}$ within week t and annualised:

$$RV_t = [52 \times \sum_d r_{t,d}^2]^{1/2} \quad (1)$$

The implied-volatility measure is the India VIX at the end of week t (VIX_t) and its weekly change ΔVIX_t . Logarithms of RV and VIX are used in regressions to reduce skewness. A COVID-19 crash dummy equals one for the weeks from 24 February to 31 May 2020.

3.3 Inclusion phases

Because high-frequency inclusion data do not exist, inclusion enters the market-level analysis through three policy-defined phases (Table 1). Phase I ends with the launch of PMJDY on 28 August 2014. Phase II covers the build-out of access infrastructure—the JAM trinity of Jan Dhan accounts, Aadhaar and mobile phones, Aadhaar-based e-KYC for securities accounts and the launch of UPI in 2016—and ends in March 2020. Phase III, from April 2020, is the period of participation deepening, in which demat accounts grew more than five-fold. The phase boundaries are set by policy events, not estimated from the data, which avoids choosing breaks that fit volatility by construction. Data-driven breaks are estimated separately (Section 3.8) and compared with these boundaries.

Table 1. Financial-inclusion phases used in the analysis

Phase	Window	Weeks	Inclusion milestones defining the phase
I. Pre-mission	14 Jan 2011 – 22 Aug 2014	189	Branch-led, supply-side inclusion; no national account-opening mission
II. Access build-out	29 Aug 2014 – 27 Mar 2020	292	PMJDY (Aug 2014), JAM trinity, Aadhaar e-KYC, UPI (2016), demat 3.6 crore by FY19
III. Participation deepening	3 Apr 2020 – 18 Sep 2026	338	Digital-only onboarding, discount brokers; demat 4.1 → 22.5 crore; FI-Index 53.9 → 70.0

Source: Author's classification based on RBI, SEBI and Ministry of Finance documents. PMJDY = Pradhan Mantri Jan Dhan Yojana; JAM = Jan Dhan–Aadhaar–Mobile; UPI = Unified Payments Interface.

3.4 Asymmetric conditional volatility

To test H1 and H3, weekly returns are modelled with a GJR–GARCH(1,1) specification (Glosten et al., 1993) with Student-t innovations:

$$r_t = \mu + \varepsilon_t, \quad \varepsilon_t = \sigma_t z_t, \quad z_t \sim t_\nu \quad (2)$$

$$\sigma_t^2 = \omega + \alpha \varepsilon_{t-1}^2 + \gamma \varepsilon_{t-1}^2 I(\varepsilon_{t-1} < 0) + \beta \sigma_{t-1}^2 \quad (3)$$

A positive γ indicates that negative shocks raise future variance more than positive shocks. Persistence is $\alpha + \beta + \gamma/2$ and the half-life of a shock is $\ln(0.5)/\ln(\alpha + \beta + \gamma/2)$. The model is estimated for the full sample and separately for each phase; a likelihood-ratio test compares the sum of the phase log-likelihoods with the full-sample log-likelihood to test parameter stability. The GJR model is compared with symmetric GARCH(1,1) and EGARCH(1,1) (Nelson, 1991) using information criteria.

3.5 Realised volatility and implied volatility

H2 is examined with a HAR model (Corsi, 2009) adapted to weekly data, in which log realised volatility depends on its weekly, monthly (four-week) and quarterly (thirteen-week) averages, augmented with lagged log VIX:

$$\ln RV_t = \beta_0 + \beta_w \ln RV_{t-1} + \beta_m \ln RV_{t-1:t-4} + \beta_q \ln RV_{t-1:t-13} + \delta \ln VIX_{t-1} + \theta' X_t + u_t \quad (4)$$

The vector X_t contains a leverage term $|r_{t-1}| \times 1(r_{t-1} < 0)$, the COVID-19 dummy, phase dummies and, in the most general model, interactions of lagged log VIX with the phase dummies. Inference uses Newey–West standard errors with eight lags (Newey & West, 1987). Out-of-sample accuracy is assessed by estimating the models on data up to December 2021 and forecasting the 246 weeks from January 2022; the Diebold–Mariano (1995) test compares squared forecast errors with a HAC variance.

3.6 The return–implied-volatility relation

To test H4, weekly changes in VIX are regressed on the positive and negative parts of the contemporaneous return, following Giot (2005):

$$\Delta VIX_t = a + b_1 r_t^+ + b_2 r_t^- + c VIX_{t-1} + \sum_k (\varphi_{1k} r_t^+ + \varphi_{2k} r_t^-) D_k + e_t \quad (5)$$

where $r_t^+ = \max(r_t, 0)$, $r_t^- = \min(r_t, 0)$ and D_k are the Phase II and Phase III dummies. A more negative b_2 than b_1 indicates that falls raise fear more than rallies reduce it. Because this relation may differ in the tails of the return distribution, quantile regressions (Koenker & Bassett, 1978) of returns on ΔVIX are estimated at $\tau = 0.05, 0.25, 0.50, 0.75$ and 0.95 , for the full sample and by phase.

3.7 Dynamic interactions

A three-variable VAR (Sims, 1980) in weekly returns, ΔVIX and $\Delta \ln RV$ is used to study lead–lag relations. The lag order is selected by the Bayesian information criterion. Granger (1969) causality tests are computed for the full sample and by phase, and orthogonalised impulse responses and forecast-error variance decompositions summarise the dynamics.

3.8 Structural breaks and regimes

To examine whether the data locate shifts in uncertainty at inclusion milestones, the study estimates multiple breaks in the mean of log RV and log VIX using the global least-squares procedure of Bai and Perron (1998, 2003), with up to six breaks, 10 per cent trimming and BIC selection. Finally, a two-state Markov-switching model (Hamilton, 1989) with regime-specific means and variances classifies each week as calm or turbulent:

$$r_t = \mu(s_t) + \sigma(s_t) \eta_t, \quad s_t \in \{\text{calm, turbulent}\}, \quad \Pr(s_t = j | s_{t-1} = i) = p_{ij} \quad (6)$$

The share of weeks in which the smoothed probability of the turbulent regime exceeds 0.5 is compared across phases.

3.9 Annual inclusion panel

For H5, weekly market measures are aggregated to fiscal years ending in March, matching the reporting period of the FI-Index and depository data. With at most eight annual observations, formal regression is not appropriate; the analysis therefore reports Spearman rank correlations as descriptive evidence and interprets them cautiously.

4. Empirical results

4.1 Descriptive statistics and preliminary tests



Figure 1. NIFTY 50 and India VIX, weekly, January 2011–September 2026

Shaded areas mark Phase I (to August 2014) and Phase III (from April 2020); the unshaded area is Phase II. Dashed lines are the four structural breaks in realised volatility estimated in Section 4.6. Source: NSE Indices Ltd via Yahoo Finance; author's calculations. Figure 1 shows the long upward trend of the NIFTY 50, which rose roughly four-fold over the sample, and the episodic nature of implied volatility. The India VIX spent most of the period between 10 and 25, with pronounced spikes in 2011, 2013, the 2014 election, early 2016, March 2020 (when it reached 70.4 on a weekly basis) and March 2026, when a West Asian conflict and an oil-price shock pushed it to 26.8.

Table 2 reports descriptive statistics. Weekly returns averaged 0.168 per cent with a standard deviation of 2.19 per cent; they are negatively skewed (-0.23) and fat-tailed (kurtosis 6.45), and the Jarque–Bera test rejects normality decisively. Mean VIX fell from 20.3 in Phase I to 15.9 in Phase II and was 16.6 in Phase III; realised volatility followed the same pattern (15.9, 12.9 and 13.2 per cent). Kruskal–Wallis tests confirm that the distributions of VIX ($H = 120.1$, $p < 0.001$) and realised volatility ($H = 44.7$, $p < 0.001$) differ across phases, whereas mean returns do not ($H = 1.81$, $p = 0.404$). Phase II has the most extreme distribution—skewness of -1.42 and kurtosis of 9.67—because it contains the COVID-19 crash, while Phase III has the highest mean return (0.293 per cent a week).

Table 2. Descriptive statistics of weekly returns, implied volatility and realised volatility

Series	Phase	N	Mean	SD	Min	Max	Skew	Kurt	JB
Weekly return (%)	Full	819	0.168	2.193	-12.96	11.97	-0.23	6.45	413.6***
	I	189	0.155	2.383	-5.22	6.97	0.22	2.76	2.0
	II	292	0.031	2.094	-12.96	6.28	-1.42	9.67	638.6***
	III	338	0.293	2.165	-6.89	11.97	0.34	6.56	184.9***
India VIX (%)	Full	819	17.21	5.95	9.15	70.39	3.05	20.91	12215.1***
	I	189	20.33	5.14	13.22	37.71	0.79	3.07	19.5***
	II	292	15.89	5.73	10.53	70.39	6.34	55.67	35710.5***
	III	338	16.59	5.96	9.15	55.30	2.35	12.14	1487.9***
Realised vol. (% ann.)	Full	819	13.75	8.72	1.25	115.20	4.10	36.09	39665.2***
	I	189	15.93	6.79	4.27	39.52	0.52	3.17	8.8**
	II	292	12.93	10.24	1.25	115.20	5.76	49.14	27517.3***
	III	338	13.23	8.06	2.50	67.40	2.53	13.31	1860.0***

Note: Kurt is raw kurtosis (normal = 3). JB = Jarque–Bera statistic. *** denotes rejection of normality at 1%. Phase I: Jan

2011–Aug 2014; Phase II: Aug 2014–Mar 2020; Phase III: Apr 2020–Sep 2026.

Returns show no significant linear autocorrelation (Ljung–Box $Q(10) = 7.67$, $p = 0.661$), but squared returns are strongly autocorrelated ($Q^2(10) = 341.9$, $p < 0.001$) and the ARCH-LM test with five lags is 211.5 ($p < 0.001$). These results justify conditional-heteroskedasticity models. Table 3 reports augmented Dickey–Fuller (Dickey & Fuller, 1979), Phillips–Perron (Phillips & Perron, 1988) and KPSS (Kwiatkowski et al., 1992) tests.

The log index level has a unit root while returns and ΔVIX are stationary under all three tests. Log VIX and log RV reject a unit root under ADF and Phillips–Perron but reject stationarity under KPSS at 1–5 per cent, the typical signature of a highly persistent but mean-reverting series; they are treated as stationary in levels, which is standard in the volatility-forecasting literature.

Table 3. Unit-root and stationarity tests

Series	ADF	p	Phillips–Perron	p	KPSS	p	Inference
ln NIFTY	-0.51	0.891	-0.46	0.900	4.284	<0.01	I(1)
Return	-27.83	<0.001	-27.90	<0.001	0.055	>0.10	I(0)
ln VIX	-4.73	<0.001	-5.26	<0.001	0.841	<0.01	I(0), persistent
ΔVIX	-11.72	<0.001	-27.14	<0.001	0.008	>0.10	I(0)
ln RV	-6.29	<0.001	-26.00	<0.001	0.592	0.023	I(0), persistent

Note: ADF lag length selected by AIC; all regressions include a constant. H_0 for ADF and PP is a unit root; H_0 for KPSS is stationarity. KPSS p-values are bounded by the tabulated range.

4.2 Asymmetric volatility clustering (H1, H3)

Table 4 reports GJR–GARCH estimates. In the full sample the asymmetry coefficient is $\gamma = 0.240$ ($p < 0.001$), while the symmetric ARCH coefficient α is estimated at the zero boundary. Thus, in weekly Indian data, past negative shocks—not shocks in general—drive the variance up. With $\beta = 0.776$, persistence is 0.896 and the half-life of a volatility shock is 6.3 weeks. The implied unconditional

volatility is 14.8 per cent a year, close to the sample mean of realised volatility. Standardised residuals show no remaining autocorrelation in levels ($p = 0.648$) or squares ($p = 1.000$), and no remaining ARCH effects ($p = 0.987$). The GJR model has lower AIC and BIC than both the symmetric GARCH and EGARCH models (Table 5), and the EGARCH leverage term is also negative and significant. H1 is supported.

Table 4. GJR-GARCH(1,1) estimates for the full sample and by inclusion phase

Parameter	Full sample	Phase I	Phase II	Phase III
μ (mean)	0.184*** (0.066)	0.161 (0.159)	0.200** (0.099)	0.243** (0.105)
ω (constant)	0.439*** (0.115)	0.619** (0.280)	0.438* (0.245)	0.328* (0.174)
α (ARCH)	0.000 (0.030)	0.000 (0.047)	0.000 (0.085)	0.000 (0.030)
γ (asymmetry)	0.240*** (0.061)	0.219** (0.094)	0.289** (0.119)	0.173* (0.104)
β (GARCH)	0.776*** (0.039)	0.783*** (0.064)	0.751*** (0.096)	0.820*** (0.079)
v (Student-t d.o.f.)	13.556** (6.003)	176.169*** (29.298)	5.780*** (2.061)	13.424 (8.746)
Persistence $\alpha + \beta + \gamma/2$	0.896	0.893	0.895	0.907
Half-life (weeks)	6.3	6.1	6.3	7.1
Unconditional vol. (% ann.)	14.8	17.3	14.8	13.5
Log-likelihood	-1720.9	-424.4	-589.2	-701.7
N	819	189	292	338

Note: Weekly NIFTY 50 log returns (%). Standard errors in parentheses. ***, **, * denote significance at 1%, 5% and 10%. The α estimate lies on the non-negativity boundary, so its standard error should be interpreted with caution.

Table 5. Comparison of conditional-volatility models (full sample)

Model (Student-t)	Log-likelihood	AIC	BIC	Leverage term
GARCH(1,1)	-1735.95	3481.90	3505.44	—
EGARCH(1,1)	-1722.68	3457.35	3485.60	-0.175***
GJR-GARCH(1,1)	-1720.89	3453.78	3482.03	0.240***

Note: Leverage term is γ in the EGARCH and GJR specifications. Lower AIC/BIC indicates a better fit

The phase estimates reveal changes in the distribution of shocks rather than in their propagation. Unconditional volatility falls from 17.3 per cent in Phase I to 14.8 per cent in Phase II and 13.5 per cent in Phase III, consistent with the descriptive evidence. Persistence, however, is almost identical across phases (0.893, 0.895 and 0.907). Asymmetry is largest in Phase II ($\gamma = 0.289$) and smaller and only marginally significant in Phase III ($\gamma = 0.173$, $p = 0.096$), where the GARCH coefficient is somewhat higher ($\beta = 0.820$). Tail thickness also differs: the degrees-

of-freedom parameter is very large in Phase I (returns close to normal), about 5.8 in Phase II and about 13 in Phase III. Despite these point differences, a likelihood-ratio test of equal parameters across the three phases gives $LR = 11.18$ with 12 degrees of freedom ($p = 0.514$). The volatility-generating process therefore cannot be said to have changed with inclusion; the lower average volatility of the later phases reflects smaller and less frequent shocks. Figure 2 shows that the model tracks realised volatility closely.




Figure 2. GJR-GARCH conditional volatility and realised volatility

Note: Both series annualised. The y-axis is truncated at 80% for readability; realised volatility peaks at 115% in the week ending 27 March 2020.

4.3 Informational content of the India VIX (H2)

Table 6 reports the HAR regressions. The baseline model (M1) confirms the long-memory structure of volatility: the monthly component is the most important ($\beta_m = 0.517$), and the three components together explain 24.5 per cent of the variation in log RV. Adding lagged log VIX (M2) raises the adjusted R^2 to 33.9 per cent; the VIX coefficient is 0.994 ($p < 0.001$), close to one, and it absorbs most of the explanatory power of the quarterly component. This means that option prices summarise

much of the slow-moving information that past volatility provides and add information of their own.

The out-of-sample test is more demanding. Estimated on data up to December 2021 and evaluated over the 246 weeks from January 2022 to September 2026, the model with VIX reduces the root-mean-squared forecast error from 0.438 to 0.409 and the mean absolute error from 0.334 to 0.322. The Diebold–Mariano statistic of 2.72 ($p = 0.006$) indicates that the improvement is statistically significant. H2 is supported, in line with evidence from developed markets (Fleming et al., 1995; Christensen & Prabhala, 1998).

Table 6. HAR regressions of weekly log realised volatility

Regressor	M1	M2	M3	M4
Constant	0.633*** (0.110)	-0.422** (0.174)	0.004 (0.183)	-0.152 (0.398)
$\ln RV_{t-1}$ (weekly)	0.037 (0.048)	-0.034 (0.038)	-0.093** (0.039)	-0.093** (0.039)
$\ln RV_{t-1:t-4}$ (monthly)	0.517*** (0.098)	0.241** (0.099)	0.249*** (0.095)	0.251*** (0.095)
$\ln RV_{t-1:t-13}$ (quarterly)	0.189* (0.104)	-0.161 (0.107)	-0.055 (0.093)	-0.051 (0.095)
$\ln VIX_{t-1}$		0.994*** (0.168)	0.777*** (0.134)	0.823*** (0.192)
$ r_{t-1} \times 1(r_{t-1} < 0)$			0.082*** (0.014)	0.081*** (0.015)
COVID crash dummy			0.232 (0.159)	0.248 (0.160)
Phase II			-0.023 (0.046)	0.080 (0.508)
Phase III			-0.018 (0.042)	0.217 (0.436)
$\ln VIX_{t-1} \times \text{Phase II}$				-0.033 (0.181)
$\ln VIX_{t-1} \times \text{Phase III}$				-0.081 (0.151)
Adjusted R ²	0.245	0.339	0.375	0.373
AIC	1021.4	915.0	874.0	877.7
N	806	806	806	806

Note: Newey–West (8 lags) standard errors in parentheses. ***, **, * denote significance at 1%, 5% and 10%. Out-of-sample (Jan 2022–Sep 2026, 246 weeks): RMSE M1 = 0.438, M2 = 0.409; Diebold–Mariano = 2.72 ($p = 0.007$).

Models M3 and M4 add the leverage term, the COVID-19 dummy and the phase variables. The leverage term is positive and highly significant (0.082, $p < 0.001$), confirming at the realised level the asymmetry found in the GARCH model. The Phase II and Phase III dummies are small and insignificant (-0.023 and -0.018); a joint Wald test cannot reject that both are zero ($\chi^2 = 0.26$, $p = 0.878$). In M4 the interactions of VIX with the phases are also insignificant, individually and jointly with the phase dummies ($\chi^2 = 0.59$, $p = 0.965$). Once implied volatility and the history of realised volatility are taken into account, there is no residual difference in volatility between inclusion phases. This, together with the GARCH stability test, leads to a rejection of the strong form of H3: the lower

average volatility of the inclusion phases is fully accounted for by lower expected volatility, not by a change in how volatility evolves.

4.4 The fear response to market declines (H4)

Table 7 reports the asymmetric regressions of ΔVIX . In the simplest specification (A1), a one-percentage-point fall in the NIFTY is associated with a rise of 0.73 points in the VIX, and a one-point rise with a fall of 0.41 points. Controlling for the level of VIX, which captures mean reversion, and for the COVID-19 crash (A2), the gap widens: the response to falls is 0.84 and to rises 0.25, and the Wald test of equal slopes rejects ($\chi^2 = 6.83$, $p = 0.009$). Fear rises more than three times as fast in falling markets as it subsides in rising ones.

Table 7. Asymmetric response of India VIX to NIFTY returns

Regressor	A1	A2	A3
Constant	-0.173 (0.166)	1.635*** (0.420)	2.497*** (0.496)
r^+ (positive return)	-0.414*** (0.086)	-0.247** (0.099)	-0.433*** (0.162)
r^- (negative return)	-0.726*** (0.170)	-0.842*** (0.149)	-0.510*** (0.161)
VIX_{t-1}		-0.121*** (0.033)	-0.125*** (0.023)
COVID crash dummy		2.023 (1.848)	1.736 (1.229)
Phase II			-1.165*** (0.417)
Phase III			-0.716** (0.346)
$r^+ \times \text{Phase II}$			0.302 (0.214)
$r^- \times \text{Phase II}$			-0.707** (0.311)
$r^+ \times \text{Phase III}$			0.161 (0.185)
$r^- \times \text{Phase III}$			-0.013 (0.250)
Adjusted R ²	0.301	0.363	0.400
N	818	818	818

Note: Dependent variable is the weekly change in India VIX (points). Newey–West (8 lags) standard errors in parentheses. ***, **, * denote significance at 1%, 5% and 10% Model A3 allows the slopes to differ by phase. The implied response to falls is -0.51 in Phase I, -1.22 in Phase II and -0.52 in Phase III; only the Phase II increment is significant (-0.707 , $p = 0.023$), and the Phase III interactions are jointly insignificant ($\chi^2 = 0.81$, $p = 0.668$). The heightened fear response of Phase II is linked to the extreme episodes it contains (the 2015–2016 sell-off and the onset of COVID-19). In the participation-deepening phase, the fear response returned to its pre-mission level, even though the investor base had grown several times over. H4 is therefore not supported in its directional form: deeper retail participation did not make

implied volatility more sensitive to declines.

The quantile regressions in Table 8 show that the negative relation between returns and ΔVIX holds across the entire return distribution and is strongest around the median (-0.610) and weakest in the upper tail (-0.414). The phase estimates add a nuance. In Phase III the slope is weakest in the left tail (-0.442 , $p = 0.097$) and strongest in the right tail (-0.619 , $p < 0.001$). In the most recent period, then, the sharpest weekly rallies have been accompanied by large falls in implied volatility, a pattern consistent with rapid dip-buying that compresses fear once markets turn. Figure 5 (Section 4.6) confirms that the rolling return– ΔVIX correlation has been more negative in Phase III (average -0.55) than in Phase I (-0.40).

Table 8. Quantile regressions of weekly returns on ΔVIX

Quantile τ	0.05	0.25	0.50	0.75	0.95
Full sample: $\beta(\Delta VIX)$	-0.544^{***}	-0.565^{***}	-0.610^{***}	-0.559^{***}	-0.414^{***}
(standard error)	(0.096)	(0.039)	(0.030)	(0.036)	(0.077)
Phase I: $\beta(\Delta VIX)$	-0.378^{***}		-0.549^{***}		-0.369^{**}
Phase II: $\beta(\Delta VIX)$	-0.671^{***}		-0.621^{***}		-0.485^*
Phase III: $\beta(\Delta VIX)$	-0.442^*		-0.686^{***}		-0.619^{***}

Note: Full-sample regressions include the COVID-19 dummy. ***, **, * denote significance at 1%, 5% and 10%.

4.5 Dynamic interactions

The BIC selects 3 lags for the VAR (AIC selects 6), and the estimated system is stable. Table 9 shows that ΔVIX Granger-causes returns ($F = 8.14$, $p < 0.001$) but returns do not Granger-cause ΔVIX ($p = 0.259$). This does not mean that returns do not affect fear—the contemporaneous relation in Table 7 is strong—but that, once the week's return is known, past returns add no

further information about next week's change in VIX, whereas past changes in VIX help predict the next week's return. The impulse responses in Figure 3 show that a one-standard-deviation rise in VIX is followed by a return about 0.35 percentage points lower in the next week, with the effect dissipating within three weeks. ΔVIX and past returns also Granger-cause changes in realised volatility.

Table 9. Granger-causality tests from the VAR, full sample and by phase

Null hypothesis (no Granger causality)	Full: F [p]	Phase I p	Phase II p	Phase III p
Return $\nrightarrow \Delta VIX$	1.34 [0.259]	0.675	0.682	0.185
$\Delta VIX \nrightarrow$ Return	8.14 [<0.001]	0.994	<0.001	0.430
$\Delta VIX \nrightarrow \Delta \ln RV$	4.37 [0.004]	0.103	0.491	0.030
$\Delta \ln RV \nrightarrow \Delta VIX$	3.26 [0.021]	0.042	0.032	0.915
Return $\nrightarrow \Delta \ln RV$	12.61 [<0.001]			
$\Delta \ln RV \nrightarrow$ Return	1.68 [0.169]			

Note: VAR(3) in weekly return, ΔVIX and $\Delta \ln RV$; lag order by BIC. The full-sample column reports F [p-value]; phase columns report p-values from VARs of the same order estimated within each phase. $\nrightarrow$ = does not Granger-cause.




Figure 3. Orthogonalised impulse responses from the VAR (95% bands)

Note: Cholesky ordering: return, ΔVIX , $\Delta \ln RV$. Responses to a one-standard-deviation shock.

The phase-wise tests show that the predictive role of VIX for returns is concentrated in Phase II ($p < 0.001$), while in Phase III VIX changes predict realised volatility ($p = 0.030$) rather than returns ($p = 0.430$). In the most recent phase, therefore, the VIX acts more as a forecast of risk than as a predictor of direction, which is what an efficient

options market would imply. The variance decompositions in Table 10 indicate that, at a ten-week horizon, returns explain about 27 per cent of the forecast-error variance of ΔVIX , while VIX shocks explain only 2.8 per cent of the variance of returns.

Table 10. Forecast-error variance decomposition at a ten-week horizon

Variance of...	explained by Return	by ΔVIX	by $\Delta \ln RV$
Return	96.7%	2.8%	0.5%
ΔVIX	27.2%	71.8%	1.0%
$\Delta \ln RV$	4.8%	3.1%	92.1%

Note: Based on the Cholesky ordering in Figure 3. Rows sum to 100%

4.6 Structural breaks and volatility regimes

If inclusion had changed the uncertainty regime, the data should place breaks near August 2014 or April 2020. Table 11 shows that they do not. For realised volatility the BIC selects four breaks, located in August 2012, June 2016, February 2020 and October 2022; for log VIX the corresponding four-break partition places the breaks within a few weeks of the same dates (August 2012, July 2016, February 2020 and October 2022). BIC gains from

adding further VIX breaks are marginal (less than 1 per cent of the criterion), which is expected for a highly persistent series. The break dates align with global rather than domestic events: the easing of euro-area stress after mid-2012, the end of the 2015–2016 global sell-off, the onset of the pandemic and the passing of the peak of global monetary tightening in late 2022. None coincides with PMJDY or the start of the retail boom.

Table 11. Multiple structural breaks in the mean of log realised volatility and log India VIX

Regime	Realised volatility: dates	Mean RV (%)	Weeks	India VIX: dates	Mean VIX (%)	Weeks
1	2011-01-14 to 2012-08-03	17.5	82	14 Jan 2011 to 27 Jul 2012	23.0	81
2	2012-08-10 to 2016-05-27	12.5	199	3 Aug 2012 to 1 Jul 2016	17.0	205
3	2016-06-03 to 2020-01-31	9.9	192	8 Jul 2016 to 20 Feb 2020	14.4	190
4	2020-02-07 to 2022-10-14	15.4	141	28 Feb 2020 to 21 Oct 2022	20.8	139
5	2022-10-21 to 2026-09-18	9.8	205	28 Oct 2022 to 18 Sep 2026	13.4	204

Note: Bai–Perron global least-squares estimation with up to six breaks and 10% trimming. RV: four breaks selected by BIC. VIX: four-break partition reported for comparability; means are geometric (exp of mean log level)

The Markov-switching model (Table 12, Figure 4) identifies a calm regime with a weekly standard deviation of 1.82 per cent and a significantly positive mean return, and a turbulent regime with a standard deviation of 4.51 per cent and a negative, though imprecise, mean. Calm

spells last on average 94 weeks and turbulent spells about 9 weeks. The turbulent regime accounted for 11.6 per cent of weeks in Phase I but only 3.8 and 4.1 per cent in Phases II and III. The inclusion phases have thus experienced turbulence less often—mainly in early 2016, briefly in late 2018, in March–June 2020 and in early 2021—while the sell-offs of 2022 and March 2026 lifted the turbulent-regime probability only to about 0.49 and 0.31 respectively. When turbulence does arrive, however, it takes the same form as before.

Table 12. Two-state Markov-switching model of weekly returns

	Calm regime	Turbulent regime
Mean weekly return (%)	0.238***	-0.594
Weekly SD (%)	1.819	4.507
Staying probability	0.989	0.888
Expected duration (weeks)	93.9	8.9
Share of weeks classified ($P > 0.5$): Phase I	88.4%	11.6%
Phase II	96.2%	3.8%
Phase III	95.9%	4.1%

Note: Switching mean and variance. ***, **, * denote significance at 1%, 5% and 10%.




Figure 4. Smoothed probability of the turbulent regime



Figure 5. Rolling 52-week correlation between weekly returns and ΔVIX

Note: Shading and dashed lines as in Figure 1.

4.7 Financial inclusion, participation and market conditions (H5)

Table 13 and Figure 6 bring together the annual data. Demat accounts grew by 13 per cent (log) in FY2020, 30 per cent in FY2021 and 49 per cent in FY2022, before

settling at 24–28 per cent a year and slowing to 15 per cent in FY2026. The FI-Index rose steadily by 2.5–4.1 points a year. Market uncertainty followed a different path, peaking in FY2021 (mean VIX 25.5) and falling to a low in FY2024 (12.1).

Table 13. Annual financial-inclusion and market-uncertainty indicators, FY2019–FY2026

FY (ending March)	Demat accounts (crore)	Demat growth (log %)	FI-Index	NIFTY return (%)	FY (log)	Mean VIX	Peak VIX	Ann. vol. (%)
FY19	3.59	—	n.p.	13.5		15.2	19.8	12.8
FY20	4.09	13.0	n.p.	-29.0		19.0	70.4	20.7
FY21	5.51	29.8	53.9	51.6		25.5	55.3	25.4
FY22	8.97	48.7	56.4	16.8		17.5	28.0	14.7
FY23	11.44	24.3	60.1	1.2		17.1	23.5	14.5
FY24	15.13	28.0	64.2	25.2		12.1	15.4	9.5
FY25	19.25	24.1	67.0	5.2		14.5	24.6	13.4
FY26	22.45	15.4	70.0	-3.0		13.7	26.8	12.4

Note: n.p. = not published. FI-Index for March 2017 was 43.4 (RBI, 2021). Demat totals are NSDL + CDSL accounts at end-March. FY return is the log change in the NIFTY between the last weekly closes of successive fiscal years. Sources: RBI; PIB; SEBI; NSE Indices via Yahoo Finance; author's calculations.




Figure 6. Participation, inclusion and uncertainty at the annual frequency

The rank correlations in Table 14 suggest two regularities. First, demat growth is strongly positively related to same-year market returns ($\rho = 0.86$, $p = 0.014$), which is consistent with return-chasing entry. Second, demat growth is positively related to the previous year's mean VIX ($\rho = 0.82$, $p = 0.023$) but not to the same year's VIX. The largest waves of account opening (FY2021 and FY2022) followed the most uncertain year of the sample, suggesting that new investors were drawn in by the

recovery from the crash rather than deterred by the preceding turbulence. The FI-Index is negatively related to mean VIX ($\rho = -0.83$, $p = 0.042$), but this correlation mainly reflects opposite trends over a six-year window—the index rising and volatility declining from its pandemic peak—and should not be read as causal. With so few observations, these results are descriptive and support H5 only tentatively.

Table 14. Rank correlations between annual inclusion and market indicators

Pair	Spearman ρ	p	n
Demat growth vs. NIFTY FY return (same year)	0.857	0.014	7
Demat growth vs. previous-year mean VIX	0.821	0.023	7
Demat growth vs. same-year mean VIX	0.179	0.702	7
Demat growth vs. previous-year return	0.036	0.939	7
FI-Index level vs. mean VIX	-0.829	0.042	6
Change in FI-Index vs. NIFTY FY return	0.100	0.873	5

Note: Spearman rank correlations; with $n \leq 7$, p-values are indicative only.

4.8 Robustness

Several checks support the main conclusions. (i) The results do not depend on the volatility model: GARCH, EGARCH and GJR agree on persistence and on the presence of asymmetry, and the HAR evidence on realised volatility is consistent with the GARCH evidence on conditional volatility. (ii) The absence of phase effects in the HAR models holds without the COVID-19 dummy (joint $\chi^2 = 0.05$, $p = 0.977$) and with or without VIX–phase interactions. When lagged VIX is dropped altogether, the Phase III dummy becomes negative and marginally significant (-0.067 , $p = 0.086$), which confirms that the lower volatility of the recent phase is transmitted through lower implied volatility. (iii) The break dates for realised and implied volatility, estimated independently, lie within four weeks of each other. (iv) The out-of-sample gain from the VIX is obtained over a period (2022–2026) that was entirely outside the estimation window and contains the 2026 geopolitical shock. (v) Weekly sampling removes day-of-week effects and asynchronous closing issues between the NIFTY and the VIX; the daily return moments reported in Section 3.1 show the same asymmetry (skewness -0.92). A limitation is that the phases are defined by policy dates; results using data-driven breaks (Table 11) lead to the same interpretation.

5. DISCUSSION

5.1 Summary of hypotheses

H1 and H2 are supported: Indian market uncertainty is asymmetric, clusters for several weeks and is anticipated by the India VIX. H3 is supported only in a weak sense—volatility levels and the frequency of turbulent regimes were lower in the inclusion phases—but rejected in its strong sense, because the parameters of the volatility process and the residual volatility after controlling for implied volatility are stable across phases. H4 is not supported: the fear response to market declines in the participation-deepening phase is statistically indistinguishable from that of the pre-mission phase. H5 receives tentative support: participation growth is procyclical with returns and rises after high-uncertainty years.

5.2 Interpretation

The results point to a market in which inclusion has widened ownership without changing the way uncertainty is generated and priced. A likely explanation is that the marginal price-setters in Indian equities remain institutional and foreign investors, whose behaviour is shaped by global risk appetite; this is why the break dates coincide with global events. Growth in domestic participation, including systematic investment through mutual funds, may have supplied a stable source of demand that made turbulent regimes less frequent, but not strong enough to alter the transmission of global shocks. These findings are consistent with the view that inclusion can support stability when it broadens the base of long-term savers (Han & Melecky, 2013; Mehrotra & Yetman, 2015), and they do not support the concern that new retail participation has made Indian markets more fragile.

At the same time, the asymmetry of uncertainty matters for the welfare of new investors. Households that entered in the post-2020 boom did so during a period of unusually low implied volatility—the VIX fell below 10 in late

2025—and the annual evidence suggests that entry follows good returns. Such investors are therefore exposed to exactly the kind of downside surprise that the GJR and Markov-switching models describe: sudden, fear-driven episodes such as March 2020 or March 2026, when volatility rises sharply and persists for about six weeks. The literature on retail investors suggests that inexperienced investors are the most likely to sell in such episodes and crystallise losses (Barber & Odean, 2000). The fact that account growth slowed to 15 per cent in FY2026, a year with a negative index return, is consistent with this sensitivity.

5.3 Implications

For regulators, the results suggest that the priority is not to limit participation on stability grounds but to strengthen its quality—the component of the FI-Index that carries the smallest weight. Investor education that explains asymmetric risk, the use of implied volatility as a risk signal and the benefits of staying invested through turbulent regimes could reduce the chance that new investors buy after rallies and sell after crashes. Disclosure of volatility conditions at the point of sale on digital platforms, and suitability rules for leveraged derivatives, are natural complements. For the RBI and SEBI, the strong predictive content of the India VIX supports its use in macroprudential monitoring dashboards alongside inclusion metrics.

For investors and advisers, the India VIX provides a timely and statistically reliable forecast of realised risk, which can be used to adjust exposure or rebalance. For asset managers designing products for first-time investors, the evidence of procyclical entry argues for automatic, rules-based investment plans that smooth the timing of purchases. For researchers, the results suggest that the effect of inclusion on markets should be studied through the composition and behaviour of investor cohorts rather than through aggregate volatility alone.

6. CONCLUSION, LIMITATIONS AND FUTURE RESEARCH

This study examined whether the rapid expansion of financial inclusion into investment in India has changed the nature of investment uncertainty in its stock market. Using fifteen years of weekly data and a set of complementary models, it found that uncertainty is asymmetric, moderately persistent and well captured by the India VIX, and that its dynamics have remained stable through the period of inclusion, with shifts driven by global shocks. Inclusion phases have been accompanied by lower average volatility and fewer turbulent weeks, but not by a different volatility process. Participation growth is procyclical and tends to follow recoveries from high-uncertainty episodes.

The study has limitations. The FI-Index is available annually and only from 2021 (with a 2017 value), so the link between inclusion and market outcomes at the annual level rests on few observations and is descriptive. Phase dummies capture the timing of inclusion, not its intensity, and they may also pick up other contemporaneous changes such as regulatory reforms or shifts in global liquidity. The analysis uses index-level data and cannot

distinguish the trading of new investors from that of other participants. Implied volatility includes a variance risk premium, so part of its predictive power may reflect risk compensation rather than pure expectations (Bekaert & Hoerova, 2014).

Future research could use account-level or investor-category data from the exchanges and depositories to measure the trading of new cohorts directly, construct monthly inclusion measures from depository, mutual-fund and payments data, and extend the analysis to derivatives, where retail participation has grown quickly. Cross-country comparisons with other emerging markets that experienced digital participation booms would help establish how general the Indian findings are. Finally, primary surveys of first-time investors could connect the market-level asymmetry documented here with the expectations and financial literacy of individual households.

REFERENCES

- Agarwalla, S. K., Barua, S. K., Jacob, J., & Varma, J. R. (2015). Financial literacy among working young in urban India. *World Development*, 67, 101–109. <https://doi.org/10.1016/j.worlddev.2014.10.004>
- Andersen, T. G., & Bollerslev, T. (1998). Answering the skeptics: Yes, standard volatility models do provide accurate forecasts. *International Economic Review*, 39(4), 885–905.
- Bai, J., & Perron, P. (1998). Estimating and testing linear models with multiple structural changes. *Econometrica*, 66(1), 47–78.
- Bai, J., & Perron, P. (2003). Computation and analysis of multiple structural change models. *Journal of Applied Econometrics*, 18(1), 1–22.
- Baker, M., & Wurgler, J. (2006). Investor sentiment and the cross-section of stock returns. *The Journal of Finance*, 61(4), 1645–1680.
- Baker, S. R., Bloom, N., & Davis, S. J. (2016). Measuring economic policy uncertainty. *The Quarterly Journal of Economics*, 131(4), 1593–1636.
- Barber, B. M., Huang, X., Odean, T., & Schwarz, C. (2022). Attention-induced trading and returns: Evidence from Robinhood users. *The Journal of Finance*, 77(6), 3141–3190.
- Barber, B. M., & Odean, T. (2000). Trading is hazardous to your wealth: The common stock investment performance of individual investors. *The Journal of Finance*, 55(2), 773–806.
- Barber, B. M., & Odean, T. (2008). All that glitters: The effect of attention and news on the buying behavior of individual and institutional investors. *The Review of Financial Studies*, 21(2), 785–818.
- Beck, T., Demirgüç-Kunt, A., & Levine, R. (2007). Finance, inequality and the poor. *Journal of Economic Growth*, 12(1), 27–49.
- Bekaert, G., & Hoerova, M. (2014). The VIX, the variance premium and stock market volatility. *Journal of Econometrics*, 183(2), 181–192.
- Black, F. (1976). Studies of stock price volatility changes. In *Proceedings of the 1976 Meetings of the American Statistical Association, Business and Economic Statistics Section* (pp. 177–181).
- Bloom, N. (2009). The impact of uncertainty shocks. *Econometrica*, 77(3), 623–685.
- Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 31(3), 307–327.
- Business Standard. (2026, July 16). RBI's FI-Index rises 4.48% in FY26, growth across all sub-indices. https://www.business-standard.com/economy/news/rbi-s-fi-index-rises-4-48-in-fy26-growth-across-all-sub-indices-126071601079_1.html
- Campbell, J. Y. (2006). Household finance. *The Journal of Finance*, 61(4), 1553–1604.
- Chandra, A., & Thenmozhi, M. (2015). On asymmetric relationship of India volatility index (India VIX) with stock market return and risk management. *Decision*, 42(1), 33–55.
- Chiah, M., & Zhong, A. (2020). Trading from home: The impact of COVID-19 on trading volume around the world. *Finance Research Letters*, 37, 101784.
- Christensen, B. J., & Prabhala, N. R. (1998). The relation between implied and realized volatility. *Journal of Financial Economics*, 50(2), 125–150.
- Corsi, F. (2009). A simple approximate long-memory model of realized volatility. *Journal of Financial Econometrics*, 7(2), 174–196.
- Deccan Chronicle. (2024, April). Demat accounts up 33% in FY24. <https://www.deccanchronicle.com/business/demat-accounts-up-33-in-fy24-889996>
- Demirgüç-Kunt, A., Klapper, L., Singer, D., & Ansar, S. (2022). The Global Findex Database 2021: Financial inclusion, digital payments, and resilience in the age of COVID-19. World Bank.
- 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.
- Diebold, F. X., & Mariano, R. S. (1995). Comparing predictive accuracy. *Journal of Business & Economic Statistics*, 13(3), 253–263.
- EquityPandit. (2022, April). Total demat accounts in FY22 up 63% to 89.7 million. <https://www.equitypandit.com/total-demat-accounts-in-fy22-up-63-to-89-7-million/>
- Engle, R. F. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica*, 50(4), 987–1007.
- Engle, R. F., & Ng, V. K. (1993). Measuring and testing the impact of news on volatility. *The Journal of Finance*, 48(5), 1749–1778.
- Fleming, J., Ostdiek, B., & Whaley, R. E. (1995). Predicting stock market volatility: A new measure. *Journal of Futures Markets*, 15(3), 265–302.
- Giot, P. (2005). Relationships between implied volatility indexes and stock index returns. *The Journal of Portfolio Management*, 31(3), 92–100.
- 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.
- Granger, C. W. J. (1969). Investigating causal relations by econometric models and cross-spectral

- methods. *Econometrica*, 37(3), 424–438.
32. Guiso, L., Sapienza, P., & Zingales, L. (2008). Trusting the stock market. *The Journal of Finance*, 63(6), 2557–2600.
33. Hamilton, J. D. (1989). A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica*, 57(2), 357–384.
34. Han, R., & Melecky, M. (2013). Financial inclusion for financial stability: Access to bank deposits and the growth of deposits in the global financial crisis (Policy Research Working Paper No. 6577). World Bank.
35. 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.
36. Koenker, R., & Bassett, G. (1978). Regression quantiles. *Econometrica*, 46(1), 33–50.
37. 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.
38. Lusardi, A., & Mitchell, O. S. (2014). The economic importance of financial literacy: Theory and evidence. *Journal of Economic Literature*, 52(1), 5–44.
39. Mehrotra, A., & Yetman, J. (2015). Financial inclusion – issues for central banks. *BIS Quarterly Review*, March, 83–96.
40. Nelson, D. B. (1991). Conditional heteroskedasticity in asset returns: A new approach. *Econometrica*, 59(2), 347–370.
41. Newey, W. K., & West, K. D. (1987). A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, 55(3), 703–708.
42. Ozik, G., Sadka, R., & Shen, S. (2021). Flattening the illiquidity curve: Retail trading during the COVID-19 lockdown. *Journal of Financial and Quantitative Analysis*, 56(7), 2356–2388.
43. Ozili, P. K. (2018). Impact of digital finance on financial inclusion and stability. *Borsa Istanbul Review*, 18(4), 329–340.
44. Phillips, P. C. B., & Perron, P. (1988). Testing for a unit root in time series regression. *Biometrika*, 75(2), 335–346.
45. Press Information Bureau. (2021, December). India’s demat account holders more than double in 3 years to 7.38 crore [Press release]. Ministry of Finance, Government of India. <https://www.pib.gov.in/PressReleaseIframePage.aspx?PRID=1780999>
46. Press Information Bureau. (2025). RBI’s Financial Inclusion Index rises to 67 in 2025 [Press note]. Government of India. <https://www.pib.gov.in/PressNoteDetails.aspx?ModuleId=3&NoteId=154980>
47. Reserve Bank of India. (2021, August 17). RBI releases Financial Inclusion Index [Press release]. https://www.rbi.org.in/Scripts/BS_PressReleaseDisplay.aspx?prid=52068
48. Reserve Bank of India. (2022, August 2). Financial Inclusion Index for March 2022 [Press release]. https://www.rbi.org.in/Scripts/BS_PressReleaseDisplay.aspx?prid=54133
49. Sarma, M. (2008). Index of financial inclusion (Working Paper No. 215). Indian Council for Research on International Economic Relations.
50. Sarma, M., & Pais, J. (2011). Financial inclusion and development. *Journal of International Development*, 23(5), 613–628.
51. Securities and Exchange Board of India. (2026). Annual report 2025–26 (Chapter 4: Secondary markets). <https://www.sebi.gov.in/reports-and-statistics/publications/aug-2026/Chapter%2004.pdf>
52. Shaikh, I., & Padhi, P. (2015). The implied volatility index: Is ‘investor fear gauge’ or ‘forward-looking’? *Borsa Istanbul Review*, 15(1), 44–52.
53. Sims, C. A. (1980). Macroeconomics and reality. *Econometrica*, 48(1), 1–48.
54. van Rooij, M., Lusardi, A., & Alessie, R. (2011). Financial literacy and stock market participation. *Journal of Financial Economics*, 101(2), 449–472.
55. Welch, I. (2022). The wisdom of the Robinhood crowd. *The Journal of Finance*, 77(3), 1489–1527.
56. Whaley, R. E. (2000). The investor fear gauge. *The Journal of Portfolio Management*, 26(3), 12–17