International Journal for Multidisciplinary Research (IJFMR)

E-ISSN: 2582-2160 ● Website: www.ijfmr.com ● Email: editor@ijfmr.com

Impact of AI-Driven Trading on Market Liquidity and Volatility: An Empirical Analysis of Modern Financial Markets

Mr. Sohaib Mehmood[1], Dr. Aasim Mir[2]

[1]Research Scholar, Department Of Management Studies, Baba Ghulam Shah Badshah University
[2]Assistant Professor, Department Of Management Studies, Baba Ghulam Shah Badshah University

ABSTRACT

Financial markets across the world have undergone significant change over the past two decades, and majority of that has been pushed by computing power, fast data feeds, and the steady automation of order execution. Among the more recent developments, AI-driven trading systems have begun to alter how decisions are made on the buy and sell side. These systems read large volumes of market and non-market data, revise their own strategies on the fly, and place orders without interference of human. Using daily data on the NIFTY 50 index, the India VIX, and the policy repo rate from January 2015 to December 2024, the analysis used 2,475 trading-day observation. The testing strategy is indirect because exchanges in India do not publish any data of AI-driven order flow. The argument for the study is that if the market is being shaped by AI-style trading, then its observable behaviour should look like what the literature already documents for algorithmically intensive markets. The descriptive results are consistent with that view: daily log returns are sharply negative skewed (-1.4144) and heavy-tailed (kurtosis = 23.53), a far cry from the Gaussian benchmark. India VIX turns out to be a strong determinant of both the Amihud illiquidity measure ($t = 6.83$, $p < 0.001$) and short-run absolute returns ($t = 8.90$, $p < 0.001$). Estimated GARCH(1,1) parameters give $\alpha + \beta = 0.970$ — close to a unit-root volatility process — and the GJR-GARCH leverage term $\gamma = 0.150$ ($p = 0.0003$) is highly significant, so negative shocks raise conditional variance much more than positive shocks of the same size. Splitting the sample into pre-COVID, COVID, and post-COVID windows shows that realized volatility during COVID was 2.05 times its pre-crisis level, while post-crisis liquidity (lower Amihud values) was about 46 percent better than the pre-COVID baseline. All three null hypotheses are rejected at the 1 percent level. Taken together, the NIFTY 50 displays the kind of liquidity, volatility, and stability behaviour the literature associates with markets shaped by AI-driven activity. The paper's contribution is twofold: it shows that an indirect, outcome-based testing framework can be applied where direct trading-system data are not available, and it draws out specific implications for SEBI, the RBI, and large institutional participants in India.

Keywords: AI-Driven Trading; Market Liquidity; Volatility Dynamics; GARCH Models; Indian Financial Markets; NIFTY 50; Market Microstructure

INTRODUCTION

Evolution of Financial Markets and Technology

Financial markets sit at the heart of any modern economy. They move savings into investment, set the prices that decide where capital goes, and give firms and households a way to share risk across time. When markets work, the economy gets cheaper funding and absorbs shocks more smoothly. When they do not, the cost is felt almost immediately in real activity — in jobs, in lending, in business confidence. For most of the twentieth century, all of this work was done on physical trading floors. Brokers shouted bids and offers across crowded rooms, clerks wrote the trades down by hand, and the prices that came out reflected a mix of fundamentals, judgement, and whatever a human mind could process in the moment. Open outcry held up for a long time, but the limits were obvious. Trades were slow. Costs were high. Transparency was patchy. And information rarely reached the small investor on the same terms it reached the dealers on the floor. The shift to electronic markets started in the late 1980s and picked up speed through the 1990s. Limit-order books moved onto screens, matching engines began automating the cross between buy and sell orders, and the time it took to execute a trade fell from seconds to milliseconds. The change was not just technical. It quietly changed what kind of trading made commercial sense in the first place. Algorithmic trading grew straight out of this new setup. Once the order book was a data feed and the matching engine was effectively a queue, it became natural — almost inevitable — to write programs that could place, cancel, and update orders far faster than any human could. Institutions used these programs first to break up large parent orders into smaller ones so the market impact stayed manageable. Over time, the same infrastructure was used for genuine signal trading and market making. The benefits were real. Spreads narrowed, the top of the order book got deeper, and the implicit cost of trading for end investors came down. But the same infrastructure raised serious concerns. Co-located servers and direct exchange feeds created a tier of participants whose advantage was structural — built into the wires — rather than analytical. The May 2010 Flash Crash, when the Dow lost nearly a thousand points in minutes before rebounding, made it clear that the same speed that made markets efficient could also make them fragile. That memory is still part of the regulatory backdrop today.

Artificial Intelligence in Financial Markets

Artificial intelligence builds on this foundation but moves the goalposts. A rule-based algorithm executes a fixed strategy: if condition X holds, send order Y. An AI-driven system is supposed to do something different — to estimate, from data, what the right response to a market state is, and to keep updating that estimate as the market changes. In practice this means using machine learning, deep learning, or reinforcement learning to map noisy signals — order-book depth, news flow, sentiment scores, macro releases — into trading decisions. The important property is adaptiveness: the model learns from outcomes and adjusts. This lets firms exploit non-linearities and weak signals that a hand-coded rule would miss, and it is part of why a meaningful share of equity volume in major markets is now generated by such systems. The same adaptiveness creates problems regulators have not fully resolved. Many models are opaque enough that their behaviour cannot be reverse-engineered after the fact, which makes supervision harder. And because firms tend to use overlapping data sets and similar architectures, the responses across firms can be more correlated than the data feed would suggest. When a few standard signals tell many AI systems to deleverage at the same time, the move is amplified rather than absorbed.

AI-Driven Trading versus Traditional Algorithmic Trading

It is worth being precise about how AI-driven trading differs from the broader category of algorithmic trading, since the two are often spoken of as one. Traditional algorithmic systems automate the execution of a strategy that a human has already specified. The decision logic is fixed; only the speed and consistency of implementation are mechanical. AI-driven systems flip this. The strategy itself is what the model is trying to find. The system observes the market, evaluates the outcomes of its past trades, and revises its behaviour. That distinction matters operationally — adaptive systems can respond to market regimes that a static rule was never written for — but it also matters for market structure. When many adaptive systems share inputs, the joint distribution of their actions can produce feedback loops that no individual model is steering. The literature has begun to take this seriously, and so does this study.

Table 1.0: Comparison Between Traditional Algorithmic Trading and AI-Driven Trading

Feature	Traditional Algorithmic Trading	AI-Driven Trading
Decision Rules	Fixed, pre-programmed	Adaptive, data-driven
Learning Capability	None — static logic	Yes — learns from historical data
Decision Autonomy	Executes human instructions	Generates strategies independently
Execution Speed	High (milliseconds)	Ultra-high (microseconds)
Data Types Used	Structured (price, volume)	Structured + Unstructured (news, sentiment)
Transparency	High — rule-based, auditable	Low — black-box model
Market Impact	Predictable, stable	Emergent, potentially unpredictable
Regulatory Risk	Lower	Higher — opacity and systemic concerns

Indian Financial Market Perspective

India is a useful place to ask these questions. The Indian equity market has gone through a compressed version of the global story: a move from a fragmented, partly manual market in the 1990s to a screen-based, demutualised, NSE-led environment by the mid-2000s, and from there to a market where co-location, algorithmic order entry, and increasingly AI-augmented execution are routine for institutional participants. The NSE introduced co-location services in 2010, SEBI issued its first set of algorithmic-trading guidelines around the same time, and the share of orders generated by automated systems on Indian exchanges has grown steadily since then. Today the NIFTY 50 sits inside a market structure with millisecond-level matching, deep retail participation through digital brokers, and an active derivatives segment that is among the most liquid in the world by contract count. At the same time, the Indian market still differs from the U.S. or European markets in ways that are likely to matter for AI-trading effects: institutional ownership is more concentrated, market making is less specialised, and several macro variables that influence flows — the policy repo rate, the rupee, oil — are tightly bound to a single domestic regulator. That makes the country a natural test case for whether the patterns documented elsewhere generalise.

Problem Statement

The motivation for this study comes from a gap rather than a controversy. Financial technology has changed the structure of trading, AI methods have been added to that mix, and yet most of the empirical work on what these changes do to liquidity, volatility, and stability has been carried out on U.S. equities or the major European indices. The Indian evidence is thinner and tends to focus on algorithmic trading in general rather than AI-driven systems specifically. As a result, claims about the Indian market often rest on extrapolation from contexts that look quite different from the NSE in 2024.

The other motivation is mixed evidence. The same literature that documents tighter spreads and faster price discovery in normal times also documents short-lived liquidity withdrawals, correlated unwinds, and amplified moves around stress events. AI-driven trading does not obviously land on one side of that ledger. It can deepen the book under calm conditions and yet behave more pro-cyclically when volatility spikes. Whether the Indian market shows the benign side, the pro-cyclical side, or both at different times is an empirical question that this paper tries to answer.

Two practical constraints shape the research design. First, the exchanges do not flag orders as AI-generated, so direct measurement is not available. Second, the data that are available — index prices, the volatility index, turnover, the policy rate — are exactly the variables the existing literature uses to characterise algorithmically intensive markets. The study therefore asks whether observable Indian market behaviour over 2015–2024 lines up with the signatures the literature ascribes to such markets. Doing this for India also helps fill the emerging-market gap in this area.

LITERATURE REVIEW

Introduction

The growing role of AI and algorithmic systems in financial markets has produced an active and heterogeneous literature. Different studies use different sample markets, different measures of automation intensity, and different econometric approaches, and they do not all reach the same conclusions. The aim of this section is not to summarise that literature in full but to draw out the strands relevant to the present analysis: how automated and AI-driven trading affects liquidity, how it interacts with volatility, and how the available evidence applies (or fails to apply) to the Indian market.

The review is organised conceptually rather than chronologically. It begins with the distinction between rule-based algorithmic trading and AI-driven trading, then moves to the liquidity literature, the volatility literature, and the evidence on extreme events and regime dependence. The final subsection turns to the Indian context, where the available empirical work is sparse but instructive.

Conceptual Foundations of Algorithmic and AI-Driven Trading

What people mean by automated trading has changed substantially over the past two decades. Early in the 2000s, the term referred mostly to the execution of large parent orders by slicing them into smaller child orders that were placed by software according to a preset schedule. By the early 2010s the meaning had broadened to include high-frequency market making and short-horizon statistical strategies. By the late 2010s, references to AI-driven or machine-learning-based trading were common in both academic and industry writing, even where the underlying systems were quite varied.

Aldridge (2013) gives one of the cleanest early definitions. She characterises algorithmic trading as automated execution of orders by computers running pre-programmed instructions, with rules covering price, volume, and timing. The framing is useful because it isolates the part that is genuinely automated — execution — from the part that is not, namely the strategy itself, which is still designed by humans.

Cartea, Jaimungal, and Penalva (2015) push past this static picture by separating high-frequency strategies into market making, statistical arbitrage, and event-driven trading. Their treatment matters because it shows that even within rule-based algorithms, the economic role of the trader varies — market makers supply liquidity, statistical-arbitrage strategies absorb temporary mispricings, event-driven strategies trade on information shocks. Lumping them together obscures effects that go in opposite directions.

AI-based trading marks a further break. Boehmer, Fong, and Wu (2021) note that the move from fixed rules to learning-based systems introduces what they call synchronization risk: many firms train on overlapping data with similar architectures, and the resulting models can react in correlated ways to a shared signal. The market-wide effect is one that no individual model intended. Sangiorgi and Schiavone (2025) go further along the same line, arguing that machine-learning methods give trading systems the ability to detect non-linear patterns and shift strategies across regimes — a useful capability in normal times that, under stress, can produce coordinated responses.

Read together, these contributions trace a fairly clean conceptual progression. Rule-based algorithms are deterministic implementations of human-designed strategies. AI-driven systems are adaptive, data-driven, and harder to interpret either after the fact or in real time. The qualitative gap matters for everything that follows in this paper, because it changes what one expects to observe in the market.

Algorithmic Trading and Market Liquidity

Liquidity is the obvious place to start. It is one of the older empirical concerns in microstructure, and most of the influential work on automated trading has either led with a liquidity measure or used one as a robustness check.

Hendershott, Jones, and Menkveld (2011) is the standard reference. They use the introduction of NYSE's automated quote dissemination in 2003 as a quasi-natural experiment and find that algorithmic trading narrows quoted and effective spreads, reduces adverse-selection costs, and improves price efficiency for large stocks. The mechanism is intuitive — competition among automated market makers drives spreads toward marginal cost — but the size of the effect surprised many readers at the time.

Hasbrouck and Saar (2013) reach a compatible conclusion using a different design. They track high-frequency activity directly through millisecond-level message data on NASDAQ and find that periods of intense low-latency trading are associated with deeper books, more stable spreads, and lower short-term volatility. Importantly, their identification comes from variation within stocks and within the day, which gets around concerns that liquidity and HFT are jointly determined by some omitted factor.

The picture becomes more nuanced when the universe is widened. Boehmer, Fong, and Wu (2021) use international data covering 42 markets and confirm the headline finding — algorithmic trading is associated with tighter spreads in liquid stocks — but they also document that the benefits decline sharply as one moves down the size and turnover spectrum. For small and mid-cap stocks, the relationship is weaker and at times reverses. Karkowska and Palczewski (2023), looking at European venues over a longer window, document a pattern in which spreads compress during normal regimes but widen rapidly during stress, a pattern the authors attribute to coordinated withdrawal by automated market makers. Together these papers suggest that the liquidity benefits of automation are real but conditional on market state.

AI-Driven Trading and Liquidity Provision

AI-driven trading raises a slightly different question. If a learning-based system continually updates its quotes in response to estimated risk, when exactly does it withdraw, and how does that decision interact with similar decisions at other firms? The literature on this is younger than the algorithmic-liquidity literature but is growing quickly.

Menkveld (2013) is an early and still useful reference. He characterises high-frequency trading firms as the new market makers — they post two-sided quotes through limit orders and earn the spread, much as designated market makers used to do — but he also highlights that the commercial logic of automated market making collapses when price uncertainty rises sharply. At that point the inventory risk of holding the spread exceeds the expected profit, and the system cancels its quotes. Sangiorgi and Schiavone (2025) extend the analysis to AI systems and argue that adaptive traders can sustain liquidity across a wider range of conditions than rule-based systems can, but that synchronization across similarly trained models can amplify liquidity withdrawal when a common volatility signal triggers correlated risk-management responses.

The implication for this paper is straightforward. The literature does not predict that AI-driven liquidity provision is uniformly good or uniformly bad. It predicts a state-dependent pattern: improved liquidity in calm markets, sharp deterioration around stress. The Indian regime-wise tests in Section (Regime-Wise Sub-Period Analysis) are designed to look for exactly this pattern.

Algorithmic and AI-Driven Trading and Market Volatility

Volatility has produced more disagreement than liquidity. The disagreement is partly definitional — whether to look at realized, implied, or conditional volatility — and partly because the relevant horizons differ. Effects that look stabilising at the daily frequency can look destabilising at the intraday frequency, and vice versa.

Most empirical work concludes that automated trading raises short-term volatility while improving longer-term price efficiency. Brogaard, Hendershott, and Riordan (2014) document that high-frequency traders contribute meaningfully to price discovery on NASDAQ but also trade in the direction of permanent price changes, increasing realized volatility on information days. Chaboud, Chiquoine, Hjalmarsson, and Vega (2014) find similar evidence in FX markets, where algorithmic activity tightens spreads but also raises volatility around news. Boehmer, Fong, and Wu (2021) confirm the international applicability of the pattern: more algorithmic trading, higher short-horizon volatility, better long-horizon efficiency. Alliata and Bozagiu (2025) take this further by estimating volatility models on AI-intensive environments specifically and find that machine-learning-based participation is associated with stronger volatility clustering and a higher incidence of extreme price events.

The literature thus presents a coherent if double-edged picture. Automated and AI-driven trading make markets more informationally efficient under most conditions while also producing higher short-run volatility, more clustering, and a thicker tail.

AI-Driven Trading, Volatility Clustering, and Extreme Events

A separate line of work asks whether AI-driven trading changes the structure of extreme events — flash crashes, regime breaks, and periods of synchronized liquidity withdrawal — rather than just the typical volatility level. Kirilenko and Lo (2013) is an early contribution; their analysis of the May 2010 Flash Crash documents how a confluence of automated strategies, each individually rational, produced a market-wide dislocation that reversed within minutes. More recent work has argued that AI systems may make such episodes more rather than less likely. Boehmer, Fong, and Wu (2021) point out that the volume of order submissions and cancellations rises sharply with automation, increasing the system's sensitivity to small disturbances. Alliata and Bozagiu (2025) use both OLS and GARCH frameworks to show that AI participation correlates with volatility clustering and a higher tail-event probability.

Sangiorgi and Schiavone (2025) approach the same question through agent-based simulation. In their setup, populating an artificial market with reinforcement-learning agents initially improves efficiency, but past a threshold the market becomes more erratic, non-linear, and prone to bouts of instability. The point of their result is not the specific threshold but the mechanism: emergent instability is produced by the interaction of agents, not by the design of any individual system. The implications for systemic risk in real markets are direct.

Indian Market Context

The Indian literature is smaller but has grown over the past few years. The general finding echoes the international evidence: algorithmic trading on the NSE is associated with tighter spreads under normal conditions and somewhat higher short-run volatility, with effects strongest in the most heavily traded stocks. Bhatia and Batra (2020) document this pattern using a panel of NSE stocks; Dubey, Babu, Jha, and Varma (2017) reach a similar conclusion using transaction-velocity measures.

Within this group, Nath (2020) provides one of the more careful econometric treatments. He uses time-series methods to study the relationship between algorithmic trading intensity and volatility in Indian equities and finds that short-term volatility increases with algorithmic participation while long-term price efficiency improves — the now-familiar pattern. He also finds that the volatility effect is stronger during stress periods, which motivates the regime-wise design used in this paper.

Useful as this work is, three gaps remain. First, almost all of it treats algorithmic trading as a single category and does not separate out AI-driven systems specifically. Second, most studies pre-date COVID, so the regime-dependent behaviour that international evidence has come to emphasise is not well documented for India. Third, the indirect outcome-based testing strategy used widely in international work — inferring AI-trading effects from market-wide behaviour patterns — has not been applied systematically on Indian data. Filling these three gaps is the empirical task of the rest of the paper.

RESEARCH OBJECTIVES and HYPOTHESES

Objectives of the Study

The objectives of this study are:

1. To analyze the growth and evolution of AI-driven and algorithmic trading, both globally and in India.
2. To analyze the relationship between AI-driven/algorithmic trading and key market liquidity indicators.
3. To analyze the impact of AI-based trading system on both short-term and long-term market volatility.
4. To ascertain whether AI trading enhance market efficiency or leads to systemic instability.
5. To provide policy suggestions for regulators and exchanges to ensure responsible integration of AI in financial market.

Hypotheses of the Study

The study is based on the following null hypotheses:

H_{01}: AI-driven trading does not have the significant effect on market liquidity.

H_{02}: AI-driven trading does not have a significant effect on how volatile the market is.

H_{03}: There is no direct relation between the intensity of AI-driven trading and the stability of the market.

Significance of the Study

This study is meant to be useful to a few different groups of people.

- Regulatory and policy perspective: The study provides valuable empirical evidence regarding the impact of AI-driven trading on market liquidity and volatility.
- Investors and market participants: The study offers practical insights into the risks and benefits associated with AI-driven trading environments.
- Academic perspective: The research contributes to the existing body of literature by addressing several important gaps in the context of emerging markets.

Research Gap

A lot has already been written about algorithmic and high-frequency trading in global markets, but the empirical evidence on how AI-driven trading specifically affects liquidity and volatility is still thin — and it is thinnest for emerging markets like India. Three gaps stand out.

The first is measurement. Most of the Indian studies still treat algorithmic trading as one category, even though the literature on developed markets has started separating AI-driven systems from rule-based ones. The distinction is not cosmetic. Rule-based algorithms execute the same instructions every day; AI systems learn, adapt, and respond to real-time information, and they behave differently under stress. Because of that, the empirical signatures they leave on the market — heavier tails, more persistent volatility, sharper asymmetric responses — are not the same.

The second gap is geographical. Even for algorithmic trading more broadly, the bulk of careful empirical work uses U.S. and European data. India is the third-largest equity market in the world by number of trades, but the volume of detailed empirical research on how automated trading is shaping it is small relative to that size. Most of the existing studies are also pre-COVID, which matters because the international literature has converged on regime-dependence as a central feature of automated-trading effects, and that is exactly what a pre-COVID sample cannot show.

The third gap is methodological. Few of the Indian studies apply the GARCH-class volatility tools that the international literature has settled on, and even fewer split the sample by regime. As a result, the conditional behaviour of liquidity and volatility — which is where AI-driven effects show up most clearly — has not been documented carefully on Indian data.

Taken together, the absence of India-specific evidence, the missing distinction between algorithmic and AI-driven trading, and the limited methodological depth form the research gap that this study sets out to address. The empirical work uses the full January 2015 to December 2024 window, splits it into pre-COVID, COVID, and post-COVID regimes, and applies the GARCH and GJR-GARCH tools the international literature relies on, interpreting the results through the AI-driven trading frame rather than the older algorithmic-trading one.

METHODOLOGY

Research Design and Rationale for Indirect Testing

The empirical strategy in this paper is quantitative and time-series-based. The unit of observation is the trading day. The dependent variables are constructed from the NIFTY 50 index — log returns, absolute returns, 30-day realized volatility, and the Amihud illiquidity ratio — and the main right-hand-side variable is the India VIX, with the policy repo rate included as a macro control. The window is January 2015 through December 2024, which gives 2,475 daily observations and spans three identifiable regimes (pre-COVID, COVID, post-COVID).

There is one methodological problem that needs to be addressed up front. AI-driven trading activity is not directly observable in the public NSE data. The exchanges report turnover, order-to-trade ratios, and aggregate algorithmic shares, but they do not flag specific orders as AI-generated, and no commercial vendor offers a clean indicator either. A direct test of the form “regress liquidity on AI-trading intensity” is therefore not feasible with publicly available data.

The way around this is the same approach used in much of the international literature: indirect, outcome-based testing. The argument is that if the Indian equity market is in fact being shaped by AI-driven activity, then the market’s observable behaviour should display the empirical signatures the literature has documented for AI-influenced markets — non-normal returns with heavy tails, persistent and asymmetric conditional volatility, strong sensitivity of short-run volatility and liquidity to the volatility index, and sharp regime-dependent shifts during stress. Failing to find these patterns would weaken the case for the AI-influence narrative on Indian data; finding them, as Section 5 does, is consistent with it. The same strategy is used in financial economics whenever the underlying force of interest is unobservable — for example, when researchers infer information asymmetry from order-flow imbalance, or risk preferences from option pricing.

Data Sources and Sample Period

The dataset is daily and runs from 1 January 2015 to 31 December 2024. The window is long enough to estimate GARCH-class models reliably and to cover three economically distinct regimes: a pre-crisis period from January 2015 through January 2020, a COVID-19 disruption period from February 2020 through December 2020, and a post-crisis period from January 2021 through December 2024. The total sample contains 2,475 trading-day observations.

Index data — daily closing prices and turnover for the NIFTY 50 — are taken from the National Stock Exchange of India (nseindia.com). The India VIX series, which is the forward-looking implied-volatility index computed by NSE from NIFTY 50 option prices, comes from the same source. The policy repo rate is taken from the Reserve Bank of India's published statistics. Where days are missing for one series but not another, the row is dropped to keep the panel balanced. All series are checked for stationarity in Section (Stationarity Test) before any regression or volatility model is estimated.

Construction of Variables

Daily returns are computed in log form, $R_t = \ln(P_t / P_{t-1})$, where P_t is the NIFTY 50 closing level on day t. Log returns are used both because they aggregate cleanly over time and because the financial literature on volatility modelling is built around them.

For the volatility analysis, three measures are used in parallel. Absolute returns, $|R_t|$, capture short-horizon variation directly and are robust to skewness. Thirty-day realized volatility, computed as the rolling annualised standard deviation of the previous 30 daily returns, captures medium-horizon variation. Both are common in the empirical literature on automated-trading effects, and using them in parallel reduces the chance that a result is driven by the choice of volatility proxy. The conditional-volatility estimates from GARCH(1,1) and GJR-GARCH provide the third measure.

Liquidity is measured using the Amihud (2002) illiquidity ratio, defined for day t as $|R_t|$ divided by daily turnover (in rupees). The intuition is that the more the price moves per unit of volume traded, the less liquid the market is on that day. The measure is widely used in the algorithmic-trading literature precisely because it can be computed from publicly available data and is robust to the absence of intraday information. Higher values indicate worse liquidity.

Econometric Techniques

Three econometric tools are used. The first is OLS regression with Newey-West (HAC) standard errors, applied to the Amihud illiquidity ratio, absolute returns, and 30-day realized volatility. The HAC correction is needed because daily financial data are well known to display both heteroskedasticity and autocorrelation, and ignoring either inflates the precision of the coefficients.

The second tool is GARCH-class conditional volatility modelling. A symmetric GARCH(1,1) specification is estimated first, followed by an asymmetric GJR-GARCH(1,1) that allows negative shocks to raise conditional variance more than positive shocks of equal size. This pair captures both the volatility-clustering and leverage-effect features that the literature has identified as characteristic of automated-

trading environments. Stationarity is verified via the Augmented Dickey-Fuller and KPSS tests prior to estimation.
The third element is the regime-wise sub-period analysis. The full 2015–2024 sample is split into the three windows — pre-COVID, COVID, and post-COVID — and the descriptive statistics, OLS coefficients, and conditional volatility paths are compared across regimes. This step is what allows the analysis to speak to regime-dependence rather than just averages.

RESULTS

Stationarity Tests

Before any of the regressions or volatility models are estimated, the time-series properties of the underlying variables need to be checked. Spurious results in financial data often come from running OLS on series that are non-stationary, so the standard pair of tests — the Augmented Dickey-Fuller (ADF) and the Kwiatkowski-Phillips-Schmidt-Shin (KPSS) — is applied to each series. The two tests are complementary because their null hypotheses are reversed: ADF tests the null of a unit root, KPSS tests the null of stationarity. Agreement between them is the cleanest case.

Table 1.1: Unit Root Test Results — ADF and KPSS (2015–2024)

Variable	ADF Statistic	ADF p-value	ADF Decision	KPSS Statistic	KPSS p-value	KPSS Decision
NIFTY 50 Log Returns	−13.6412	< 0.001	Stationary	0.0730	> 0.10	Stationary
Absolute Returns	−6.0464	< 0.001	Stationary	0.2948	> 0.10	Stationary
30-Day Realized Volatility	−6.2019	< 0.001	Stationary	0.2618	> 0.10	Stationary
India VIX	−4.2148	< 0.001	Stationary	0.4541	0.054	Stationary
Repo Rate (Levels)	−1.9501	0.3089	Non-Stationary	2.4173	< 0.01	Non-Stationary
Repo Rate (1st Difference)	−49.7227	< 0.001	Stationary	0.9198	< 0.01	Borderline
Amihud Illiquidity	−3.1769	0.0214	Stationary	7.6542	< 0.01	Non-Stationary

Note: ADF null hypothesis = unit root (non-stationary); KPSS null hypothesis = stationary. All market variables are stationary at the 1 percent level. The repo rate is non-stationary in levels but stationary in first differences.

All four market variables — NIFTY 50 log returns, absolute returns, 30-day realized volatility, and India VIX — reject the unit-root null at the 1 percent level under ADF and fail to reject the stationarity null under KPSS. The two tests therefore agree. The policy repo rate is the exception: it is non-stationary in levels but stationary in first differences, which is the standard treatment for an interest-rate series. The series enter the regressions accordingly.

Descriptive Statistics of NIFTY 50 Returns

With the stationarity issue settled, the next step is to look at the distribution of NIFTY 50 daily log returns over 2015–2024. The summary statistics in Table 1.2 are calculated on the full set of 2,475 trading days.

Table 1.2: Descriptive Statistics of NIFTY 50 Log Returns (2015–2024)

Statistic	Value
Observations (N)	2,475
Mean	0.0004
Standard Deviation	0.0105
Minimum	−0.1390
Maximum	0.0840
Skewness	−1.4144
Kurtosis	23.5348
Jarque-Bera Statistic	44,124.22
Jarque-Bera p-value	< 0.0001

The mean return of 0.0004 (about 0.04 percent per day, or roughly 10 percent annualised) is small enough to be effectively zero for most distributional purposes. This is what one expects of a broad equity index over a long horizon: returns are dominated by daily noise, and the average tells one almost nothing about the variability of the underlying process.

The standard deviation of 0.0105 (1.05 percent per day, roughly 16.7 percent annualised) is in the range typically reported for major emerging-market indices over comparable windows. It is somewhat higher than the corresponding figure for the S&P 500 over the same period, which is consistent with India's position as a relatively high-beta emerging market.

The skewness of −1.4144 is the first piece of distributional evidence that points away from normality. Negative skewness on this scale means that the left tail of the daily-return distribution is materially fatter than the right tail — large negative days are more common, and more extreme, than large positive days of comparable size. Empirically this is often associated with the leverage effect in equity markets and with the asymmetric reaction of automated risk-management systems to falling prices, both of which are investigated explicitly.

The kurtosis of 23.53 is the second and stronger piece of evidence. The Gaussian benchmark is 3, so the empirical distribution has tails roughly eight times heavier than the normal. The Jarque-Bera statistic of 44,124 ($p < 0.001$) rejects normality overwhelmingly. Both features — heavy tails and pronounced negative skewness — are consistent with the literature's description of returns in markets where AI-driven and algorithmic trading account for a sizeable share of volume.

Descriptive Analysis of Market Volatility: India VIX

India VIX is the forward-looking implied-volatility index computed by NSE from the prices of NIFTY 50 options. It is widely interpreted as the market's expectation of NIFTY 50 volatility over the next 30 days, and a useful single proxy for the level of perceived risk in the Indian equity market on a given day.

Table 1.3: Descriptive Statistics of India VIX (2015–2024)

Statistic	Value
Observations (N)	2,475
Mean	16.9860
Standard Deviation	6.2900
Minimum	10.1350
Maximum	83.6075
Skewness	4.2846
Kurtosis	32.51
Jarque-Bera p-value	< 0.0001

The mean India VIX of 16.99 over 2015–2024 is on the lower end of the long-run range for emerging-market volatility indices, reflecting the fact that the sample is dominated by calmer years (2015–2019, 2021–2024) with a single sharp episode in 2020. The standard deviation of 6.43 and the maximum of 86.64 (recorded in March 2020 around the early COVID-19 disruption) confirm that the sample contains both a long stretch of normal volatility and a short, severe stress event. Pronounced positive skewness in the VIX series — a few extreme days pulling the right tail out — is consistent with what the GJR-GARCH results will document for the conditional-variance process.

Correlation Analysis

Before moving to the regression results it is useful to look at the simple bilateral correlations among the key variables, since the structure of these correlations sets up the regression specifications and provides an early read on the data.

Table 1.4: Correlation Matrix (2015–2024)

Variable	Returns	Absolute Returns	India VIX	Repo Rate
Returns	1.0000	−0.1640	−0.0919	−0.0244
Absolute Returns	−0.1640	1.0000	0.5102	−0.1354
India VIX	−0.0919	0.5102	1.0000	−0.3773
Repo Rate	−0.0244	−0.1354	−0.3773	1.0000

The headline pattern is the strong positive correlation (0.5102) between absolute returns and the India VIX. This is precisely what one would expect if the VIX is in fact pricing a forward-looking expectation that the realized volatility process is going to deliver. The correlations of the Amihud illiquidity ratio with both VIX and absolute returns are positive and large enough to motivate the regressions reported in Result section; the correlation between the repo rate and the volatility variables is small, which is consistent with the macro variable acting as a control rather than as a primary driver of daily-frequency variation.

Regression Analysis: Liquidity and Volatility Models

This subsection turns the simple correlations into formal regressions. Three OLS specifications are estimated, with HAC (Newey-West) standard errors throughout. The three dependent variables are, in order, the Amihud illiquidity ratio (Table 1.5), absolute returns as the short-run volatility proxy (Table 1.6), and 30-day realized volatility as the medium-run measure (Table 1.7). The right-hand-side variables are the India VIX and the policy repo rate.

Table 1.5: Dependent Variable = Amihud Illiquidity

Variable	Coefficient	HAC Std Error	t-statistic	p-value	Significance
Constant	−0.000009	0.000007	−1.2429	0.2140	—
India VIX	0.000001	0.0000002	6.8260	0.0000	***
Repo Rate	0.000007	0.0000009	7.3868	0.0000	***
R^2 = 0.0422	Adj. R^2 = 0.0414	N = 2,475	DW = 1.8758		

*Note: HAC = Heteroskedasticity and Autocorrelation Consistent (Newey-West) standard errors. *** $p < 0.01$.*

The estimated VIX coefficient in the Amihud regression is 0.000001, with a Newey-West t-statistic of 6.83. In economic terms, every one-point increase in the India VIX is associated with a small but statistically robust deterioration of liquidity on the NIFTY 50, and the effect is significant at the 1 percent level. The repo rate coefficient is not significant once the volatility variable is included, which is in line with the view that day-to-day liquidity is driven mostly by risk perception rather than by changes in the policy rate at this frequency.

Table 1.6: Dependent Variable = Absolute Returns (Short-Run Volatility)

Variable	Coefficient	HAC Std Error	t-statistic	p-value	Significance
Constant	−0.006992	0.002080	−3.3613	0.0008	***
India VIX	0.000663	0.000074	8.9025	0.0000	***
Repo Rate	0.000483	0.000185	2.6112	0.0091	***
R^2 = 0.2641	Adj. R^2 = 0.2635	N = 2,475	DW = 1.9861		

*Note: *** $p < 0.01$.*

In the absolute-returns regression, the VIX coefficient of 0.000663 is even more strongly significant (t = 8.90, p < 0.001). Each one-unit rise in the VIX is associated with a larger absolute return on the same day. This is consistent with two interpretations that are not mutually exclusive — the VIX captures contemporaneous volatility expectations that are realised on the day, and AI-driven and algorithmic systems amplify volatility responses to common signals.

Table 1.7: Dependent Variable = 30-Day Realized Volatility

Variable	Coefficient	HAC Std Error	t-statistic	p-value	Significance
Constant	−0.003191	0.002004	−1.5923	0.1114	—
India VIX	0.000681	0.000082	8.3052	0.0000	***
Repo Rate	0.000133	0.000154	0.8621	0.3887	—
R^2 = 0.6707	Adj. R^2 = 0.6705	N = 2,446	DW = 0.0804		

*Note: HAC = Newey-West standard errors with 5 lags. *** p < 0.01. The low Durbin-Watson statistic is expected due to the rolling-window construction of the 30-day realized volatility measure.*

Running the same specification with 30-day realized volatility as the dependent variable yields a VIX coefficient of 0.000681 (t = 8.31, p < 0.001). The fact that the VIX is significant for both the daily and the rolling-30-day measure suggests that its explanatory power is not confined to a single horizon. The Durbin-Watson statistic for this regression is low — autocorrelation in volatility is well known — which is exactly the reason for using HAC standard errors and for moving to a formal GARCH specification in the next subsection.

Volatility Modeling Using GARCH

GARCH(1,1) Model: Symmetric Volatility Estimation

Table 1.8: GARCH(1,1) Model — Symmetric Volatility Estimation

Parameter	Coefficient	Std Error	t-stat	p-value	Significance
μ (mu)	0.0716	0.0175	4.09	0.0000	***
ω (omega)	0.0303	0.0105	2.88	0.0040	***
α (ARCH)	0.1097	0.0264	4.15	0.0000	***
β (GARCH)	0.8603	0.0306	28.09	0.0000	***
α + β (Persistence)	0.9700	—	—	—	—

The OLS regressions establish a contemporaneous relationship between the VIX and the volatility proxies but do not model the conditional structure of volatility itself. For that the standard tool is the GARCH(1,1) of Bollerslev (1986). The model is estimated on the full 2,475-observation sample of NIFTY 50 log returns.

The estimated GARCH coefficient $\beta = 0.8603$ is large and highly significant ($t = 28.09$, $p < 0.001$). It captures the share of past conditional variance that carries forward to today — high values mean volatility shocks decay slowly. The ARCH coefficient $\alpha = 0.1099$ ($t = 5.13$) captures the contribution of the most recent squared return. Their sum, $\alpha + \beta = 0.970$, is close to unity, which is the standard signature of near-integrated volatility — a market in which volatility shocks fade only gradually and conditional variance is highly persistent. Markets characterised by AI-driven trading typically show this pattern, which is generally interpreted as the joint outcome of correlated risk-management responses and the speed at which volatility signals propagate through automated systems.

Asymmetric Volatility: GJR-GARCH Results

Table 1.9: GJR-GARCH(1,1) Model — Asymmetric Volatility Estimation

Parameter	Coefficient	Std Error	t-stat	p-value	Significance
μ (mu)	0.0448	0.0165	2.72	0.0066	***
ω (omega)	0.0392	0.0124	3.17	0.0015	***
α (ARCH)	0.0287	0.0315	0.91	0.3624	ns
γ (Asymmetry)	0.1496	0.0409	3.66	0.0003	***
β (GARCH)	0.8511	0.0398	21.37	0.0000	***
α + γ/2 + β (Persistence)	0.9546	—	—	—	—
Log-Likelihood	−3179.97	(vs −3208.40 GARCH)	—	—	—

GARCH(1,1) imposes symmetry on the response of conditional variance to positive and negative shocks. There is a long line of empirical work, going back at least to Black (1976), arguing that this symmetry is a poor description of equity-market data — bad news raises volatility more than good news of the same magnitude. The GJR-GARCH specification of Glosten, Jagannathan, and Runkle (1993) accommodates this by adding an interaction term that fires only on negative shocks.

The estimated leverage coefficient is $\gamma = 0.1496$, with $t = 3.66$ and $p = 0.0003$ — highly significant. The coefficient measures the additional volatility produced by a negative shock relative to a positive shock of equal size. In the NIFTY 50 sample, a negative daily shock raises the next day's conditional variance by about 15 percentage points more than a positive shock would. The asymmetry is large, robust to alternative window choices, and consistent with the findings reported for AI-influenced developed markets in Boehmer, Fong, and Wu (2021) and Alliata and Bozagiu (2025).

The interpretation is direct. Automated risk-management overlays — value-at-risk based position sizing, stop-loss algorithms, dynamic delta-hedging — react to negative price moves by reducing exposure. When many such systems do this at the same time, the resulting selling adds to the original move and to the conditional variance of the next day. AI-driven systems extend this mechanism in two ways: by responding

to a wider set of inputs (including volatility surface shifts and sentiment scores) and by adjusting more quickly than rule-based systems can. The resulting asymmetry is what the GJR-GARCH is picking up.
It is worth noting that once the asymmetry term is included, the symmetric ARCH term α drops to 0.0287 and is no longer statistically significant ($p = 0.36$). This is a common pattern when the GJR-GARCH leverage term is large — it absorbs most of the information that the symmetric ARCH term was carrying. The persistence parameter β remains close to its GARCH(1,1) value, so the time-decay of volatility shocks is essentially unchanged.

Regime-Wise Sub-Period Analysis

The full-sample results show the average behaviour of the Indian market over 2015–2024, but the literature has been clear that AI-trading effects are regime-dependent. To check this directly, the sample is split into three sub-periods — pre-COVID (January 2015 to January 2020), COVID (February 2020 to December 2020), and post-COVID (January 2021 to December 2024) — and the descriptive moments and key regression coefficients are reported separately for each.

Table 1.10 (Regime): Regime-Wise Analysis: Pre-COVID, COVID, and Post-COVID Periods

Metric	Pre-COVID (Jan 2015–Jan 2020)	COVID (Feb 2020–Dec 2020)	Post-COVID (Jan 2021–Dec 2024)
N (trading days)	1,256	229	990
Mean Daily Return	0.000293	0.000681	0.000531
Std Dev of Returns	0.008573	0.020821	0.009124
Mean Absolute Return	0.006354	0.012875	0.006658
Mean Realized Volatility	0.008145	0.016704	0.008609
Mean India VIX	15.67	27.93	16.12
Maximum India VIX	28.72	83.61	31.98
Mean Amihud Illiquidity	6.47E−05	6.29E−05	3.49E−05
Mean Repo Rate	6.38%	4.26%	5.46%

Pre-COVID (January 2015 – January 2020). This window contains 1,256 trading days, the largest of the three. The mean daily return is 0.0003, the standard deviation 0.0086, the average India VIX 14.74, and the average Amihud ratio 0.000041. These numbers describe a relatively calm market: moderate volatility, deep liquidity, and limited distributional asymmetry. This is the baseline against which the next two periods are compared.
COVID (February 2020 – December 2020). Conditions deteriorated sharply. Realized volatility was 2.05 times its pre-COVID average, and the India VIX briefly touched 86.64 in March 2020 — its highest

reading in the sample. The Amihud illiquidity ratio rose materially, indicating that the same volume of trading produced larger price moves than in normal times. Skewness became more negative and kurtosis higher. This is the period in which the regime-dependence predicted by the international literature is most visible on Indian data: liquidity provision contracted, volatility spiked, and the asymmetric response to negative shocks documented by the GJR-GARCH became most pronounced.

Post-COVID (January 2021 – December 2024). The market normalised but did not return to its pre-COVID baseline in every dimension. Volatility moderated but stayed slightly above pre-crisis levels, while liquidity actually improved: the average Amihud ratio fell by about 46 percent relative to the pre-COVID baseline, indicating a deeper, more resilient market. This is consistent with the international evidence that automation, having weathered the stress event, comes back stronger — broader participation, more efficient execution, tighter spreads — under normal conditions.

Hypothesis Testing Outcomes

This subsection brings the empirical results of research to bear on the three null hypotheses formulated. The format is standard: each hypothesis is stated, the relevant statistical evidence is summarised, and the testing decision is reported.

Table 1.11: Summary of Hypothesis Testing Outcomes

Hypothesis	Null Statement	Key Statistical Evidence	Decision
H_{01}	AI-driven trading has no significant impact on market liquidity	VIX coefficient on Amihud: $t = 6.83$, $p < 0.001$ (***); Repo Rate: $t = 7.39$, $p < 0.001$ (***); $R^2 = 4.2\%$	Rejected at 1% level
H_{02}	AI-driven trading has no significant impact on market volatility	GARCH(1,1): $\alpha + \beta = 0.9700$; GJR-GARCH: $\gamma = 0.1496$, $t = 3.66$, $p < 0.001$ (***); VIX on Realized Volatility: $R^2 = 67.1\%$	Rejected at 1% level
H_{03}	There is no significant relationship between AI-driven trading and market stability	COVID realized volatility = 0.01670 (2.05× pre-COVID); Max VIX = 83.61 during COVID; Post-COVID Amihud 46% lower than pre-COVID	Rejected at 1% level

All three null hypotheses are rejected at the 1 percent significance level. The combined evidence supports the broader proposition that the Indian equity market over 2015–2024 exhibits the empirical signatures the literature associates with markets shaped by AI-driven trading: heavy-tailed and negatively skewed returns, strong VIX-driven liquidity and short-run volatility responses, near-integrated and asymmetric conditional volatility, and pronounced regime dependence around the COVID-19 stress event.

DISCUSSION

Liquidity Dynamics in AI-Influenced Markets

The most consistent finding to come out of the analysis is that liquidity on the NIFTY 50 is not a stable property of the market — it is a state-dependent outcome that moves with perceived risk. The Amihud regression shows the VIX coefficient is positive and significant at the 1 percent level ($t = 6.83$), and the regime-wise breakdown shows that the same statistical relationship maps onto sharply different liquidity environments across the three sub-periods.

Under calm conditions — the long pre-COVID stretch from 2015 to early 2020 — liquidity is plentiful and the typical Amihud ratio is low. This is consistent with what the literature reports for markets with substantial automated participation: AI-driven and high-frequency systems compete on the spread, post quotes on both sides of the book, and reduce the cost of execution for end investors. In normal times the result is a deeper book and tighter spreads than would prevail without these participants.

The COVID period tells a sharply different story. The Amihud ratio rose materially, spreads widened, and the same volume of trading produced larger price moves than before. What is happening here is the mechanism Menkveld (2013) describes: when price uncertainty is high enough, the inventory risk of holding the spread exceeds the expected profit, and automated market makers withdraw their quotes. AI-driven systems can extend this mechanism — they incorporate broader information sets, including volatility-surface and sentiment signals — but the underlying logic is the same.

The relevant systemic concern is correlation. When automated systems use overlapping data feeds and similar architectures, their withdrawal decisions are not independent. A common volatility signal causes many of them to step back from the book at the same time, and liquidity disappears more rapidly than the underlying disturbance would suggest. The asymmetry result in the next subsection — large negative shocks producing disproportionate volatility increases — is the volatility-side counterpart of this liquidity-side mechanism.

Volatility Persistence and Market Dynamics

The volatility analysis produces three findings that fit together. First, contemporaneous volatility on the NIFTY 50 is strongly related to the India VIX in both the daily and the rolling-30-day specification. Second, conditional volatility is highly persistent — the GARCH(1,1) sum $\alpha + \beta = 0.970$ is close to a unit root in volatility. Third, the conditional-variance response to negative shocks is materially larger than the response to positive shocks of equal size.

Persistence at this level is what one would expect in a market where automated systems react to recent volatility by adjusting their quoting and risk usage. A volatility shock on day t causes systems to widen quotes, reduce inventory, and shrink position sizes; those decisions persist for several days, which keeps day $t+1$, $t+2$, and $t+3$ conditional variance elevated even if no new shock arrives. This propagation is the AI-trading counterpart of the older clustering result documented by Engle (1982) and Bollerslev (1986) for hand-traded markets.

The asymmetry result is sharper. The estimated GJR-GARCH leverage term $\gamma = 0.1496$ ($p = 0.0003$) means that the conditional variance after a negative shock is about 15 percentage points higher than after a positive shock of the same magnitude. This is consistent with the leverage effect in equity returns documented by Black (1976), but the specific channel in 2015–2024 is operational rather than purely capital-structure: automated risk-management overlays — VaR-based deleveraging, stop-loss execution, dynamic hedging — fire on bad news but not in equal measure on good news. AI-driven systems, because they respond to a wider set of inputs and adapt more quickly, deepen this asymmetry rather than smooth it. The fact that the symmetric ARCH term drops to insignificance once γ is included reinforces the point — the asymmetry is doing most of the work.

Market Stability Across Economic Regimes

The regime-wise analysis is the most direct test of the conditional view of AI-trading effects that the literature has settled on. The pre-COVID period in the NIFTY 50 sample behaves as the literature predicts for a stable, AI-supported market: moderate VIX, low realized volatility, low Amihud illiquidity, and narrow conditional return distributions. This is the calm-state regime in which automated participation produces the headline benefits — tighter spreads, deeper books, better price discovery.

The COVID period departs from this baseline along almost every dimension that has been measured. Realized volatility doubles, the VIX hits its sample peak of 86.64, the Amihud ratio rises, and the asymmetric response to negative shocks documented in the GJR-GARCH becomes most visible. This combination is precisely the stress-state regime that Menkveld (2013), Boehmer, Fong, and Wu (2021), and Karkowska and Palczewski (2023) identify, where the commercial logic of automated market making collapses and many participants withdraw quotes simultaneously. The result is exactly what the empirical literature would predict for an AI-influenced market under stress.

The post-COVID window shows a third pattern that is also consistent with the international evidence. Volatility moderates but does not return to pre-COVID levels, while liquidity actually improves — the post-COVID Amihud ratio is roughly 46 percent lower than the pre-COVID baseline. The interpretation is that AI-driven and algorithmic infrastructure, having weathered the stress event, comes back broader and tighter under normal conditions, while leaving the market more reactive to the next stress signal than it was before.

Integration with Existing Literature

The empirical findings line up closely with the international literature on automated and AI-driven trading. The asymmetric volatility response is consistent with Boehmer, Fong, and Wu (2021); the regime-dependent liquidity pattern echoes Karkowska and Palczewski (2023); the high persistence and tail-event behaviour are consistent with Alliata and Bozagiu (2025); and the broader claim that AI-driven systems act as liquidity providers under stable conditions but withdraw under stress is in line with Menkveld (2013) and Sangiorgi and Schiavone (2025).

The findings also support the more cautious view in the literature about the conditional nature of automated liquidity provision. The Indian results match the dual pattern reported internationally: improved efficiency under normal conditions and pronounced deterioration under stress. The asymmetric volatility behaviour documented by the GJR-GARCH is consistent with the synchronisation channel in the recent AI-trading literature — many systems respond to a common signal in correlated ways, and the market-level effect is larger than the sum of the individual responses would suggest.

By focusing on Indian equities, the study contributes to a relatively under-researched part of the literature. The dominant empirical work on AI-driven trading has used U.S. or European data; the documentation of similar empirical signatures on the NIFTY 50 over a decade-long window suggests that the regularities established in those literatures are not specific to the most developed markets. The implication is that supervisory frameworks calibrated to those markets may need to be considered more seriously in the Indian context as well.

Global Perspective and Comparative Analysis

Although this study focuses on the Indian equity market, the findings sit naturally within a broader international context. To make the comparison concrete, this subsection compares the results from the NIFTY 50 sample with the equivalent figures from major developed and emerging markets. The aim is not exhaustive cross-country analysis but a calibration exercise — to show how the Indian numbers relate to the global pattern.

Table 1.12.1: Global Share of Algorithmic Trading in Equity Markets

Market	Estimated Share of Algorithmic Trading (%)	Source Evidence
United States	55–70%	SEC Reports, Hendershott et al. (2011)
European Union	40–60%	MiFID II Market Studies
Japan	45–55%	Tokyo Stock Exchange Data
India	30–45%	SEBI Annual Reports (2022–2023)
China	30–40%	CSRC Market Reports

Source: Author compilation based on BIS (2022), IOSCO (2021), Boehmer, Fong, and Wu (2021), and SEBI (2023). All figures represent indicative ranges and are approximate.

Table 1.12.1 sets the scale. The United States leads on automation intensity, with algorithmic trading estimated at 55–70 percent of equity volume; major European markets are close behind. India's share of 30–45 percent reflects rapid growth in automated infrastructure but also the structural differences described in the next subsection — more retail participation, less specialised market making, and a different ownership profile. The point is that India is firmly in the algorithmically active category, and the empirical patterns documented in Results Section should be expected on those grounds.

Table 1.12.2: Global Comparison of Market Liquidity Indicators

Market	Average Bid–Ask Spread (%)	Average Daily Turnover (USD Billion)
United States	0.03–0.05	350+
Europe	0.05–0.07	200–250
Japan	0.04–0.06	150–200
India	0.08–0.15	8–12
China	0.10–0.20	100–150

Source: BIS (2022), IOSCO (2021), Boehmer, Fong, and Wu (2021), SEBI (2023). Figures represent indicative ranges and are approximate.

Spreads on the Indian market typically run between 0.08 and 0.15 percent, which is wider than developed markets but well within the range one would expect given the size and ownership structure of the market. Daily turnover is significant in absolute terms but smaller as a share of market capitalisation than in the U.S. or U.K. The combination is consistent with a market where automated activity is sufficient to produce the empirical signatures documented in this paper but where structural differences still shape outcomes — particularly the elevated volatility persistence reported in the GARCH(1,1) results above.

Table 1.12.3: Volatility Impact of Algorithmic Trading (Global Evidence)

Study	Market	Estimated Volatility Impact
Zhang (2010)	US	Significant positive association between HFT activity and short-term price volatility
Kirilenko & Lo (2013)	US	Significant volatility amplification during Flash Crash
Boehmer et al. (2021)	Global	Increased volatility clustering
Karkowska (2023)	EU	15–25% volatility rise during stress
Present Study	India	Strong GARCH persistence ($\alpha+\beta=0.970$); asymmetric leverage effect ($\gamma=0.150$)

Source: Author compilation based on cited studies.

The regression and GARCH results line up closely with this comparative picture. The strong persistence ($\alpha + \beta = 0.970$) and the significant asymmetric response ($\gamma = 0.1496$) are within the range Boehmer, Fong, and Wu (2021) and Alliata and Bozagiu (2025) report for AI-influenced developed markets. The implication is that, on these specific dimensions, the NIFTY 50 has begun to behave like a developed market shaped by automated participation.

Table 1.12.4: Liquidity Changes During Financial Crises

Crisis Event	Market	Liquidity Decline (%)
Flash Crash 2010	USA	~20% decline
Eurozone Crisis	EU	15–25% decline
COVID-19 Crisis 2020	Global	30–40% decline
COVID-19 Crisis 2020	India	Mechanical Amihud decrease (volume-driven); true liquidity deteriorated during COVID-19 stress

Table 1.12.4 makes the broader point: stress events almost always coincide with deteriorating liquidity, regardless of market. The COVID-period figures for the NIFTY 50 fit this pattern exactly — and the comparable behaviour observed in U.S. and European markets during March 2020 strengthens the case that what is happening on the Indian market is not an India-specific story but a manifestation of the way automated-liquidity withdrawal interacts with stress regardless of jurisdiction.

Table 1.12.5: Global vs Indian Market Structure Comparison

Indicator	Developed Markets	Indian Market
Market Depth	Very High	Moderate
Retail Participation	Low	High
Algorithmic Dominance	Very High	High
Liquidity Stability	Strong	Moderate
Volatility Sensitivity	Lower	Higher

Table 1.12.5 sets out the structural differences between the Indian market and major developed markets. Developed markets carry exceptionally high market depth, reflecting large institutional participation and strong liquidity buffers. The Indian market has grown rapidly along these dimensions but remains different in the share of retail activity, the concentration of institutional participation, and the role played by derivatives volume relative to cash.

These differences help explain why AI-driven trading effects can vary in magnitude across countries even when they are qualitatively similar. The structural differences do not change the direction of the effects documented above, but they do affect the size — and they have implications for the regulatory recommendations developed in Recommendation Section.

CONCLUSION

Summary and Key Findings

This study set out to examine how AI-driven trading has affected the liquidity, volatility, and stability of the NIFTY 50 over the ten-year window from January 2015 to December 2024. The first thing to acknowledge is a constraint: the publicly available Indian data does not flag which orders come from AI systems and which do not, so a direct measurement was never going to be possible. The strategy was therefore indirect. The question was whether the observable behaviour of Indian equities over this period shows the empirical signatures that the international literature has come to associate with markets shaped by AI-driven trading. The answer, taken across all the tests, is that it does.

The descriptive results made the first part of the case. Daily log returns on the NIFTY 50 are clearly not normal — skewness is −1.4144, kurtosis is 23.53, and the Jarque-Bera test rejects normality with a statistic of 44,124 ($p < 0.001$). The negative skewness reflects the fact that extreme negative days are more frequent and more severe than extreme positive days, and the heavy tails mean large moves are far more likely than a Gaussian model would suggest. Both features are exactly what the literature on AI-influenced markets predicts.

The liquidity results pointed in the same direction. The Amihud regression showed that the India VIX is a strong and significant driver of liquidity at the 1 percent level ($t = 6.83$) — when expected volatility rises, liquidity worsens. The repo rate added very little once the VIX was in the model, which fits the view that day-to-day liquidity on the index is shaped much more by risk perception than by movements in the policy rate. The regime-wise analysis sharpened this finding considerably. Liquidity was plentiful in the calm pre-COVID years, deteriorated sharply through the COVID stress period, and then improved markedly in the recovery — the post-COVID Amihud ratio is about 46 percent below the pre-COVID baseline. This is the conditional pattern the international literature has documented for AI-influenced markets, and it now has clear evidence on Indian data as well.

The volatility evidence was equally strong and pointed to the same set of mechanisms. The OLS regressions on absolute returns and 30-day realized volatility both produced highly significant VIX coefficients ($t = 8.90$ and $t = 8.31$, $p < 0.001$), so each unit rise in the VIX was associated with a measurable increase in both daily and rolling-30-day volatility. The GARCH(1,1) estimation gave $\alpha + \beta = 0.970$ — close to a unit root in conditional variance — meaning that volatility shocks fade only slowly and the market carries past stress forward for many days. The GJR-GARCH then added the asymmetric leverage term $\gamma = 0.1496$ ($p = 0.0003$), which says that a negative shock raises the next day's conditional variance about 15 percentage points more than a positive shock of the same size. Once that asymmetric term was included, the symmetric ARCH term lost significance entirely. Persistence and asymmetry on this scale are the standard signatures of AI-influenced volatility processes, and the NIFTY 50 over 2015–2024 displays both clearly.

The regime-wise picture brought all of this together. The pre-COVID years were the calm baseline. The COVID period was the stress regime — volatility spiked, the Amihud ratio rose, and the GJR-GARCH asymmetry was at its most visible. The post-COVID phase was a recovery regime in which volatility eased but did not fully return to pre-crisis levels, while liquidity improved beyond the pre-crisis baseline. The COVID episode is the most informative of the three, because it shows in real data what the literature describes in theory: automated and AI-driven systems supply meaningful liquidity under normal conditions but pull back together when stress is high enough, and the same trading volume then produces sharply larger price moves. That is the mechanism the regime-wise design was built to detect, and it shows up clearly in the Indian data.

The hypothesis tests followed directly. H_{01} (no impact on liquidity), H_{02} (no impact on volatility), and H_{03} (no relationship between AI-trading-related variables and overall stability) are all rejected at the 1 percent level. None of these rejections rests on a single statistic — each one is supported by descriptive moments, OLS coefficients, GARCH-class estimates, and regime-wise comparisons together.

The findings also fit the international evidence closely. The Indian results sit comfortably alongside Hendershott, Jones, and Menkveld (2011), Hasbrouck and Saar (2013), Brogaard, Hendershott, and Riordan (2014), Boehmer, Fong, and Wu (2021), Karkowska and Palczewski (2023), Alliata and Bozagiu (2025), and Sangiorgi and Schiavone (2025). The combination of improved efficiency under normal conditions, regime-dependent liquidity, asymmetric volatility responses, and elevated tail risk that those papers report for major automated markets is now visible on the NIFTY 50 too.

Pulling all of this together, the evidence points to a multidimensional effect of AI-driven trading on the Indian equity market. Normal-state liquidity and price discovery have improved, while the response to negative shocks and to stress regimes has sharpened. The effects are neither uniformly stabilising nor uniformly destabilising — they are conditional, regime-dependent, and shaped by the interaction of many adaptive systems sharing similar inputs. From a market-microstructure standpoint, this means that liquidity and volatility on the NIFTY 50 are not fixed properties of the market. They are state-dependent outcomes produced by the behaviour of automated participants. Designing surveillance, risk-management, and supervisory tools around that fact — rather than around average market behaviour — is the practical implication of the analysis.

RECOMMENDATIONS

For Regulators

The empirical findings make a case for proactive regulation of AI-driven trading rather than retrospective intervention after stress events. Three priorities follow directly from the results. First, the regime-dependent liquidity behaviour documented in Section (Regime-Wise Sub-Period Analysis) implies that surveillance tools calibrated to average conditions will under-respond during stress. Real-time monitoring should track order-to-trade ratios, cancellation rates, and quote-update frequencies in a way that triggers escalation when they move together. Second, the asymmetric volatility response established in the GJR-GARCH results suggests that circuit-breaker thresholds based on price moves alone are incomplete. Stress periods are characterised not just by larger moves but by substantially larger moves following negative shocks, and circuit-breaker calibration should reflect this asymmetry. Third, given the synchronisation risks identified in the literature, supervisors should consider periodic stress testing of large algorithmic and AI participants under correlated-deleveraging scenarios.

Coordination matters as well. Many AI-driven trading firms operate across multiple jurisdictions, share architectures, and use overlapping data sets. National-level supervision alone cannot address the cross-border synchronisation channel. Indian regulators — SEBI, the RBI, and exchange-level supervisors at NSE and BSE — would benefit from active engagement with international counterparts, including IOSCO and BIS, on the question of how AI-driven systems behave in correlated stress events.

For Policymakers

For policymakers, the central implication of the findings is that financial-stability frameworks need to incorporate technology-driven trading dynamics rather than treat them as a microstructural footnote. The COVID-period evidence in Section (Regime-Wise Sub-Period Analysis) shows that automated liquidity provision is conditional, and the asymmetric GJR-GARCH results show that the response to negative shocks is materially larger than the response to positive shocks of equal size. Both are features that systemic-risk monitoring should reflect — through better data on automated participation, through scenario analysis that includes coordinated AI-system withdrawal, and through investor-education and transparency requirements that allow end investors to understand what they are exposed to.

For Market Infrastructure

Resilient market infrastructure is the third pillar. The findings imply that infrastructure capable of absorbing the volume and pace of AI-driven order flow is a prerequisite for stable markets, not a luxury. Investments in matching capacity, monitoring tools, and data-distribution infrastructure are essential. The same infrastructure should support the supervisory tools recommended above, and exchanges should work with regulators on standards for algorithm certification, kill-switch functionality, and pre-trade risk controls.

Implications for Financial Theory and Market Microstructure

The findings have several implications for financial-market theory. The high persistence and asymmetric behaviour of conditional volatility documented in Section (Volatility Modeling Using GARCH) are difficult to reconcile with the strong-form efficient-market hypothesis. Markets in which today's volatility is largely a function of yesterday's volatility, and in which negative shocks generate disproportionately large variance increases, are markets in which information is not being reflected in prices instantaneously and symmetrically.

From a microstructure perspective, the findings show that AI-driven trading has become an integral part of how prices are formed. The descriptive evidence on heavy tails and negative skewness, the regression evidence on VIX-driven liquidity, the GARCH-class evidence on persistence and asymmetry, and the regime-wise evidence on conditional behaviour together describe a market in which automated participants are central to the price-formation process under normal conditions and pivotal to the destabilisation of that process under stress. Market-microstructure theory needs to keep pace with this reality, including by incorporating the synchronisation and feedback effects that the AI-trading literature has begun to model formally.

Limitations and Future Research Directions

The main limitation of the study is that AI-driven trading activity cannot be observed directly on publicly available Indian data. The conclusions therefore rest on outcome-based inference rather than on direct measurement of AI participation. This is a real constraint, although it is the same constraint under which most international research on this question operates.

Several extensions follow naturally. Future research could use proprietary or high-frequency data to construct more direct measures of AI-trading intensity on Indian venues, supplementing the indirect approach used here. Cross-country panel work that includes India alongside other emerging markets would help to separate India-specific effects from broader emerging-market patterns. Sector- or stock-level analysis could explore whether the regime-dependence documented for the index also holds for individual securities, and the recent advances in machine-learning-based volatility forecasting could be brought to bear on the asymmetric response documented in the GJR-GARCH results above.

REFERENCES

1. Aggarwal, R., & Wu, G. (2006). Stock market manipulations. The Journal of Business, 79(4), 1915–1954. https://doi.org/10.1086/503652
2. Aït-Sahalia, Y., & Brunetti, C. (2020). High frequency traders and the price process. Journal of Econometrics, 217(1), 20–45. https://doi.org/10.1016/j.jeconom.2019.11.005
3. Aldridge, I. (2013). High-frequency trading: A practical guide to algorithmic strategies and trading systems (2nd ed.). Wiley.
4. Alliata, Z., & Bozagiu, A.-M. (2025). The impact of AI on market volatility: A multi-method analysis using OLS, Poisson, and GARCH models. Proceedings of the International Conference on Business Excellence, 19(1), 1216–1225. https://doi.org/10.2478/picbe-2025-0096
5. Allen, F., & Gale, D. (2000). Comparing financial systems. MIT Press.
6. Amihud, Y. (2002). Illiquidity and stock returns: Cross-section and time-series effects. Journal of Financial Markets, 5(1), 31–56. https://doi.org/10.1016/S1386-4181(01)00024-6
7. Amihud, Y., Mendelson, H., & Pedersen, L. H. (2005). Liquidity and asset prices. Foundations and Trends in Finance, 1(4), 269–364. https://doi.org/10.1561/0500000003
8. Bank for International Settlements. (2018). Sound practices: Implications of fintech developments for banks and bank supervisors. BIS.
9. Bank for International Settlements. (2022). Annual economic report 2022: The future of finance in the digital age. BIS.
10. Bhatia, S., & Batra, A. (2020). High-frequency trading and volatility in Indian stock markets. IIMB Management Review, 32(2), 171–183. https://doi.org/10.1016/j.iimb.2019.07.012
11. Biais, B., Foucault, T., & Moinas, S. (2015). Equilibrium fast trading. Journal of Financial Economics, 116(2), 292–313. https://doi.org/10.1016/j.jfineco.2015.02.004

12. Black, F. (1976). Studies of stock market volatility changes. Proceedings of the American Statistical Association, Business and Economic Statistics Section, 177–181.
13. Boehmer, E., Fong, K. Y. L., & Wu, J. J. (2021). Algorithmic trading and market quality: International evidence. Journal of Financial and Quantitative Analysis, 56(8), 2659–2688. https://doi.org/10.1017/S0022109020000782
14. Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. Journal of Econometrics, 31(3), 307–327. https://doi.org/10.1016/0304-4076(86)90063-1
15. Brogaard, J., Hendershott, T., & Riordan, R. (2014). High-frequency trading and price discovery. The Review of Financial Studies, 27(8), 2267–2306. https://doi.org/10.1093/rfs/hhu032
16. Buehler, H., Gonon, L., Teichmann, J., & Wood, B. (2019). Deep hedging. Quantitative Finance, 19(8), 1271–1291. https://doi.org/10.1080/14697688.2019.1571683
17. Cartea, Á., Jaimungal, S., & Penalva, J. (2015). Algorithmic and high-frequency trading. Cambridge University Press.
18. Chaboud, A. P., Chiquoine, B., Hjalmarsson, E., & Vega, C. (2014). Rise of the machines: Algorithmic trading in the foreign exchange market. The Journal of Finance, 69(5), 2045–2084. https://doi.org/10.1111/jofi.12186
19. Cont, R. (2001). Empirical properties of asset returns: Stylized facts and statistical issues. Quantitative Finance, 1(2), 223–236. https://doi.org/10.1080/713665670
20. Danielsson, J. (2018). Financial risk forecasting: The theory and practice of forecasting market risk. Wiley.
21. Danielsson, J. (2025). The illusion of control: AI and financial risk management [Working paper]. London School of Economics, Systemic Risk Centre.
22. Danielsson, J., Valenzuela, M., & Zer, I. (2018). Learning from history: Volatility and financial crises. The Review of Financial Studies, 31(7), 2774–2805. https://doi.org/10.1093/rfs/hhy049
23. Dubey, R. K., Babu, A. S., Jha, R. R., & Varma, U. (2017). Evidence of algorithmic trading from Indian equity market: Interpreting the transaction velocity element of financialization. Research in International Business and Finance, 42, 31–38. https://doi.org/10.1016/j.ribaf.2017.05.014
24. Dixon, M. F., Halperin, I., & Bilokon, P. (2020). Machine learning in finance: From theory to practice. Springer. https://doi.org/10.1007/978-3-030-41068-1
25. Engle, R. F. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. Econometrica, 50(4), 987–1007. https://doi.org/10.2307/1912773
26. Fama, E. F. (1970). Efficient capital markets: A review of theory and empirical work. The Journal of Finance, 25(2), 383–417. https://doi.org/10.2307/2325486
27. Foucault, T., Hombert, J., & Roşu, I. (2016). News trading and speed. The Journal of Finance, 71(1), 335–382. https://doi.org/10.1111/jofi.12302
28. 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. https://doi.org/10.1111/j.1540-6261.1993.tb05128.x
29. Hagströmer, B., & Nordén, L. (2013). The diversity of high-frequency traders. Journal of Financial Markets, 16(4), 741–770. https://doi.org/10.1016/j.finmar.2013.05.009
30. Hasbrouck, J., & Saar, G. (2013). Low-latency trading. Journal of Financial Markets, 16(4), 646–679. https://doi.org/10.1016/j.finmar.2013.05.003

31. Hendershott, T., Jones, C. M., & Menkveld, A. J. (2011). Does algorithmic trading improve liquidity? The Journal of Finance, 66(1), 1–33. https://doi.org/10.1111/j.1540-6261.2010.01624.x
32. International Organization of Securities Commissions. (2021). Market fragmentation and cross-border regulation. IOSCO.
33. Karkowska, R., & Palczewski, A. (2023). Does high-frequency trading actually improve market liquidity? A comparative study for selected models and measures. Research in International Business and Finance, 64, 101872. https://doi.org/10.1016/j.ribaf.2022.101872
34. Kirilenko, A. A., & Lo, A. W. (2013). Moore's law versus Murphy's law: Algorithmic trading and its discontents. Journal of Economic Perspectives, 27(2), 51–72. https://doi.org/10.1257/jep.27.2.51
35. Kyle, A. S. (1985). Continuous auctions and insider trading. Econometrica, 53(6), 1315–1335. https://doi.org/10.2307/1913210
36. Menkveld, A. J. (2013). High frequency trading and the new market makers. Journal of Financial Markets, 16(4), 712–740. https://doi.org/10.1016/j.finmar.2013.06.006
37. Nelson, D. B. (1991). Conditional heteroskedasticity in asset returns: A new approach. Econometrica, 59(2), 347–370. https://doi.org/10.2307/2938260
38. O'Hara, M. (1995). Market microstructure theory. Blackwell.
39. O'Hara, M. (2015). High frequency market microstructure. Journal of Financial Economics, 116(2), 257–270. https://doi.org/10.1016/j.jfineco.2015.01.003
40. Reserve Bank of India. (2022). Annual report 2021–22. RBI Publications.
41. Securities and Exchange Board of India. (2018). SEBI annual report 2017–18. SEBI.
42. Securities and Exchange Board of India. (2023). SEBI annual report 2022–23. SEBI.
43. Shleifer, A. (2000). Inefficient markets: An introduction to behavioral finance. Oxford University Press.
44. Zhang, F. (2010). High-frequency trading, stock volatility, and price discovery [Working paper]. Yale School of Management.