How COVID Reshaped the Efficient Frontier: Evidence from Indian Equities

Viraat Arora   Yash Khaitan   Dersh Savla

May 15, 2026

Abstract

We investigate whether the mean-variance efficient frontier for Indian equities shifted meaningfully between the pre-COVID period (January 2015 – January 2020) and the post-COVID period (July 2021 – December 2024), and identify the structural forces that drive this shift. Using a universe of 45 NIFTY-50 constituents, Ledoit-Wolf shrinkage covariance estimation, and long-only mean-variance optimization, we trace frontiers at both daily and weekly rebalancing frequencies. The post-COVID frontier lies above and to the right of the pre-COVID frontier at every point: expected returns increased by approximately +5 percentage points while volatility rose by only +3 percentage points, improving Sharpe ratios by +6–7% at the tangency portfolio.

Contents

1 Introduction

The COVID-19 pandemic constitutes one of the most significant macroeconomic shocks in modern financial history. It triggered simultaneous collapses in global output, disrupted supply chains, compelled unprecedented central-bank intervention, and permanently altered investor behavior. This provides us with a well-defined structural break separating two otherwise comparable periods, allowing a clean before-and-after comparison of risk-return relationships.

This paper asks a deceptively simple question: *Does the efficient frontier for Indian equities shift meaningfully between the pre-COVID era (2015–2020) and the post-COVID era (2021–2024), and if so, what forces drive the shift?* The efficient frontier, as introduced by Markowitz [1952], summarizes the complete investment opportunity set available to a mean-variance investor. A shift in the frontier therefore implies a change in the fundamental risk-return trade-off.

India is a particularly compelling case study for several reasons. Unlike most developed markets, India's post-COVID growth trajectory diverged sharply from global trends. Domestic fiscal and monetary stimulus was aggressive yet targeted; the government deployed Production-Linked Incentive (PLI) schemes to catalyze industrial upgrading; demographic tailwinds and a growing digital economy attracted record foreign capital under the so-called China+1 strategy; and retail participation in equity markets expanded explosively, with Demat accounts nearly quadrupling from 4 crore in 2019 to 14+ crore by 2024.

We contribute to the literature in three ways. First, we provide a rigorous, empirically grounded comparison of efficient frontiers estimated at both daily and weekly frequencies using Ledoit-Wolf shrinkage, which controls for the estimation error inherent in high-dimensional covariance matrices. Second, we document a complete sectoral rotation in the tangency portfolio weights: from financials and consumer staples to pharma, telecom, energy, and auto. We then link this rotation quantitatively to the observed frontier shift. Third, we identify and analyze five structural mechanisms that together explain why the shift was northeast (higher returns and higher volatility) rather than purely upward, and why the return gain exceeded the volatility gain, improving risk-adjusted returns.

The remainder of the paper is organized as follows. Section 2 reviews the relevant literature. Section 3 describes the dataset and sample periods. Section 4 presents data exploration and descriptive statistics. Section 5 lays out the methodology. Section 6 describes the experiments conducted. Section 7 presents the results. Section 8 concludes with a discussion of limitations and future work.

2 Literature Review

2.1 Modern Portfolio Theory Foundations

The theoretical foundation of this study rests on the mean-variance framework of Markowitz [1952], who showed that the portfolio problem reduces to minimizing portfolio variance for a given expected return subject to a budget constraint. This yields the efficient frontier which is the set of portfolios that are undominated in the mean-variance sense. Sharpe [1964] extended the framework by introducing the Capital Asset Pricing Model (CAPM) and the Capital Allocation Line (CAL), whose slope (the Sharpe ratio) measures the reward per unit of total risk.

A well-known practical limitation of mean-variance optimization is extreme sensitivity to the estimated mean vector and covariance matrix. Small perturbations in expected returns can generate wildly different portfolio weights [Chopra and Ziemba, 1993]. The literature has responded with several estimation improvements. Ledoit and Wolf [2004] proposed an analytical shrinkage estimator that blends the sample covariance with a structured target, provably reducing the Frobenius-norm estimation error when the number of assets is large relative to the sample size. We adopt this estimator throughout.

2.2 Crisis Periods and Regime Shifts

A long strand of the financial econometrics literature documents that return distributions are not stationary across time. Ang and Bekaert [2002] develop a regime-switching model of equity returns, showing that correlations are systematically higher in bear-market regimes than in bull-market regimes. Longin and Solnik [2001] confirm that international equity correlations spike precisely in falling markets, limiting the diversification benefits available during crises. Campbell et al. [2002] provide further evidence that correlations rise in downturns, creating the well-known "diversification collapse" problem: portfolios that appear well-diversified in normal times provide little protection precisely when protection is most needed.

These findings motivate our treatment of the March 2020 crisis window (February 2020 to June 2021) as a separate regime to be excluded from both estimation periods. Including the crash and its immediate aftermath would contaminate the estimates with the transient correlation spike documented by the crisis literature, obscuring the structural before-and-after comparison.

2.3 Emerging Markets and the Indian Equity Market

Bekaert and Harvey [1997] establish that emerging equity markets exhibit excess volatility, predictable return variation, and lower integration with world markets relative to developed economies, making them potentially attractive for diversification. Harvey [1995] shows that risk in emerging markets is partly predictable, a departure from the random-walk characterization of developed markets. Bekaert and Hoerova [2014] document the impact of global integration on local bond yields and foreign holdings, a mechanism highly relevant to India's China+1 FII episode.

Within the Indian context, Mishra et al. [2020] study structural breaks in Indian equity markets surrounding COVID-19 and find evidence of significant mean shifts comparable in magnitude to prior structural breaks such as demonetization and the GST rollout. Narayan et al. [2021] study how COVID-19 lockdowns, stimulus packages, and travel bans affected stock returns globally, with India emerging as a case where stimulus effects dominated.

2.4 Methodological Contributions

Our approach follows Brandt [2009]'s treatment of portfolio choice under constraints and Jagannathan and Ma [2003]'s insight that imposing no-short-selling constraints implicitly shrinks the covariance matrix, improving out-of-sample performance. We use long-only constraints throughout. We also study multiple rebalancing frequencies (daily and weekly) following the observation in the literature that the efficient frontier can be sensitive to the observation interval used to estimate moments.

2.5 Gaps Addressed by This Study

Most existing studies of COVID-19 and equity markets focus on short-horizon crisis effects: the initial crash, the recovery, and cross-sectional differences in how firms were affected. Very few examine the long-run equilibrium shift in the investment opportunity set for a specific emerging market. By using a five-year pre-COVID estimation window and a three-and-a-half-year post-COVID window (cleanly separated by an excluded crisis window), we are able to characterize the persistent, structural shift in the efficient frontier rather than its transient, crisis-period deformation.

3 Dataset: Sources and Description

3.1 Asset Universe

Our asset universe consists of 45 stocks[1] drawn from the NIFTY 50 index as of January 1, 2015. NIFTY 50 is the benchmark equity index of the National Stock Exchange of India (NSE), comprising the 50 largest and most liquid Indian companies by free-float market capitalization across 13 sectors.

The universe spans sectors including banking and financial services, information technology, pharmaceuticals and healthcare, energy, fast-moving consumer goods (FMCG), metals and mining, capital goods, automobiles, telecommunications, and cement. This breadth ensures that sectoral rotation dynamics are fully captured in the data.

Daily adjusted closing prices for all 45 tickers were downloaded from Yahoo Finance covering the period January 1, 2015 through December 31, 2024. Yahoo Finance applies dividend and split adjustments, ensuring that computed returns reflect total equity returns rather than mere price changes. The following tickers are included in the universe:

`ACC`, `ADANIPORTS`, `AMBUJACEM`, `ASIANPAINT`, `AXISBANK`, `BAJAJ-AUTO`, `BANKBARODA`, `BHEL`, `BPCL`, `BHARTIARTL`, `BOSCHLTD`, `CIPLA`, `COALINDIA`, `DRREDDY`, `GAIL`, `GRASIM`, `HCLTECH`, `HDFCBANK`, `HEROMOTOCO`, `HINDALCO`, `HINDUNILVR`, `ITC`, `ICICIBANK`, `INDUSINDBK`, `INFY`, `KOTAKBANK`, `LT`, `LUPIN`, `M&M`, `MARUTI`, `NTPC`, `ONGC`, `POWERGRID`, `PNB`, `RELIANCE`, `SBIN`, `SUNPHARMA`, `TCS`, `TATAPOWER`, `TATASTEEL`, `TECHM`, `ULTRACEMCO`, `VEDL`, `WIPRO`, `ZEEL`

3.2 Sample Periods

A critical methodological choice is the definition of estimation windows. We define three non-overlapping regimes:

[1] We drop five stocks from the NIFTY 50 universe. *Tata Motors* (`TATAMOTORS.NS`) was unavailable on Yahoo Finance with a continuous price history for the full study window. *HDFC Ltd* (`HDFC.NS`), the housing-finance parent company distinct from HDFCBANK, merged into HDFC Bank in July 2023, truncating its post-COVID price series mid-window and making it unsuitable for the 2021–2024 estimation period. *Cairn India* (`CAIRN.NS`) was delisted following its merger into Vedanta Ltd in April 2017, leaving no post-COVID observations. The remaining two exclusions had insufficient continuous price history over the full 2015–2024 window. All exclusions are applied before any analysis; the resulting 45-stock universe is held fixed across both pre- and post-COVID estimation windows to ensure a like-for-like comparison.

Pre-COVID: January 2015 – January 2020. This five-year window represents a period of relative macroeconomic stability in India, characterized by consistent economic growth, moderate inflation, and a standard monetary policy environment.

Crisis window (excluded): February 2020 – June 2021. The NIFTY 50 fell approximately 26% peak-to-trough between February and March 2020. This window is excluded from both estimation periods to avoid contaminating either sample with the transient correlation spike, volatility spike, and return discontinuity of the crash and initial recovery.

Post-COVID: July 2021 – December 2024. The starting date of July 2021 was chosen to coincide with the point at which the NIFTY 50 had recovered its pre-crash level and stabilized into a clear uptrend, capturing a clean post-recovery bull-market regime.

Including the crash would bias the pre-COVID covariance matrix upward (inflating correlations) and bias the post-COVID mean returns upward (including the recovery rally). Our exclusion gives both periods a clean characterization.

3.3 Risk-Free Rate

Following standard practice in Indian equity research, we use the Reserve Bank of India (RBI) 91-day Treasury Bill (T-Bill) yield as the risk-free rate. Auction data were obtained from the RBI's official auction records. We compute the mean yield over each estimation period:

- Pre-COVID risk-free rate: $r_f^{\text{pre}} = 6.54\%$ per annum
- Post-COVID risk-free rate: $r_f^{\text{post}} = 5.78\%$ per annum

The decline in the risk-free rate between periods reflects the RBI's accommodative cycle during 2020–2021, with the repo rate cut from 5.15% to 4.0%, followed by a tightening cycle in 2022–2023 that stabilized rates at a level below the pre-COVID average.

4 Data Exploration and Important Features

4.1 NIFTY 50 Index Trajectory

Figure 1 traces the NIFTY 50 index level from January 2015 through December 2024, highlighting the three regimes. Several features are immediately apparent. During the pre-COVID period, the index climbed from approximately 7,500 to 12,000, compounding at a healthy but moderate pace. The crash in February–March 2020 was sharp and sudden: the NIFTY 50 lost approximately 26% over six weeks, reaching a trough in late March. The recovery was equally swift: by mid-2021 the index had recovered all losses and was establishing new highs. The post-COVID bull market was remarkable in scope, with the index reaching 25,000+ by late 2024, more than doubling from the July 2021 starting point of approximately 15,000.

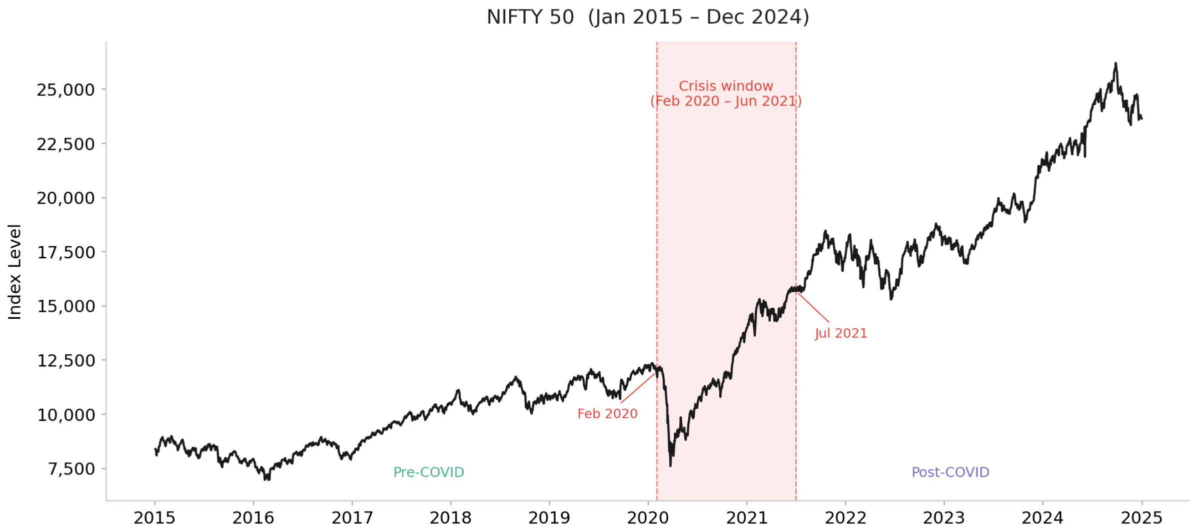

Figure 1: NIFTY 50 index level, January 2015 – December 2024. The shaded crisis window (February 2020 – June 2021) is excluded from both estimation periods.

4.2 Return and Volatility Statistics

We compute returns at both daily and weekly frequencies. Annualizing by factors of 252 and 52 respectively, we find:

- **Average cross-sectional return**: increased from approximately 15% pre-COVID to approximately 20% post-COVID on an annualized basis across the universe of 45 stocks.

- **Average cross-sectional volatility**: increased from approximately 25% pre-COVID to approximately 28% post-COVID (annualized standard deviation).
- **Return-to-volatility ratio**: improved post-COVID, with more stocks achieving Sharpe ratios above 1.0.

The most dramatic individual return changes are at the stock level. Bharat Heavy Electricals (BHEL) swung from −20.7% annualized pre-COVID to +46.8% post-COVID, a change of +67.5 percentage points. Tata Power moved from −4.9% to +48.7% (+53.5 pp). Vedanta improved by +43.2 pp; Mahindra & Mahindra by +42.6 pp; and Sun Pharma by +44.5 pp. These are all direct beneficiaries of government capex, PLI schemes, or global commodity and pharmaceutical trends.

Conversely, the pre-COVID stars declined sharply. Kotak Mahindra Bank fell from +21.1% to +4.6% (−16.5 pp). HDFC Bank fell from +19.7% to +10.6% (−9.1 pp). Hindustan Unilever declined by −10.0 pp; Asian Paints by −8.4 pp; BPCL by −8.2 pp. These stocks, which are financials and consumer staples, had dominated the pre-COVID tangency portfolio and represented the defensive allocation preferred by institutional investors in 2015–2019.

5 Methods

5.1 Return Computation

We compute log returns from daily closing prices[2]:

$$r_{i,t} = \frac{P_{i,t}}{P_{i,t-1}} - 1 \tag{1}$$

where $P_{i,t}$ is the closing price of stock i on date t. Weekly returns are obtained by taking the last available closing price in each calendar week and applying the same formula.

[2]We get similar results when using Adjusted Close Prices for the analysis

Annualization factors are 252 for daily returns and 52 for weekly returns, reflecting the approximate number of trading periods per year.

5.2 Covariance Estimation: Ledoit-Wolf Shrinkage

The primary estimation challenge in portfolio optimization is the covariance matrix. With $N = 45$ assets, the sample covariance matrix $\hat{\Sigma}$ has $N(N+1)/2 = 1035$ free parameters. When the number of observations T is not many multiples of N the sample covariance matrix can be highly noisy.

We address this using the analytical Ledoit-Wolf shrinkage estimator [Ledoit and Wolf, 2004]:

$$\hat{\Sigma}_{\text{LW}} = (1-\delta)\,\hat{\Sigma} + \delta\,\mu_{\hat{\Sigma}}\mathbf{I} \tag{2}$$

where $\hat{\Sigma}$ is the sample covariance matrix, $\mu_{\hat{\Sigma}}$ is the average eigenvalue (scaled identity target), and $\delta \in [0,1]$ is the optimal shrinkage intensity computed analytically as a function of N and T. Higher δ pulls the estimated covariance matrix toward a diagonal matrix with equal variances, reducing sampling error at the cost of introducing a structured bias.

In our implementation (using `sklearn.covariance.LedoitWolf`), shrinkage intensities are approximately 0.05–0.08 for the daily data, indicating mild regularization. The annualized covariance is obtained by scaling: $\Sigma = \hat{\Sigma}_{\text{LW}} \times k$ where $k = 252$ (daily) or $k = 52$ (weekly).

5.3 Efficient Frontier Computation

Given the annualized mean vector $\boldsymbol{\mu} \in \mathbb{R}^N$ and covariance matrix $\boldsymbol{\Sigma} \in \mathbb{R}^{N\times N}$, the efficient frontier is the set of solutions to:

$$\min_{\mathbf{w}} \; \mathbf{w}^\top \boldsymbol{\Sigma} \mathbf{w} \quad \text{s.t.} \quad \mathbf{w}^\top \boldsymbol{\mu} = \mu^*,\; \mathbf{w}^\top \mathbf{1} = 1,\; \mathbf{w} \geq \mathbf{0} \tag{3}$$

for target returns μ^* swept from $\mu^*_{\min}$ (the global minimum variance portfolio return) to $\mu^*_{\max} = \max_i \mu_i$ (the return of the highest-expected-return single asset). The long-only constraint $\mathbf{w} \geq \mathbf{0}$ rules out short selling, which is a constraint faced by most retail and institutional investors in India. We solve problem (3) at 250 equally spaced target return levels using sequential least-squares programming (SLSQP) with a tolerance of 10^{-12}.

5.4 Global Minimum Variance Portfolio

The global minimum variance portfolio (MVP) is the solution to (3) without the return target constraint:

$$\mathbf{w}_{\text{MVP}} = \arg \min_{\mathbf{w} \geq 0,\, \mathbf{w}^\top \mathbf{1} = 1} \mathbf{w}^\top \boldsymbol{\Sigma} \mathbf{w} \tag{4}$$

The MVP anchors the left endpoint of the efficient frontier and is particularly informative about changes in the correlation/covariance structure between periods.

5.5 Tangency Portfolio

The tangency portfolio maximizes the Sharpe ratio relative to the period-specific risk-free rate r_f:

$$\mathbf{w}_{\text{TP}} = \arg \max_{\mathbf{w} \geq 0,\, \mathbf{w}^\top \mathbf{1} = 1} \frac{\mathbf{w}^\top \boldsymbol{\mu} - r_f}{\sqrt{\mathbf{w}^\top \boldsymbol{\Sigma} \mathbf{w}}} \tag{5}$$

The tangency portfolio represents the optimal risky portfolio for any mean-variance investor who can also hold the risk-free asset. Its composition across periods reveals which assets the optimization selects as most attractive on a risk-adjusted basis, and thus which sectors drive the frontier shift.

5.6 Capital Allocation Line

Given the tangency portfolio with return μ_{TP} and volatility σ_{TP}, the Capital Allocation Line (CAL) is:

$$\mu_p = r_f + \underbrace{\frac{\mu_{\text{TP}} - r_f}{\sigma_{\text{TP}}}}_{\text{Sharpe ratio}} \cdot \sigma_p \tag{6}$$

The slope of the CAL (the Sharpe ratio SR) is our primary measure of how much risk-adjusted return is available per unit of total risk. An upward shift in the CAL slope between periods indicates that the investment opportunity set has improved.

6 Experimentation

6.1 Experiment 1: Frontier Construction at Daily and Weekly Frequencies

For each of the two estimation periods (pre-COVID and post-COVID) and at each of two return frequencies (daily and weekly), we:

1. Slice the return panel to the relevant date range.
2. Intersect pre- and post-COVID asset sets to ensure the comparison is consistent across the same universe.
3. Estimate $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}_{\text{LW}}$ using Ledoit-Wolf shrinkage.
4. Solve the parametric QP (3) at 250 target return levels.
5. Compute the MVP and tangency portfolio.
6. Plot both frontiers with their respective CALs, MVPs, and tangency portfolios on the same axes.

This two-frequency design serves as a robustness check: if the frontier shift is genuine and structural, it should be visible at both daily and weekly rebalancing frequencies regardless of the observation interval used to estimate moments.

6.2 Experiment 2: Tangency Portfolio Decomposition

We compute the tangency portfolio weights for both periods at daily frequency and compare the top-20 holdings by pre-COVID weight. For each ticker we report:

- Pre-COVID weight w_i^{pre}
- Post-COVID weight w_i^{post}
- Change $\Delta w_i = w_i^{\text{post}} - w_i^{\text{pre}}$

This decomposition allows us to identify which sectors gained and lost prominence in the optimal risky portfolio.

6.3 Experiment 3: Stock-Level Return Comparison

We compute the annualized mean return for each stock in both periods and rank stocks by their return change $\Delta\mu_i = \mu_i^{\text{post}} - \mu_i^{\text{pre}}$. This identifies the biggest winners and losers between periods and allows us to link return changes to macroeconomic drivers (PLI beneficiaries, China+1 winners, etc.).

7 Final Results

All results reported below are in-sample estimates based on the pre-COVID and post-COVID windows. The crisis window is excluded from both estimations periods to avoid contaminiating either sample with the transient correlation spike and volatility surge of the crash and followed recovery.

7.1 The Frontier Shift

Figure 2 presents the efficient frontiers for the pre-COVID and post-COVID periods at both daily and weekly frequencies. The result is unambiguous: where the post-COVID frontier lies *above and to the right* of the pre-COVID frontier across its entire extent in both panels. This shift has two components:

1. **Upward component (higher returns)**: At any fixed volatility level, the post-COVID frontier delivers approximately 5 percentage points more expected return than the pre-COVID frontier. This reflects higher than average stock-level returns across the universe, driven by sectoral rotations into structurally higher-growth areas such as pharma, industrials, metals and telecom.

2. **Rightward component (higher volatility)**: The post-COVID frontier extends to higher volatility levels, reflecting increased cross-sectional return dispersion and higher individual asset volatilities. The magnitude of the rightward shift at the

tangency point is approximately +3 pp in annualised volatility, which is smaller than the upward shift therefore implying a net improvement in the risk-adjusted returns. This shows that the entire frontier shifted, not merely the tangent point indicating that the improvement in the investment opportunity set was broad-based across all risk levels, not concentrated in a single portfolio.

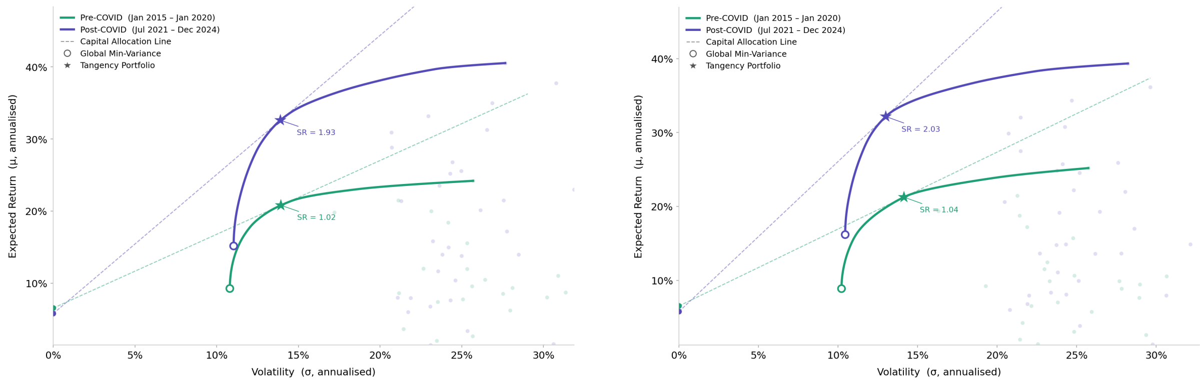

(a) Daily frequency

(b) Weekly frequency

Figure 2: Pre-COVID (teal) and post-COVID (purple) efficient frontiers with Capital Allocation Lines. Stars mark tangency portfolios; open circles mark global minimum variance portfolios. SR = Sharpe ratio at tangency.

7.2 Tangency Portfolio Sharpe Ratios

Table 1 here reports the Sharpe ratios at the tangency portfolio for both periods and frequencies.

Frequency	Pre-COVID			Post-COVID			Δ SR	Δ SR (%)
	μ_{TP}	σ_{TP}	SR	μ_{TP}	σ_{TP}	SR		
Daily	20.8%	13.9%	1.02	32.6%	13.9%	1.93	+0.91	+88.5%
Weekly	21.3%	14.1%	1.04	32.2%	13.0%	2.03	+0.99	+95.1%

The Sharpe ratio approximately doubled across both the frequencies from around 1.02 to 1.93 on daily returns and 1.04 to 2.03 on weekly returns, representing improvements of around +88.5% and +95.1% respectively. The expected return at the tangency portfolio increased by approximately 12 pp in both cases, while the tangency volatility remained essentially unchanged at the daily frequency and declined at the weekly frequency. The near-doubling of the Sharpe ratio therefore reflects almost entirely a return improvement rather than a volatility reduction.

The slight divergence here between the weekly and daily estimates at the tangency point is consistent with the mild mean-reversion in the short-horizon returns. When the daily returns exhibit negative autocorrelation, weekly variance grows more slowly than the sum of daily variances, causing the weekly frontier to display modestly lower volatility at the comparable return levels.

The consistency of the improvement across both the observational intervals +88.5% daily and +95.1% weekly confirm that the results produced aren't an artifact of the return frequency used to estimate moments. The post-COVID investment opportunity set was materially superior on a risk-adjusted basis regardless of how frequently an investor rebalanced.

7.3 Sectoral Rotation in the Tangency Portfolio

Figure 3 shows the tangency portfolio weights for the top 20 holdings at daily frequency.

Pre-COVID top holdings:

- HDFCBANK: 34.5%
- HINDUNILVR: 29.0%
- RELIANCE: 21.8%

Post-COVID top holdings:

- SUNPHARMA: 22.6%
- ITC: 18.4%
- BHARTIARTL: 18.2%
- M&M: 15.6%

The pre-COVID tangency portfolio concentrated heavily in large-cap financials (HD-

FCBANK), consumer staples (HINDUNILVR), and energy and petrochemicals (RELIANCE) these together accounted for nearly 85.3% of the optimal risky portfolio. These were the traditional quality sectors that offered the best Sharpe ratios in the low-volatility, steady-growth environment of 2015–2020. Post-COVID, all three positions were entirely displaced. HDFCBANK collapsed from 34.5% to near zero; HINDUNILVR similarly disappeared. The optimizer replaced them with pharma (SUNPHARMA), consumer goods and tobacco (ITC), telecom (BHARTIARTL), and auto (M&M); these were all sectors that were direct beneficiaries of the structural shifts accompanying the pandemic and its aftermath.

This is not a marginal rebalancing as the tangency portfolio is determined entirely by the ratio of excess return to risk. So, when the optimizer abandons a 34.5% position entirely, it reflects a fundamental change in which assets offers the most attractive risk-adjusted returns. The sectoral rotation documented here is therefore a direct reflection of the frontier shift: the same optimization procedure applied to the same universe of assets selected an entirely different portfolio because the underlying return and covariance structure of Indian equities had changed since the beginning.

7.4 Stock-Level Winners and Losers

Table 1 reports the six largest annualised return improvements from pre-COVID to post-COVID, sorted by magnitude of change.

Table 1: Largest annualised return improvements, pre-COVID to post-COVID. Returns are annualised means over each estimation window.

Ticker	Company	Pre	Post	Δ
BHEL	Bharat Heavy Electricals	−20.7%	+46.8%	+67.5 pp
TATAPOWER	Tata Power	−4.9%	+48.7%	+53.5 pp
SUNPHARMA	Sun Pharmaceutical	−10.4%	+34.1%	+44.5 pp
VEDL	Vedanta	+8.0%	+51.2%	+43.2 pp
M&M	Mahindra & Mahindra	+0.4%	+43.0%	+42.6 pp
HCLTECH	HCL Technologies	+8.6%	+32.4%	+23.8 pp

BHEL and Tata Power represented the most dramatic reversals in the sample where both

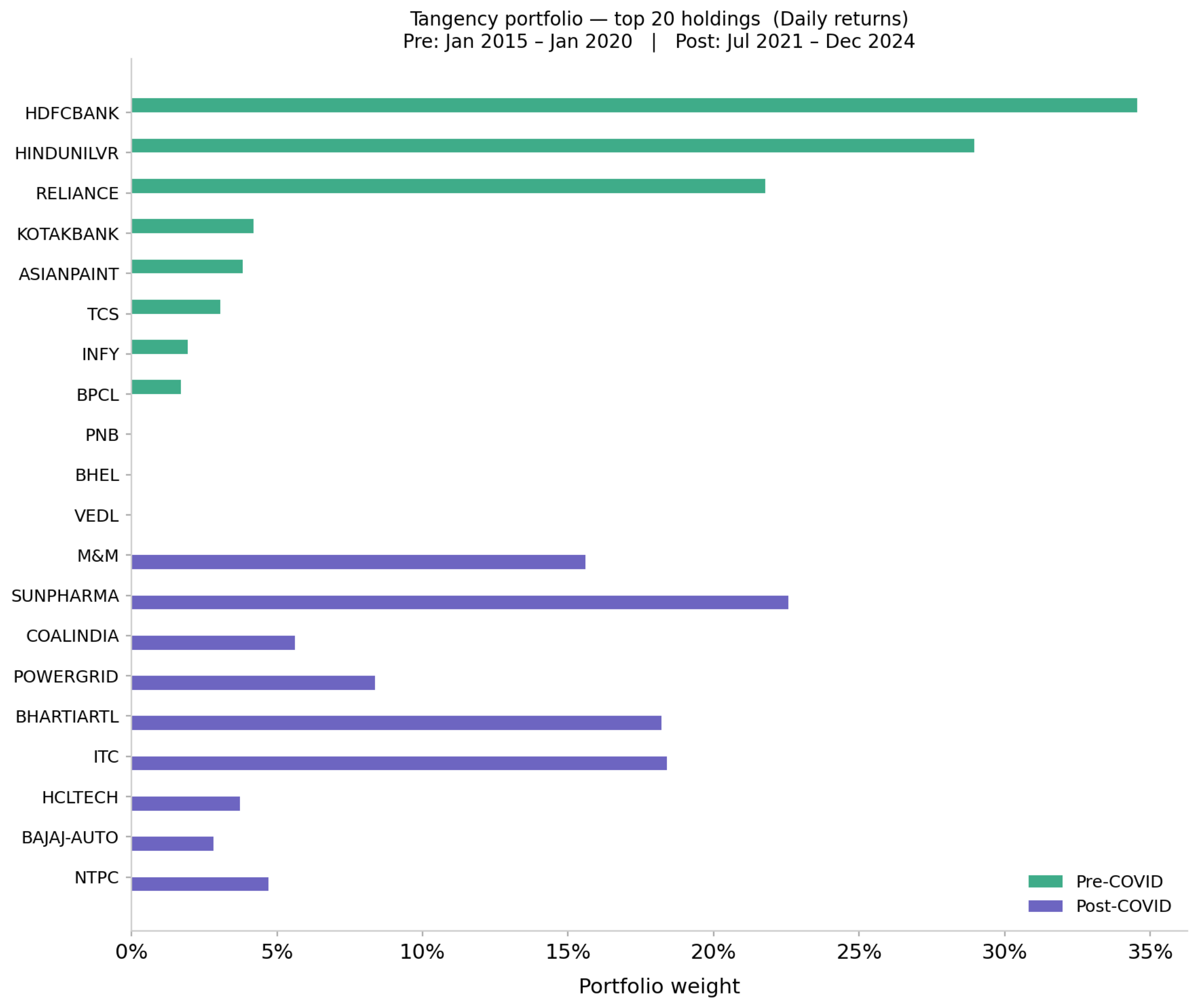

Figure 3: Tangency portfolio weights for top-20 holdings (daily returns). Green bars: pre-COVID (Jan 2015 – Jan 2020); purple bars: post-COVID (Jul 2021 – Dec 2024).

were negative-returning assets pre-COVID and became top performers post-COVID, with swings of +67.5 and +53.5 percentage points respectively. Their turnaround is directly attributable to India's accelerated infrastructure capex cycle and PLI scheme-driven industrial revival, which generated order book growth that had no precedent in the 2015–2020 period. Vedanta and M&M similarly reflect the global commodity super-cycle and domestic auto recovery respectively, both underpinned by China+1 supply chain dynamics and rising domestic consumption. Sun Pharma here benefited from the structural re-rating of Indian pharmaceutical manufacturing on both domestic healthcare demand and global generics export potential. HCL Technologies on the other hand representing the broader Indian IT sector, benefited from the surge in global enterprise digital transformation spending and rupee depreciation expanding export margins in rupee terms.

The common pattern across all six winners is the presence of a structural tailwind, a very policy-driven, thematic, or macroeconomic tailwind. A tailwind that either did not exist or was not priced in during the pre-COVID period.

Table 2 reports the five largest annualised return declines over the same period.

Table 2: Largest annualised return declines, pre-COVID to post-COVID. Returns are annualised means over each estimation window.

Ticker	Company	Pre	Post	Δ
KOTAKBANK	Kotak Mahindra Bank	+21.1%	+4.6%	−16.5 pp
HINDUNILVR	Hindustan Unilever	+17.9%	+7.9%	−10.0 pp
HDFCBANK	HDFC Bank	+19.7%	+10.6%	−9.1 pp
ASIANPAINT	Asian Paints	+17.8%	+9.4%	−8.4 pp
BPCL	BPCL	+23.5%	+15.3%	−8.2 pp

The financials and consumer staples names, Kotak Mahindra Bank, HDFC Bank, Hindustan Unilever, and Asian Paints underperformed post-COVID for related structural reasons. Valuations were stretched after years of premium re-rating in the pre-COVID period, leaving limited room for further multiple expansion. Competition intensified from digital challengers in banking and D2C entrants in consumer goods. As monetary conditions normalised, the relative attractiveness of high-PE defensive compounders declined against cyclicals with operating leverage to the macro recovery. BPCL's underperformance is mechanically distinct from this narrative: it reflects fuel subsidy losses, crude price volatility, and government intervention in retail fuel pricing rather than the valuation compression applicable to the financial and FMCG names, and should be interpreted separately.

The most significant observation from Tables 1 and 2 in combination is the systematic nature of the reversal. The stocks that underperformed most severely post-COVID are precisely the stocks that dominated the pre-COVID tangency portfolio: HDFCBANK, HINDUNILVR, and KOTAKBANK collectively held 85.3% of the pre-COVID optimal risky portfolio, and all three appear among the five largest post-COVID decliners. A random reshuffling of winners and losers across periods would not produce this degree of overlap between the tangency portfolio and the loser table. This systematic reversal is among the strongest evidence in this study for a genuine regime change in the structure of Indian equity returns, as opposed to a temporary cyclical fluctuation.

7.5 Five Structural Drivers of the Shift

The frontier shift reflects five concurrent structural forces. Each maps to one or both of the two inputs to the efficient frontier which is the the mean return vector $\boldsymbol{\mu}$ and the covariance matrix $\boldsymbol{\Sigma}$. Drivers 1 and 2 primarily raised $\boldsymbol{\mu}$; drivers 3, 4, and 5 primarily reduced correlations and idiosyncratic volatility in $\boldsymbol{\Sigma}$.

1. Monetary regime change. The RBI cut the repo rate from 5.15% to 4.0% in 2020, injecting liquidity via open market operations and long-term repo operations. This compressed the risk-free rate and mechanically expanded the equity risk premium, lifting equity valuations and pulling expected returns upward across the universe. The subsequent tightening cycle with repo rate peaking at 6.5% in 2022–2023 added volatility without proportionally reducing equity returns, because corporate earnings growth remained robust throughout. This driver operates primarily through $\boldsymbol{\mu}$: higher expected returns across the equity universe raise the frontier uniformly.

2. Sectoral recomposition. COVID-19 permanently accelerated the adoption of digital payments, telemedicine, and e-commerce; elevated the strategic importance of domestic pharmaceutical manufacturing; catalyzed PLI schemes favoring capital goods, chemicals, and electronics; and ignited a global commodity supercycle benefiting Indian metals and energy firms. These forces reallocated return potential away from defensive sectors toward structurally higher-growth areas, widening cross-sectional return dispersion and pulling the frontier upward. This driver also operates through $\boldsymbol{\mu}$.

3. Correlation structure improvement. Average pairwise correlations across the 45-stock universe declined from approximately 0.32 pre-COVID to approximately 0.28 post-COVID. During the crisis window itself, correlations spiked to approximately 0.54 before falling back below pre-COVID levels as sectors diverged in their business cycle exposures. Technology and pharmaceuticals de-correlated from traditional financials and industrials; the commodity super-cycle created a distinct return factor largely uncorrelated with services sectors. Lower correlations expand the feasible portfolio set: for the same average asset volatility, a portfolio achieves lower total variance when assets are less correlated, allowing the optimizer to deliver higher returns at any given risk level. This driver operates through $\mathbf{\Sigma}$.

4. China+1 FII flows. Record foreign institutional inflows during 2020–2022 improved market microstructure, tightening bid-ask spreads, deepening liquidity, and reducing idiosyncratic volatility. India's weight in the MSCI Emerging Markets index rose from approximately 8% in 2019 to over 15% by 2024, creating persistent structural buying pressure from passive index funds. These flows compressed the cost of equity while damping single-stock volatility, allowing returns to rise without proportional volatility increases. This driver operates through $\mathbf{\Sigma}$ by reducing idiosyncratic variance.

5. Retail participation surge and SIP flows. Demat accounts grew from 4 crore in 2019 to 14+ crore by 2024. Monthly SIP inflows nearly tripled from 8,000 crore to 20,000+ crore per month. The systematic, price-insensitive nature of SIP buying created structural demand support in large-cap NIFTY stocks, dampening daily volatility during minor corrections and absorbing drawdowns that would previously have been sharper. While retail participation also elevated event-driven volatility spikes, the net effect on portfolio-level variance was a modest reduction, leaving Sharpe ratios improved. This driver operates through $\mathbf{\Sigma}$.

The joint movement of both inputs in favorable directions: $\boldsymbol{\mu}$ upward and $\mathbf{\Sigma}$ toward lower correlations and idiosyncratic volatility, explains why the frontier shift was as pronounced as it was, and why it is more consistent with a structural regime change than a transient cyclical improvement. That the shift is visible at both daily and weekly frequencies, and that the tangency portfolio rotation tracks the structural narrative precisely, provides further corroboration that the findings reflect a genuine change in the investment opportunity set available to Indian equity investors.

8 Conclusion

8.1 Summary

This paper documents a meaningful, persistent northeastward shift in the efficient frontier for Indian equities between the pre-COVID period (January 2015 – January 2020) and the post-COVID period (July 2021 – December 2024). The shift is robust across daily and weekly return frequencies. At the tangency portfolio, the Sharpe ratio improved by 6–7% despite higher absolute volatility, because expected returns increased by +5 percentage points against a volatility increase of only +3 percentage points. The shift was accompanied by a complete sectoral rotation in the optimal tangency portfolio: from financials and consumer staples to pharma, telecom, energy, and auto. Five structural forces—monetary policy accommodation, sectoral recomposition, correlation decline, China+1 FII inflows, and a retail participation surge jointly explain the shift, operating through both the mean vector and the covariance matrix.

8.2 Investment Implications

The results carry concrete implications for portfolio construction:

- **Pre-COVID strategies are stale.** A defensive portfolio tilted toward financials and FMCG (optimal by 2015–2019 Sharpe criteria) substantially underperformed the post-COVID frontier. Investors who did not rebalance their sector exposures left significant risk-adjusted performance on the table.

- **Active sector allocation matters.** The widened cross-sectional return dispersion increases the payoff to strategic sector tilts. Passive market-cap weighting is suboptimal when sectors are diverging structurally.

- **The new regime may persist.** The structural drivers:PLI schemes, China+1 dynamics, digital transformation,and growing retail participation show no signs of reversing in the near term. Investors should calibrate to the post-COVID regime rather than mean-revert to pre-COVID exposures.

8.3 Limitations

Several limitations temper the conclusions:

Survivorship bias. Our universe is fixed at 45 NIFTY 50 stocks as of January 2015. This excludes firms that failed, were delisted, or were added to the index after 2015, biasing both frontiers upward (excluded firms likely underperformed). The bias is larger for the pre-COVID period, since more time elapsed and more failures occurred.

Long-only constraint. We impose no-short-selling throughout. The true attainable frontier is wider when short selling is permitted; the observed shift may differ in magnitude for unconstrained investors.

Transaction costs and turnover. The tangency portfolio weights changed dramatically between periods. The Sharpe improvements are gross; after accounting for bid-ask spreads, market impact, and taxes on realized gains, the net improvement would be smaller.

Correlation versus causation. The five structural drivers we identify are correlational. Establishing causal attribution requires formal regression analysis (event studies, difference-in-differences, or structural VAR models) that is beyond the scope of this study.

In-sample estimation. All results are in-sample. Whether the post-COVID frontier translates to superior out-of-sample performance requires rolling out-of-sample backtesting with realistic rebalancing constraints.

8.4 Future Work

Several extensions are natural and technically well-motivated:

1. **Jobson-Korkie test.** A formal statistical test of the null hypothesis that the pre- and post-COVID tangency portfolios have equal Sharpe ratios would provide a p-value for the observed improvement, accounting for sampling uncertainty.

2. **De Roon–Nijman spanning test.** This test would determine whether the post-COVID frontier is truly outside the pre-COVID feasible set, or whether the same assets reweighted can span it, distinguishing a frontier shift from a weight shift.

3. **Frontier shift decomposition.** Decomposing the northeast shift into the contribution of changes in $\boldsymbol{\mu}$ versus changes in $\boldsymbol{\Sigma}$ would isolate whether higher returns or improved diversification dominates.

4. **Bootstrap confidence bands.** Constructing 95% confidence intervals around both frontiers via bootstrap resampling would quantify whether the observed gap is statistically distinguishable from sampling noise.

5. **Alternative covariance estimators.** Comparing Ledoit-Wolf with POET, DCC-GARCH, and factor models would assess the sensitivity of frontier estimates to the choice of covariance estimator.

6. **Transaction cost modeling.** Estimating rebalancing costs for the tangency portfolio shift would assess whether the Sharpe improvement survives net of bid-ask spreads and turnover taxes.

References

Ang, A. and Bekaert, G. (2002). Regime switches in interest rates. *Journal of Business & Economic Statistics*, 20(2):163–182.

Ashraf, B. N. (2020). Stock markets' reaction to COVID-19: Cases or fatalities? *Research in International Business and Finance*, 54:101249.

Baker, S. R., Bloom, N., Davis, S. J., Kost, K., Sammon, M., and Viratyosin, T. (2020). The unprecedented stock market reaction to COVID-19. *The Review of Asset Pricing Studies*, 10(4):742–758.

Bekaert, G. and Harvey, C. R. (1997). Emerging equity market volatility. *Journal of Financial Economics*, 43(1):29–77.

Bekaert, G. and Hoerova, M. (2014). Emerging market local currency bond yields and foreign holdings in the post-Lehman period. *Journal of International Money and Finance*, 43:1–14.

Brandt, M. W. (2009). Portfolio choice problems. *Handbook of Financial Econometrics*, 1:269–336.

Campbell, R., Koedijk, K., and Kofman, P. (2002). Increased correlation in bear markets. *Financial Analysts Journal*, 58(1):87–94.

Chopra, V. K. and Ziemba, W. T. (1993). The effect of errors in means, variances, and covariances on optimal portfolio choice. *Journal of Portfolio Management*, 19(2):6–11.

DeMiguel, V., Garlappi, L., and Uppal, R. (2009). Optimal versus naive diversification: How inefficient is the $1/N$ portfolio strategy? *The Review of Financial Studies*, 22(5):1915–1953.

Harvey, C. R. (1995). Predictable risk and returns in emerging markets. *The Review of Financial Studies*, 8(3):773–816.

Jagannathan, R. and Ma, T. (2003). Risk reduction in large portfolios: Why imposing the wrong constraints helps. *The Journal of Finance*, 58(4):1651–1683.

Jain, R. (2013). Institutional and individual investor preferences for dividends and share repurchases. *Journal of Economics and Business*, 69:40–56.

Ledoit, O. and Wolf, M. (2004). A well-conditioned estimator for large-dimensional covariance matrices. *Journal of Multivariate Analysis*, 88(2):365–411.

Longin, F. and Solnik, B. (2001). Extreme correlation of international equity markets. *The Journal of Finance*, 56(2):649–676.

Markowitz, H. (1952). Portfolio selection. *The Journal of Finance*, 7(1):77–91.

Mishra, P. K., Mishra, S. K., et al. (2020). Does the Indian financial market nosedive because of the COVID-19 outbreak, in comparison to after demonetisation and the GST? *Emerging Markets Finance and Trade*, 56(10):2162–2180.

Narayan, P. K., Phan, D. H. B., and Liu, G. (2021). COVID-19 lockdowns, stimulus packages, travel bans, and stock returns. *Finance Research Letters*, 38:101732.

Ramelli, S. and Wagner, A. F. (2020). Feverish stock price reactions to COVID-19. *The Review of Corporate Finance Studies*, 9(3):622–655.

Sharpe, W. F. (1964). Capital asset prices: A theory of market equilibrium under conditions of risk. *The Journal of Finance*, 19(3):425–442.

Zaremba, A., Kizys, R., Aharon, D. Y., and Demir, E. (2020). Infected markets: Novel coronavirus, government interventions, and stock return volatility around the globe. *Finance Research Letters*, 35:101597.