



The Structural De-Risking and Financialization of Indian Capital Markets: A Quantitative Macro and Fundamental Valuation Framework

EXYST VENTURES | Institutional Equity Strategy | 14 October 2026

1. The Macro-Fiscal Transmission Mechanism & Capex Velocity

1.1 Understanding GDP growth and what it means for investors

India's latest GDP data suggest that both spending on everyday goods and services and investment in long-lived assets expanded in the opening quarter of FY27. Investment grew faster than consumption. For Exyst's investment framework, this provides a reason to investigate businesses exposed to capacity creation and infrastructure spending. It does not, by itself, establish how much any listed company's earnings will grow.

What the four growth numbers tell us

Q1 FY27 means April to June 2026. Year-on-year growth compares this period with April to June 2025, rather than with the immediately preceding quarter. The official release estimates real GDP growth at 7.8%, nominal GDP growth at 10.3%, real private final consumption growth at 7.1%, and real gross fixed capital formation growth at 11.9%. ^[S01]

Measure	Reported growth	Meaning in everyday language
Real GDP	7.8%	The economy produced more goods and services after removing the effect of price changes.
Nominal GDP	10.3%	The rupee value of output increased, reflecting both changes in output and changes in prices.
Private final consumption	7.1% real	Households and relevant nonprofit institutions consumed more goods and services.
Gross fixed capital formation	11.9% real	Investment in assets such as machinery, buildings and infrastructure increased. This includes public, private and household investment.

A simple way to understand real and nominal growth is to imagine last year's GDP indexed to 100. On a constant-price basis, it rises to 107.8; on a current-price basis, it rises to 110.3. The second number incorporates price movements as well as output growth. It should not be described as 10.3% more physical production.

Using the rounded headline rates, the implied increase in the GDP deflator is approximately 2.32%: 1.103 divided by 1.078, minus 1. This is a calculated economy-wide price measure, not the consumer inflation rate. Simply subtracting 7.8 from 10.3 gives a rough 2.5 percentage point gap, rather than the exact implied price change.

Data cut-offs are stated beside observations. Sourced observations, algebraic derivations and illustrative scenarios are distinguished throughout. Currency is INR unless indicated; ₹1 lakh crore = ₹1 trillion. Returns in valuation scenarios are nominal, before fees and taxes.



These GDP figures are official estimates of activity that has already occurred. They can be revised as more information arrives. They are different from forecasts of future GDP or forecasts of corporate profits. A 7.8% increase in real GDP does not mean that company revenues, earnings or share prices must rise by 7.8%.

Why the base year matters

The revised national accounts use 2022–23 as the base year. Updating the base year also incorporates changes in data sources and measurement methods. Consequently, the current quarter and its comparison quarter must come from the same revised series and, wherever possible, the same release. Comparing a new-series number with an old-series number can create a misleading growth rate. ^[501]

Think of this as measuring two objects with the same ruler. If the measurement system changes, the earlier observation must be put on the new basis before calculating how much growth occurred. A release vintage means the version of the data published on a particular date, including revisions to earlier periods.

How spending is added up without counting it twice

The expenditure approach asks who bought the economy's final goods and services. It separates private consumption, government consumption, investment and net exports, while showing any statistical discrepancy. The simplified accounting identity is:

$$\text{GDP} = \text{Private consumption} + \text{Government consumption} + \text{Gross capital formation} + \text{Exports} - \text{Imports} + \text{Statistical discrepancy}$$

Gross capital formation includes fixed investment, changes in inventories and acquisitions less disposals of valuables. Fixed investment creates or replaces assets used over time. It is not the same as buying shares or mutual funds. Government consumption covers the delivery of current public services; it does not include every item in the government budget.

For example, a newly constructed public road is fixed investment. The fact that the government pays for it does not mean it should also be added to government consumption. Counting the same road in both categories would exaggerate GDP. Transfer payments similarly do not directly represent new production; the recipient's subsequent spending may enter consumption or investment.

Imports require another adjustment. If an Indian business purchases an imported machine for ₹100 crore, that purchase enters investment, but the imported equipment also enters imports. Subtracting imports prevents foreign production from being counted as Indian production. Domestic installation, transport and related services can still add Indian value. The machine may also improve domestic productivity in future years; that future benefit is separate from its immediate accounting treatment.

Inventories are goods produced but not yet sold or used. Domestic production placed into stock can contribute to GDP even before a final customer buys it. This may reflect planned preparation for stronger demand, or it may reflect weak sales. Faster inventory accumulation can temporarily support GDP growth, but investors should check whether stock subsequently sells and cash is collected.

The statistical discrepancy is the reconciliation gap between GDP and the separately measured expenditure components. It arises because different sources and estimates do not match perfectly. It must remain visible; it should not be presented as household spending, investment demand or another identified growth driver.

Why investment can contribute as much as consumption

An expenditure category's contribution to GDP growth depends on two things: how large it was in the starting period and how quickly it grew. A large category growing moderately and a smaller category growing rapidly can therefore make similar contributions.



For a simple illustration, assume last year's real GDP was 100, with private consumption of 55 and fixed investment of 32. These are teaching assumptions, not verified Q1 expenditure weights. Applying the reported growth rates gives:

Component	Assumed prior-year share	Real growth	Illustrative contribution
Private consumption	55%	7.1%	$55\% \times 7.1\% = 3.905$ percentage points
Fixed investment	32%	11.9%	$32\% \times 11.9\% = 3.808$ percentage points
Combined contribution	87%	Not applicable	7.713 percentage points

Consumption adds 3.905 units to our starting GDP of 100, while fixed investment adds 3.808 units. Investment starts from a smaller base but grows faster, so its contribution is almost as large. This is why investment can matter greatly to the growth outcome even when consumption remains the larger part of the economy.

If we mechanically compare this assumed contribution of 7.713 points with headline growth of 7.8%, the balancing amount is 0.087 point, or about 0.09. This is only an arithmetic residual in the example. It does not verify the actual combined contribution of government consumption, net exports, inventories, valuables and discrepancies. Their individual contributions could also be substantial and offset one another.

The calculations retain three decimal places to avoid a rounding inconsistency: 3.91 plus 3.81 is 7.72, whereas the unrounded sum, 7.713, rounds to 7.71. The extra digits illustrate the arithmetic and do not imply greater precision in the official growth rates.

For an actual decomposition, use the official component levels for both comparable quarters, on a consistent real-price basis. With additive constant-price levels, a component's contribution equals its real change divided by prior-year real GDP. Imports enter negatively; changes in inventory investment and discrepancies should be calculated from their levels. Multiplying current nominal shares by real growth rates mixes different measurement bases and is not an exact decomposition.

What would confirm a wider capital expenditure cycle

Capital expenditure, or capex, means spending on productive assets. Rising fixed investment is consistent with stronger asset creation, but it does not identify who is investing or why. It can include public infrastructure, household housing, replacement machinery and new private industrial capacity.

Construction growth might come from public roads or residential buildings rather than new corporate factories. Machinery output might serve exports or replacement demand. These are economically useful activities, but they provide different revenue opportunities for listed companies. The interpretation should therefore be tested against evidence beyond the GDP headline.

Evidence to examine	What it helps establish
Project financing and actual drawdowns	Whether announced investment is receiving funding and expenditure is beginning.
Equipment orders and deliveries	Whether demand is turning into procurement and supplier revenue.



Evidence to examine	What it helps establish
Construction progress and commissioning	Whether projects are advancing and new assets are becoming operational.
Customer and end-use details	Whether spending represents public works, housing, replacement assets, exports or private capacity expansion.
Receivables and operating cash flow	Whether recognised revenue is being collected rather than absorbed by working capital.

How Exyst connect the macro data to equities

The investment interpretation is selective. Faster investment can create demand for equipment, materials, engineering and related services. Which businesses benefit depends on the projects being executed, domestic sourcing, their market position and the timing of revenue recognition. This is an investment hypothesis to test against company disclosures, rather than a direct conclusion from national accounts.

The analytical sequence is to identify relevant spending, verify company orders, track delivery and revenue recognition, assess profitability, and then examine cash collection. An order announcement alone does not complete this sequence. Execution delays, low-margin contracts or slow customer payments can interrupt it.

For example, assume a supplier's revenue increases from ₹1,000 crore to ₹1,100 crore. If its operating margin remains at 10%, operating profit rises from ₹100 crore to ₹110 crore. If the margin improves to 12%, operating profit reaches ₹132 crore, an increase of 32%. If the margin falls to 8%, operating profit declines to ₹88 crore despite higher sales. These are hypothetical company calculations, not earnings forecasts.

Net profit also depends on depreciation, interest and tax. A company undertaking its own capex may initially face higher financing costs and depreciation before the new assets generate adequate sales. Even growing accounting profit may produce weak cash flow if receivables and inventories absorb funds. Shareholder returns then depend on the price paid and the valuation the market assigns to those earnings.

For Exyst, the GDP release supports a constructive starting point for research into consumption and investment exposure. Portfolio conviction should strengthen when the same company demonstrates relevant demand, profitable execution, cash conversion and acceptable returns on the capital it employs. The macro data identify where to investigate; company evidence establishes whether the opportunity can create shareholder value.

1.2 Project phasing and the measurable capex acceleration

RBI Bulletin authors estimate comparable pipeline capex across bank/FI finance, external commercial borrowing and IPO channels at **₹3.2 lakh crore for FY27**, against **₹2.6 lakh crore for FY26**. The bank/FI channel alone moves from approximately ₹2.0 lakh crore to ₹2.3 lakh crore. These are envisaged phasing amounts for identified projects, not realised economy-wide private investment. The publication explicitly attributes its analysis to the authors rather than an official RBI policy position. [S02]

The all-channel increment is ₹0.6 lakh crore, or ₹60,000 crore; growth on the rounded base is 23.08%. The bank/FI component grows 15.00%, implying that non-bank/FI channels contribute to the faster aggregate expansion. Because the source amounts are rounded, these derived growth rates have greater displayed precision than the underlying observations. A capex "velocity" defined as planned spending growth is therefore +23.08%; it is not



monetary velocity, an observed rate of project completion, or an acceleration in the mathematical second derivative of investment.

Capex measurement	FY26	FY27	Derived change	Interpretation
Comparable pipeline across three financing channels [S02]	₹2.6 lakh crore	₹3.2 lakh crore	+₹0.6 lakh crore; +23.08%	Planned phasing; execution remains conditional
Comparable bank/FI pipeline [S02]	₹2.0 lakh crore	₹2.3 lakh crore	+15.00%	Covered financing subset
Illustrative equal 85% realisation in both years	₹2.21 lakh crore	₹2.72 lakh crore	+23.08%	Ratio unchanged by equal completion rates
Illustrative FY26 90% / FY27 75% realisation	₹2.34 lakh crore	₹2.40 lakh crore	+2.56%	Execution slippage nearly removes headline growth

Let planned phasing be P and execution realisation be a . Realised expenditure is aP . If FY26 achieves 90% of its phasing, FY27 must achieve 73.13% of its own plan merely to maintain the same absolute expenditure: $0.90 \times 2.6 / 3.2$. Procurement delays, land acquisition, grid connection, environmental clearance and sponsor equity availability can therefore dominate the observable investment outcome even when commitments rise. New sanctions during FY27 may add expenditure beyond the pre-existing pipeline; cancellation and rescheduling may subtract it. Comparing unlike pipeline vintages would bias the growth estimate.

For portfolio attribution, an engineering supplier's addressable revenue is the project expenditure multiplied by domestic procurement share, relevant equipment share, attainable market share and revenue-recognition timing. An increase in national project finance intentions does not establish an equal increase in any company's order book. Deposits and mobilisation advances, performance guarantees, retention receivables and milestone certification can separate order intake, revenue and operating cash flow by several quarters. The analytical sequence is sanction → financing close → drawdown → procurement → execution → billing → cash collection → commissioned utilisation.

1.3 Public investment multipliers and private crowding-in

Public investment can affect private investment through three separate channels: contemporaneous demand for construction and equipment; a reduction in operating bottlenecks; and changes in the financing environment. The first raises supplier output, the second can increase asset turnover and potential output, while the third can either support or displace private credit. Aggregate government borrowing may raise risk-free yields even while better infrastructure raises expected private project cash flows. The sign of crowding-in depends on the net effect on project net present value.

$$\Delta Y = \mu_G \Delta I_G + \mu_P \Delta I_P, \quad \Delta I_P = \kappa \Delta I_G$$

The composite multiplier is $\mu_G + \mu_P \kappa$ when κ measures induced incremental private investment and μ_P captures its output effect over the stated horizon. This expression is an accounting scenario, not an identified econometric estimate. If μ_G already includes private crowding-in, adding $\mu_P \kappa$ would double-count that transmission. An empirical estimate should specify the fiscal shock, horizon, price basis, treatment of imports and monetary-policy response, preferably through local projections with controls for the business cycle.



In an illustrative model, an additional ₹1 lakh crore of public investment, a direct multiplier of 1.5, induced private investment of 0.4 per public rupee and a private-investment multiplier of 1.2 produce ₹1.98 lakh crore of incremental output over the model horizon. If private crowding-in falls to 0.1 and the direct multiplier to 0.8, the same impulse produces ₹0.92 lakh crore. The difference measures execution and financing sensitivity. Neither parameter set is asserted as an observed Indian fiscal multiplier.

The supply-side channel is governed by the capital-output ratio and utilisation. An approximate incremental capital-output ratio is investment share divided by real output growth, under restrictive steady-state assumptions. A 32% investment share and 7% output growth imply an ICOR of 4.57. If the capital required per unit of output falls to 4.2 at an unchanged investment share, implied growth rises to 7.62%. This is comparative statics, not a growth forecast: depreciation, sector composition, gestation periods and capacity utilisation prevent a one-period interpretation.

Private project approval requires expected incremental ROIC to exceed WACC. For a non-financial company, ROIC can be decomposed into after-tax operating margin multiplied by invested-capital turnover. Reducing logistics time or inventory requirements improves turnover even without selling-price increases. Conversely, a new plant can initially dilute ROIC because invested capital rises before capacity utilisation. Under an illustrative 20% tax rate, a 12% operating margin and capital turnover of 1.5, ROIC is 14.4%. If turnover is only 1.0 during ramp-up, ROIC is 9.6%. At an assumed WACC of 11%, the mature configuration creates economic value while the ramp-up configuration initially does not.

1.4 Tax buoyancy and the limits of the formalisation inference

EY's FY27 April–July fiscal analysis reports **gross tax revenue growth of 11.4% YoY** and **corporate tax growth of 20.8% YoY**. These exact figures align with the supplied inputs, but the tax growth is a nominal collection measure and includes timing effects. It is evidence consistent with a wider taxable base and stronger assessed profits; it does not isolate formalisation causally. ^[503]

Tax buoyancy is percentage growth in tax receipts divided by percentage growth in nominal GDP over a comparable interval. Using an illustrative 10% nominal GDP growth denominator gives gross-tax buoyancy of 1.14 and corporate-tax buoyancy of 2.08. A quarterly GDP print cannot be mechanically substituted for a four-month denominator without acknowledging the mismatch. Corporate-tax growth can reflect profit concentration, loss carryforwards, advance-tax instalments, refund timing, enforcement, effective tax rates and base effects. Structural formalisation should be assessed using constant-policy tax productivity and a multi-year decomposition, not four months of gross receipts.

$$T = \tau \Pi b, \quad \Delta \ln T \approx \Delta \ln \tau + \Delta \ln \Pi + \Delta \ln b$$

Here τ is the effective collection rate, Π taxable profit and b a coverage/compliance factor. The decomposition is conceptual because the factors are not directly identified from aggregate receipts. A collection increase driven by b is potentially persistent; one driven by a temporarily high commodity profit pool is cyclical. Refund-adjusted receipts, cash-versus-accrual differences and devolution to states determine what portion ultimately becomes a central fiscal cushion. Gross tax receipts and net revenue available to the Centre must therefore remain separate variables.

If a comparable annual tax base were ₹40 lakh crore, growth at 11.4% rather than 10.0% would create an additional ₹0.56 lakh crore of gross revenue. This hypothetical amount is close to the ₹0.6 lakh crore increment in the rounded private pipeline, but no financing linkage follows: private projects are funded through their sponsors and financial channels, while incremental taxes can fund interest, transfers, subsidies, deficits or public investment. Treating all tax outperformance as capex capacity would overstate the transmission.



1.5 Debt anchoring, sovereign duration and the corporate discount rate

The FY27 FRBM statement sets central government debt at 55.6% of GDP, versus 56.1% in FY26 revised estimates, and a fiscal deficit at 4.3%, versus 4.4%. Its medium-term debt objective is $50 \pm 1\%$ by FY31, with the fiscal deficit functioning as the operational target. These ratios describe central government debt, not consolidated general-government debt. [S04]

$$d_t = \frac{1+i_t}{1+g_t} d_{t-1} + p_t + SFA_t$$

Debt ratio d evolves with the effective nominal interest rate i , nominal GDP growth g , the primary deficit ratio p and stock-flow adjustments SFA . At an illustrative starting ratio of 55.6%, $i = 7\%$ and $g = 10\%$, the existing debt contributes a 1.52 percentage-point decline through the denominator/interest-growth effect. Adding a primary deficit of 0.7% leaves an approximately 0.82-point decline before stock-flow adjustments. If nominal growth falls to 7%, the denominator advantage disappears and debt increases with the primary deficit. These are model calculations, not forecasts of the published glide path.

A credible debt path can reduce expected sovereign supply pressure and the term premium, lowering corporate WACC. The pass-through is not one-for-one: inflation risk, state borrowing, bank demand for government securities, global yields and currency pressure can offset fiscal improvement. Cost of equity is modelled as risk-free yield plus market beta times equity risk premium, with any additional country or liquidity adjustment explicitly defined to prevent double-counting. A 50-basis-point decline in the risk-free rate lowers cost of equity by 50 basis points only if beta and risk premia are held constant. If the equity risk premium simultaneously rises 100 basis points for a beta-one portfolio, the net discount-rate change is +50 basis points.

The macro-to-equity transmission should therefore be tested through realised capex, incremental ROIC, cash conversion and discount rates jointly. A project boom funded at rising real rates can increase revenue while reducing economic profit. Strong collections absorbed by subsidies during an energy shock can preserve the deficit while failing to lower bond yields. These are identifiable failure channels for the structural de-risking hypothesis.

2. Valuation Mechanics, Earnings Yields, and Forward Curve Analysis

2.1 Valuation observations and date-consistent arithmetic

The supplied October dataset, reproduced by IndexPE for 9 October 2026, reports Nifty 50 trailing P/E of **19.27**, a five-year median of **21.95**, P/B of **2.77**, dividend yield of **1.52%**, and five-year derived EPS growth of **12.34% annualised**. These are provider-reported observations and calculations, not an independently reconstructed daily NSE series. The official NSE factsheet for 30 September reports P/E 19.36, P/B 2.78 and dividend yield 1.22%. The differing date and dividend input remain explicit; the 1.52% yield is retained as the supplied valuation scenario rather than silently substituted into the official factsheet. [S05, S06]

At 19.27 times trailing earnings, the earnings yield is 5.189%. At 21.95 times, it is 4.556%. The current multiple is 12.21% below the supplied historical median; the earnings-yield increment is 63.36 basis points. Reversion from 19.27 to 21.95 would generate 13.91% price appreciation at constant EPS. The percentage discount and the percentage recovery are different because their denominators differ. A relative discount is evidence of a lower entry price per unit of trailing earnings; it does not establish absolute undervaluation or an unconditional buy zone.



$$EY = \frac{1}{PE}; \quad \text{Discount} = 1 - \frac{19.27}{21.95}; \quad \text{Reversion} = \frac{21.95}{19.27} - 1$$

At P/B 2.77, the book yield is 36.10%. If P/E and P/B use consistent aggregate numerator/denominator definitions, P/B divided by P/E implies earnings-to-book of 14.37%. This is a cross-sectional ROE proxy based on the relevant book denominator, not necessarily an average-book accounting ROE. At a 1.52% dividend yield, the implied payout ratio is 29.29%; the corresponding retained share is 70.71%. Multiplying that retained share by the ROE proxy gives a stylised sustainable growth rate of 10.16%. Using the official 1.22% dividend yield with the same scenario P/E for sensitivity gives a payout ratio of 23.51% and sustainable growth of 10.99%.

These linked ratios are useful consistency checks, but they are not three independent statistical confirmations. P/E and P/B share price in the numerator, while dividend yield uses price in the denominator. One price decline lowers the first two and raises the third even if earnings, book value and dividends do not change. Calling their alignment “rare” requires a joint historical distribution under consistent methodology and a defined percentile threshold. No frequency estimate is inferred from the snapshot alone.

Derived valuation measure	Supplied October inputs	Calculation	Interpretation
Trailing earnings yield	P/E 19.27	5.189%	Earnings/price; not distributable cash yield
Historical relative discount	Median P/E 21.95	12.21%	Relative to the supplied median
Constant-EPS median reversion	19.27 → 21.95	+13.91%	Conditional price effect
Earnings-to-book proxy	P/B 2.77 / P/E 19.27	14.37%	Requires consistent aggregation

2.2 Earnings growth, multiple compression and the price identity

For the supplied five-year earnings CAGR, the EPS growth factor is 1.1234^5 , approximately 1.789. An unchanged multiple would produce the same price growth factor, before dividends. To analyse realised consolidation, the starting multiple and price index must also be observed. The provider reports a starting P/E of 27.40 and ending P/E of 19.27 for this window. Their ratio is 0.7033, representing 29.67% cumulative multiple compression. [S05]

$$\frac{P_T}{P_0} = \frac{E_T M_T}{E_0 M_0}; \quad \Delta \ln P = \Delta \ln E + \Delta \ln M$$

Combining the reported earnings growth with those multiples implies price CAGR of approximately 4.70%. The annualised multiplicative multiple effect is –6.80%; adding simple annual percentage rates would be incorrect because the components compound. The exact relationship is $(1 + \text{price CAGR}) = (1 + \text{EPS CAGR}) \times (\text{terminal multiple} / \text{initial multiple})^{(1/5)}$. This demonstrates that the five-year provider-defined window can contain earnings expansion and weak price growth simultaneously.

It does not prove that the latest correction contains no earnings degradation. A positive five-year endpoint CAGR can coexist with negative revisions during the last two quarters, margin pressure in a particular sector, or deteriorating cash conversion. Nor is index EPS a same-company aggregate: reconstitution, free-float changes and corporate actions affect the index's economic exposure. The provider derives EPS from index level divided



by published P/E; that series is not a constituent-by-constituent growth audit. The investment inference is narrower: historical price underperformance relative to the reported EPS series is algebraically compatible with multiple compression.

For a one-period attribution, a 5% price fall and 3% earnings decline imply a multiple change of $0.95 / 0.97 - 1 = -2.06\%$. A 5% price fall with unchanged earnings implies -5% multiple compression. A 5% price fall with 5% earnings growth implies -9.52% multiple compression. The same price move can therefore represent different fundamental outcomes. A research dashboard should track trailing earnings, next-year earnings revisions, normalised margins and cash earnings separately.

2.3 Forward valuation curve and return sensitivity

Trailing P/E is not forward P/E. If earnings grow at g for h years while price remains fixed, forward P/E equals $19.27 / (1 + g)^h$. At the historical 12.34% rate used solely as a scenario, one-, two- and three-year forward multiples are approximately 17.15, 15.27 and 13.60. No analyst-consensus forecast is embedded in those calculations. At slower EPS growth, the apparent forward cheapness decays more slowly; with an earnings decline, forward P/E increases despite an unchanged index level.

The holding-period return surface should therefore specify earnings growth and terminal multiple jointly. For a price-only return, the exact five-year CAGR is $(1 + g) \times (\text{terminal P/E} / 19.27)^{1/5} - 1$. A stylised total-return calculation assumes dividends grow at the same constant g as earnings and are reinvested annually at the same annual price growth rate. Under those assumptions, the gross annual wealth factor is the sum of the annual price factor and starting dividend yield multiplied by $(1 + g)$. This is a transparent scenario construction, not an additive forecast or reproduction of the historical TRI study.


$$R_{P,5} = (1 + g) \left(\frac{M_5}{19.27} \right)^{1/5} - 1$$

Illustrative five-year scenario	EPS CAGR	Terminal P/E	Price CAGR	Total-return CAGR using 1.52% starting yield
Earnings stagnation / de-rating	0.00%	16.00	-3.65%	-2.13%
Slower earnings / de-rating	6.00%	16.00	+2.13%	+3.74%
Moderate earnings / unchanged multiple	9.00%	19.27	+9.00%	+10.66%
Historical EPS rate / unchanged multiple	12.34%	19.27	+12.34%	+14.05%
Historical EPS rate / median multiple	12.34%	21.95	+15.30%	+17.01%

The median-reversion case is upside sensitivity, not the base required to justify allocation. At a 12.34% EPS rate, a terminal P/E of 16 still produces approximately 8.24% price CAGR because earnings offset de-rating over five years. In contrast, a one-year 10% EPS contraction combined with a terminal multiple of 16 produces a price loss of 25.27%. Long-horizon compounding and short-horizon drawdown are distinct risk dimensions.



Five-year price CAGR surface



Scenario calculations from starting P/E 19.27. Dividends, fees and taxes excluded.

Figure 1. Earnings growth and terminal multiple jointly determine price returns.

The cost-of-equity hurdle determines whether a return surface is economically attractive. A portfolio requiring 12% nominal annual return cannot interpret the slower-growth scenario as adequate merely because its terminal return is positive. Discounting cash distributions and terminal value at the investor's hurdle is equivalent to testing the scenario's internal rate of return against that hurdle, after allowing for fees, tax and implementation. Foreign investors must additionally translate local-currency returns through the exchange-rate path.

2.4 Earnings yields, residual income and justified multiples

Comparing the 5.189% trailing earnings yield with an assumed 7% government-bond yield gives a -1.811 percentage-point raw yield gap. That difference is not the equity risk premium: earnings are risky, retained earnings can compound, bond coupons are contractual, and the maturity and tax bases differ. A meaningful comparison requires expected distributable cash flows, growth duration and an explicit discount rate. A negative raw yield gap can coexist with attractive equities if reinvestment earns sufficient returns, or with expensive equities if reported earnings overstate durable cash generation.

A constant-growth dividend model gives a justified trailing P/E of $\text{payout} \times (1 + g) / (\text{cost of equity} - g)$, provided cost of equity exceeds perpetual growth. Growth must be consistent with retained earnings and sustainable ROE. At the scenario-implied payout of 29.29%, earnings/book of 14.37% and sustainable growth of 10.16%, a P/E of 19.27 corresponds to a stylised cost of equity of approximately 11.84%. This is reverse engineering of a restrictive model, not an observed market discount rate.

$$PB_{\text{justified}} = \frac{ROE - g}{k_e - g}; \quad g = b ROE$$

The residual-income relationship shows why P/B below a historical median is not automatically value. Book value creates economic value when sustainable ROE exceeds cost of equity. A bank carrying impaired loans or a manufacturer capitalising costs can report book value that overstates recoverable equity. Asset quality,



provisioning, intangibles, acquisition goodwill, minority interests and dilution should be normalised before linking P/B to expected returns. A portfolio with higher book multiples can still be economically cheaper if incremental ROE and cash conversion are materially superior.

The steady-state discount rate must not be applied to a temporary high-growth phase indefinitely. A two-stage framework discounts explicit five-year cash flows, then switches to a terminal growth rate compatible with long-run nominal output and the firm's competitive position. Sensitivity becomes extreme when cost of equity approaches terminal growth. At an assumed 12% cost of equity and 5% terminal growth, raising the discount rate to 13% reduces the terminal denominator-based value by 12.5% before other effects. That is a discount-rate shock, not an earnings shock.

2.5 Historical conditional returns and statistical identification

IndexPE's original published study reports that the **17.3–20.2 P/E band** has a **13.17% median five-year forward annualised total return** and a displayed **100% positive historical frequency** for daily overlapping windows starting since 1999. The reported worst five-year CAGR in that band is +3.68%. The same study explicitly states that windows overlap and are not independent. It also reports +13.20% median five-year return for the most expensive quintile, so its table does not establish a strictly monotonic five-year valuation-return relationship. [S07]

The historical statistic is a conditional sample frequency, not a 100% probability of future profit. Only starts with a full five-year future return can enter a completed-window analysis: with a 9 October 2026 cut-off, starts after approximately 9 October 2021 are right-censored. A dataset described as extending to 2026 may therefore contain valuation observations that cannot contribute to the five-year outcome sample. The published matrix does not disclose enough raw observations here to independently reproduce the sample or estimate confidence intervals.

$$\hat{p} = \frac{\sum_t \mathbf{1}\{17.3 \leq PE_t \leq 20.2\} \mathbf{1}\{R_{t,t+5} > 0\}}{\sum_t \mathbf{1}\{17.3 \leq PE_t \leq 20.2\}}$$

The denominator must exclude censored and missing outcomes. Adjacent daily five-year windows share almost all their observations, so treating several thousand starts as independent Bernoulli trials radically overstates precision. An approximate effective sample size can be based on outcome autocorrelation, but binary conditional outcomes and regime clustering require care. For illustration only, five independent successful observations would give a one-sided 95% binomial lower bound of $0.05^{1/5} = 54.93\%$; twenty independent successes would give 86.09%. Neither calculation estimates this study's actual probability because its independent sample size is not established.

NSE changed index P/E to consolidated earnings where available from 31 March 2021, replacing the preceding standalone basis; its dividend-yield methodology also changed to a rolling ex-dividend basis. [S08] A valuation history spanning 1999 therefore crosses a measurement regime. An institutional replication must either reconstruct a comparable earnings basis or estimate separate regimes. Other controls include contemporaneous constituents, point-in-time data, fixed versus full-sample percentile boundaries, sector weights, starting inflation, real rates and returns in both INR and USD. Ex-post quintiles can introduce information unavailable to the investor at the historical entry date.

Evidence supporting a valuation-conditioned allocation framework is consequently stronger than evidence supporting a deterministic entry rule. Validation should report median, lower decile, expected shortfall, maximum interim drawdown, recovery time and frequency of failing the investor's hurdle. It should use block



bootstrap or non-overlapping cohorts, and out-of-sample tests after parameter selection. No target index level or formal buy recommendation is implied by the supplied band.

3. Micro-Structure Alterations & The Financialization of Savings

3.1 Asset stocks, recurring flows and the investable liquidity denominator

The mutual fund industry has exceeded the supplied ₹72 lakh crore scale threshold: the AMFI–Crisil Factbook reports ₹73.73 lakh crore at March 2026 against ₹31.43 lakh crore at March 2021, an 18.60% five-year CAGR. It reports SIP assets of ₹14.83 lakh crore versus ₹4.25 lakh crore, with SIP assets representing 20.1% of industry assets versus 13.5%. AMFI reports August 2026 gross SIP collections of ₹32,297 crore. These observations have different cut-offs and must not be combined into a single same-date balance sheet. [S09, S10]

Total industry AUM is a stock containing equity, debt, hybrid and other schemes; it is not cash available to buy Indian equities. SIP collections are gross subscriptions through an investment method, not net equity subscriptions and not identical to DII cash-market purchases. Redemptions from existing holdings can exceed fresh instalments. Mutual fund equity purchases can also arise from non-SIP subscriptions, internal cash deployment or rebalancing. Insurance and pension institutions form additional domestic channels whose purchases cannot be inferred from AMFI totals alone.

The correct flow identity separates contributions from returns. AUM at the end of the period equals beginning assets plus net subscriptions plus investment gains and other adjustments. In return-accounting terms, valuation gains on beginning assets and on intra-period flows must be handled at their respective dates. Counting mark-to-market AUM growth as new liquidity exaggerates financialisation precisely when market prices rise fastest.

$$A_t = A_{t-1} + \text{Subscriptions}_t - \text{Redemptions}_t + \text{Investment gains}_t + \text{Adjustments}_t$$

Annualising August's gross SIP rate gives ₹3.876 lakh crore, assuming an unchanged monthly collection level. That is a run-rate calculation, not a forecast, and it does not measure net equity demand. On an assumed equity-eligible share of 80%, the monthly allocation capacity before redemptions would be ₹25,838 crore. If only 60% of that amount remains after associated withdrawals and cash retention, available net demand falls to ₹15,503 crore. The eligibility and retention parameters are illustrative and must be calibrated from scheme-level data.

The SIP asset stock grows by approximately 3.49 times over the Factbook's five years, implying a derived CAGR near 28.4%. This combines subscriptions and investment performance, and it does not mean new contributions grew at that rate. SIP's industry-asset share rises 6.6 percentage points. To measure the structural savings transition, households' net acquisition of financial assets must be compared with income, physical savings and liabilities; a larger mutual fund stock alone is insufficient to infer a higher household saving rate.

3.2 Price impact, foreign flow absorption and state-dependent beta

Price formation depends on net aggressive order flow relative to available risk-bearing capacity. Define q_F as signed foreign demand, q_D as domestic institutional demand, q_H as other household/non-institutional demand and q_S as new security supply. An illustrative impact equation scales flows by a relevant float or liquidity denominator and includes a common global-risk factor. New issues absorb cash without necessarily appearing in secondary-market flow statistics; turnover is a gross two-sided transaction measure, not net capital supply.

$$r_t = \lambda_t(q_{F,t} + q_{D,t} + q_{H,t} - q_{S,t}) + \gamma_t Z_t + \varepsilon_t$$




Here λ captures impact per unit of normalised net demand, z a global factor and ϵ residual return. Domestic absorption can be expressed as $qD = q0 - \phi qF + u$, where $q0$ is autonomous domestic demand and ϕ measures countercyclical response. Substitution makes sensitivity to foreign flow $\lambda(1 - \phi)$. At $\phi = 0.5$, the foreign-flow impact is half its value under zero domestic reaction, holding λ constant. This parameterisation explains a mechanism; it is not an empirical estimate of India's present beta.

If foreign flows depend on a global factor as $qF = az + v$, return loading on that factor is $\gamma + \lambda(1 - \phi)a$, assuming u is independent of z . When domestic demand also falls during global stress, that independence fails. If risk-bearing capacity deteriorates and λ doubles, an absorption coefficient of 0.5 merely offsets the liquidity deterioration; impact need not decline. Stable monthly SIP contributions can improve $q0$ while leaving the acute foreign-flow coefficient or global fundamental loading substantially unchanged.

A stylised monthly stress sets foreign selling at ₹60,000 crore and gross SIP inflow at the reported August rate. At 80% equity eligibility, gross SIP capacity covers 43.06% of the selling. If 20% of eligible contributions is lost to a retention/cash adjustment, capacity falls to ₹20,670 crore, leaving ₹39,330 crore of selling to be absorbed elsewhere. Adding an assumed ₹15,000 crore of other domestic demand reduces the imbalance to ₹24,330 crore. If issuance simultaneously absorbs ₹10,000 crore, the residual imbalance increases to ₹34,330 crore. These figures describe flow accounting, not predicted index losses.

Illustrative monthly flow absorption



August SIP observation with assumed 80% eligibility and 80% retention. Other flows assumed.

Figure 2. Gross recurring subscriptions are only one component of net demand.

The distinction matters because stronger domestic participation may reduce the price impact of a given flow shock while foreign investors still price the marginal risk premium. Currency depreciation, global discount-rate increases and earnings revisions affect all holders regardless of nationality. Domestic buyers can substitute ownership without reversing the change in discounted fundamental value. Capital-market financialisation lowers a particular dependence on foreign financing only to the extent that domestic savings are persistent and available at stress prices.



3.3 Liquidity persistence and the endogenous redemption channel

Recurring instalments smooth contribution timing and reduce the need for households to choose an entry day. Persistence remains behavioural and income-dependent. Unemployment, lower disposable income, leverage, emergency spending and disappointing returns can cause cancellations or redemptions. Account registrations, active contributing accounts, stopped mandates and unique investors measure different populations; none should be substituted for another. The relevant stress variable is rupees of net deployable equity inflow, not a headline count of mandates.

A linearised feedback model writes domestic demand as autonomous flow plus br , with b positive for procyclical flows and negative for contrarian deployment. Solving the impact equation gives a denominator $1 - \lambda b$. Positive b amplifies shocks as that denominator declines. Negative b dampens them. A market supported by steady instalments can nevertheless suffer amplification if redemptions from its accumulated stock become strongly procyclical. The size of that stock makes even a small redemption percentage economically material.

$$r = \frac{\lambda(q_F + q_0 - q_S) + \varepsilon}{1 - \lambda b}$$

For example, assume an equity/hybrid redeemable stock of ₹40 lakh crore solely for stress modelling. A 1% monthly net redemption generates ₹40,000 crore of withdrawal demand, exceeding the illustrated equity-eligible gross SIP amount of ₹25,838 crore. If the net withdrawal is 2%, the demand doubles. This stock is an assumed model input, not a measured category total. The exercise identifies the denominator risk: contribution persistence should be tested against a larger accumulated liability to redeeming investors.

Open-ended funds face a mismatch when shareholders can redeem promptly but underlying securities cannot be liquidated quickly at quoted prices. The effect is greatest where ownership concentration is high and executable depth is low. Managers may initially sell liquid large caps to meet redemptions, transferring stress into the index even if the original performance shock originated in small caps. Alternatively, selling illiquid holdings can increase impact and recovery time. Portfolio construction must consider redemption coverage, cash buffers and stress liquidation costs at security level.

3.4 Market beta, drawdown floors and empirical testing

Market beta is covariance of local returns with the chosen global benchmark divided by global return variance. It is benchmark-, currency- and horizon-dependent. A lower INR beta need not imply a lower USD beta when currency losses covary with global stress. Domestic flows can reduce return variance attributable to foreign order imbalance while the covariance generated by energy, rates or earnings shocks remains high. Therefore, a claim that financialisation has “fundamentally altered beta” requires measured pre/post estimates and uncertainty, not an AUM level alone.

An empirical specification should include global equity returns, changes in US yields, USD strength, oil returns, foreign and domestic flows scaled by lagged free float, valuation and domestic macro surprises. Interactions with a stress indicator identify whether domestic absorption persists in the lower return tail. Flow and return are jointly determined: contemporaneous DII buying can be a reaction to lower prices. Causal interpretation requires instruments, institutional timing or carefully lagged variables, rather than a regression coefficient alone.

Rolling windows should be accompanied by confidence intervals and regime tests. A structural break attributed to domestic flows may instead reflect changing index sector weights, a lower oil price regime, different monetary policy or a shorter observation sample. A comparison of selected crises confounds shock magnitude with market



structure. Counterfactual drawdowns should condition on comparable global shocks, domestic leverage and initial valuations.

Maximum drawdown is the largest peak-to-trough loss on a defined total-return wealth path; it cannot be bounded mechanically by monthly SIP volume. A smaller negative-tail impact coefficient would support a higher conditional price floor in a calibrated model, but no guaranteed floor is justified. Stronger participation can also sustain richer starting multiples, increasing valuation duration. If earnings disappoint, that additional duration can offset liquidity stabilisation. Financialisation changes who bears risk and how flows transmit; it does not eliminate fundamental repricing.

3.5 Portfolio implementation under domestic flow concentration

The liquidity denominator for a position is executable volume at the proposed participation rate, adjusted for crowding, spreads and adverse impact. An illustrative ₹500 crore position with ₹100 crore ADTV and a 10% participation limit requires 50 normal-volume trading days to liquidate. If volume halves under stress, it requires 100 days before adjusting for further price impact. The position's inclusion in an index or ownership by domestic institutions does not establish stress liquidity.

A square-root impact model can be used for sensitivity, with impact proportional to volatility times the square root of traded position divided by market volume. Its scale coefficient must be calibrated to the security and venue. Splitting a parent order lowers daily participation but exposes the portfolio to execution risk across more days. Risk budgeting should combine beta, single-name fundamentals, ownership overlap, liquidation horizon and mandate redemption needs.

Domestic flow support is most directly relevant to secondary-market absorption; it is less direct evidence for a company's growth or governance. The fundamental underwriting variables remain incremental ROIC, diluted EPS growth, working-capital intensity and free-cash-flow conversion. Portfolio-level exposure to financialisation can be expressed through asset managers, distribution platforms and market infrastructure, but their earnings sensitivity also depends on fee compression, product mix, regulation and market levels. A growing industry stock does not guarantee operating leverage for every intermediary.

4. Asymmetric Risk Matrix & Tail Risk Quantification

4.1 AI capital reallocation and benchmark-relative earnings exposure

MSCI India's 30 September 2026 factsheet assigns Financials 30.23%, Consumer Discretionary 12.80% and Consumer Staples 5.20%, a combined **48.23%**, validating the approximate 50% concentration premise. Information Technology represents 6.75%, but that category is not equivalent to pure-play semiconductor or AI hardware exposure. A dedicated constituent-level classification is necessary to measure AI revenue, capex and profit sensitivity. ^[S11]

Relative performance against an AI-heavy benchmark can weaken even if domestic earnings remain positive. The relevant decomposition separates industry allocation, stock selection within industries, currency and valuation change. If an assumed global comparator holds 30% in AI hardware and India holds 2% under a bespoke exposure definition, a 20% hardware outperformance contributes approximately –5.6 percentage points of allocation effect, before interaction terms. Both weights are hypothetical; this is a factor scenario, not measured MSCI attribution.



Foreign reallocation is a portfolio-demand channel. If global investors raise AI weights by reducing other emerging-market allocations, India's local multiple can compress without a domestic earnings contraction. The domestic economy can simultaneously benefit from digital adoption while the local index lacks ownership of the dominant global hardware profit pool. Distinguishing technology consumption from rent capture is essential: electricity demand, data-centre construction and software productivity gains do not imply that the listed Indian index earns the semiconductor scarcity margin.

The countervailing risk is AI capex retrenchment. Lower global hardware earnings could impair technology benchmarks while easing their competition for portfolio capital; it could also lower enterprise technology budgets and threaten Indian services revenue. Portfolio hedges using global AI securities can therefore diversify one relative-performance risk while introducing valuation, concentration and currency risks of their own. A correlated global risk-off shock can reduce both markets, so the hedge must be evaluated across regimes rather than through a single growth narrative.

4.2 Crude dependency, the balance-of-payments identity and policy burden sharing

The supplied **~80% crude import dependence** understates the more recent official consumption-basis measure. A Ministry of Petroleum and Natural Gas parliamentary reply dated 30 March 2026 places it broadly at **~88%** during FY24–FY26, and cites current oil demand around 5.5 million barrels per day. [S12] The 80% premise is retained as a lower-dependence sensitivity; the 88% case governs the higher-exposure illustration. Import dependence is not itself the share of GDP spent on oil, and consumption-basis dependence is not identical to refinery throughput dependence.

Assume total oil demand of 5.5 million barrels/day, import-linked exposure equal to the stated dependence ratio, no initial volume response, and a uniform delivered-price shock. The mechanical gross import-bill increment at 88% exposure is $5.5 \times 0.88 \times 365 \times \text{price change}$, in millions of dollars. This approximation does not reproduce the customs crude import series because product trade, refinery yields, stock changes and price discounts require separate treatment.

$$\Delta B_{\text{oil}} \approx Q d 365 \Delta P$$


Illustrative sustained oil shock	Gross bill increase at 80% dependence	Gross bill increase at 88% dependence	Share of assumed USD 4.5tn GDP at 88%
+USD 10/barrel for one year	USD 16.06bn	USD 17.67bn	0.393%
+USD 20/barrel for one year	USD 32.12bn	USD 35.33bn	0.785%
+USD 40/barrel for one year	USD 64.24bn	USD 70.66bn	1.570%
+USD 20/barrel for one quarter	USD 8.03bn	USD 8.83bn	0.196% of annual GDP

The GDP denominator of USD 4.5 trillion is a model assumption. These amounts are gross external expenditure sensitivities, not estimated current-account deterioration. Product export receipts, volume substitution, remittances, service exports and other trade flows can offset the impact. Higher refining margins can support some exporters while domestic transport, aviation, chemicals and consumer spending face higher costs. A



geopolitical disruption can also affect freight, insurance and physical availability, making a price-only model incomplete.

Oil price shock exposure at fixed demand



Assumed demand 5.5mb/d; no volume response. Gross expenditure sensitivity, not a CAD forecast.

Figure 3. Gross oil expenditure sensitivity rises with import dependence.

The burden is allocated across households, firms and the public sector through pass-through decisions. Full retail pass-through protects fiscal revenue and distributor margins but raises CPI and reduces real disposable income. Excise reductions or subsidies limit household transmission but weaken fiscal space. Delayed pass-through can compress marketing margins or generate receivables and contingent fiscal claims. Domestic consumption is a growth buffer against export weakness; it is not a hedge against a terms-of-trade loss because domestic households still purchase energy directly and through goods prices.

To quantify rather than assume pass-through, an illustrative coefficient of 0.015–0.035 percentage point of CPI inflation per USD 1/barrel sustained price increase yields 0.30–0.70 point for a USD 20 shock. An illustrative growth coefficient of –0.01 to –0.02 percentage point per USD 1 gives a 0.20–0.40-point growth drag. These coefficients are analyst-selected stress ranges, not RBI estimates. They should be re-estimated under current fuel taxes, exchange rates and administered prices before entering an investment forecast.

If oil import spending is assumed at 3% of GDP and prices increase 25%, the first-order gross external burden is 0.75% of GDP at unchanged quantities. Absorbing 40% fiscally would imply a 0.30%-of-GDP burden, subject to the actual tax/subsidy mechanism. This can erode the tax cushion described in Section 1. Nonlinearities arise if currency depreciation compounds dollar-priced energy costs or if fiscal support itself raises borrowing needs and sovereign yields.

4.3 Currency translation, inflation and cross-asset correlation

For an unhedged USD investor, local equity return must be divided by the increase in INR per USD. A 12% INR return with 5% INR depreciation yields 6.67% in USD, before fees and withholding. A 20% local equity decline



with 8% depreciation yields –25.93%. Currency and equity losses compound, so subtracting percentage changes understates the adverse result.

$$1 + R_{USD} = \frac{1 + R_{INR}}{1 + \delta_{INR/USD}}$$

Hedging does not remove risk costlessly. Forward points reflect interest-rate differentials and market basis, while collateral and rollover requirements create liquidity demands. A hedge sized to starting NAV can become excessive after an equity drawdown unless rebalanced. For companies, the relevant exposure is net foreign-currency cash flow and debt service, not exports alone. Import-intensive exporters can have smaller natural hedges than gross revenue labels imply; translation gains may differ from operating cash gains.

Equity/bond diversification can fail in an oil-inflation shock. Consider an assumed 60% equity/40% bond portfolio with annualised volatilities of 20% and 6%. At equity/bond correlation –0.2, implied portfolio volatility is 11.76%; at +0.5 it becomes 13.36%. These are covariance calculations under hypothetical inputs, not observed Indian correlations. The rise occurs because both asset classes can be exposed to higher inflation and yields. A nominal bond allocation should therefore be stress-tested by duration and inflation regime, rather than treated as invariant protection.

Gold or foreign assets can provide different shock exposures, but their hedge effectiveness depends on real yields, dollar strength, local currency and liquidity. A hedge should be selected against a named risk channel and evaluated by conditional covariance and carry. Buying an asset merely because it has historically risen in one crisis does not quantify its hedge ratio or opportunity cost.

4.4 Monetary transmission lags, repricing asymmetry and banking earnings

RBI policy changes affect overnight funding, deposit pricing, floating-rate assets, fixed-rate assets and credit demand on different schedules. An illustrative distributed-lag model writes the change in loan yield as a weighted sum of present and past policy-rate changes; the corresponding deposit response has different weights. Cumulative pass-through and the timing of that pass-through are separate parameters. No uniform banking-system lag is assumed from a single policy announcement.

$$\Delta y_{L,t} = \sum_{j=0}^J a_j \Delta r_{t-j}; \quad \Delta c_{D,t} = \sum_{j=0}^J b_j \Delta r_{t-j}$$

For a simplified bank with deposits equal to 90% of earning assets, a 50-basis-point loan-yield decline and a 20-basis-point deposit-cost decline imply approximately 32 basis points of NIM compression: $-50 + 0.9 \times 20 = -32$. If average earning assets are assumed at ₹10 lakh crore, that change reduces annualised net interest income by ₹3,200 crore before balance-sheet growth and other repricing. For an assumed ₹20,000 crore pre-tax profit base, it is a 16% adverse sensitivity, before tax, fee income and credit costs.

If deposit costs later decline by a further 30 basis points while asset yields stop falling, NIM recovers by 27 basis points under the same simplified funding ratio. This demonstrates a lagged J-curve in earnings without requiring a deterioration in loan demand. During tightening, deposit costs can rise faster than asset yields where competitive funding pressure dominates; during easing, asset yields can reset faster than term deposits. The direction depends on liability mix, reset dates and hedging.

A second banking risk is credit deterioration after an income shock. Rate cuts can reduce borrower debt service while contemporaneous NIM compression reduces lender earnings. New loans can initially improve growth metrics even as older cohorts develop stress. Analysis should use vintage delinquency, unsecured concentration,




stage migration, collateral coverage and provision adequacy. A lower policy rate is not a sufficient condition for higher bank EPS, particularly in an index with substantial financial-sector exposure.

A simple sensitivity test on the sovereign curve should also accompany equity duration. At an assumed equity cash-flow duration of 12, a 100-basis-point increase in discount rate implies an approximate 12% valuation decline before convexity and earnings effects. The duration is a model parameter, not a Nifty estimate. Combining that loss with an earnings downgrade must use a multiplicative or full DCF calculation so that the same discount-rate effect is not counted again through an independently selected P/E shock.

4.5 Joint shocks, expected shortfall and reverse stress testing

The scenario matrix below applies the exact price identity to the supplied starting P/E. EPS shocks, terminal multiples and currency changes are assumptions. They are not probabilities, target levels, VaR estimates or forecasts. Dividend receipts are excluded from the one-year price-loss calculation to isolate the valuation channel.

One-year conditional stress	EPS change	Terminal P/E	INR price return	INR depreciation	USD price return
Foreign reallocation / stable earnings	0%	18.00	-6.59%	3%	-9.31%
Oil and rate pressure	-8%	17.00	-18.84%	6%	-23.43%
Joint earnings and valuation shock	-15%	15.00	-33.83%	10%	-39.85%
Severe liquidity and earnings shock	-25%	13.00	-49.40%	15%	-56.00%



Deterministic stress assumptions. No probabilities assigned; dividends excluded.

Figure 4. Currency depreciation magnifies USD equity losses.



The severe case requires roughly a 97.6% subsequent INR price gain to recover the original price level, before dividends. It can occur in a scenario without contradicting a positive median five-year historical return. Interim funding needs, collateral calls or forced redemptions can make that path unacceptable even if eventual recovery is possible. Horizon length should be matched to liabilities and liquidity tolerance rather than invoked as a universal risk mitigant.

For loss variable L , VaR at confidence α is the α -quantile; expected shortfall is the average loss in the corresponding tail, with appropriate treatment of atoms in a discrete distribution. At an assumed zero-mean Gaussian annual return with 20% volatility, one-year 95% VaR is 32.90% and expected shortfall is 41.25%. These are distributional illustrations. Normality omits jumps, volatility clustering and changing correlations, and the four deterministic scenarios above do not constitute a probability distribution from which market VaR can be inferred.

$$\text{VaR}_\alpha = F_L^{-1}(\alpha), \quad \text{ES}_\alpha = \frac{1}{1-\alpha} \int_\alpha^1 \text{VaR}_u du$$

An institutional implementation should combine historical simulation, filtered volatility, block resampling and explicit event overlays. Backtests should distinguish exception frequency from exception severity, and recognise that a regime transition can invalidate the calibration sample. Scenario losses can enter a probability-weighted framework only after probabilities are justified independently; arbitrary weights would create false precision. Maximum drawdown, liquidity-adjusted loss and expected shortfall answer different questions and should all remain visible.

Reverse stress testing starts with the investor's loss tolerance. A 25% price loss with a 10% EPS contraction corresponds to terminal P/E of approximately 16.06: $19.27 \times 0.75 / 0.90$. With stable earnings, the same loss requires 14.45. If an institution has a 12%-of-NAV budget for the -33.83% equity stress and assumes all other assets unchanged, the maximum equity weight is approximately 35.47%. That calculation is a risk-budget illustration, not a proposed client allocation; losses in other assets, leverage and hedging costs would reduce the allowable weight.

4.6 Risk controls and falsification thresholds

Risk mitigation should be connected to the mechanism being underwritten. Capex exposure is controlled through project completion and cash collection; energy exposure through pricing power, input hedges and balance-sheet capacity; financial exposure through liability repricing and credit quality; foreign-flow exposure through liquid implementation and currency management. A valuation discount cannot substitute for governance diligence, and a domestic flow cushion cannot substitute for a liquidity reserve.

The following thresholds are proposed monitoring rules for this framework, not observed breaches or regulatory requirements. Consecutive misses should trigger a documented reassessment, with sector-specific calibration rather than automatic execution.

Risk channel	Proposed monitoring trigger	Quantitative review	Portfolio response to assess
Capex execution	FY27 pipeline realisation below 73.13%, assuming FY26 at 90%	Realised spend would fail to exceed the illustrative FY26 base	Reassess suppliers' revenue conversion and working capital



Risk channel	Proposed monitoring trigger	Quantitative review	Portfolio response to assess
Earnings revision	Next-year EPS estimates fall >10% over two quarters	Rebuild forward P/E and hurdle-return surface	Reduce earnings-duration exposure where thesis deteriorates
Domestic liquidity	Net equity subscriptions turn negative for three months	Separate SIP instalments from stock redemptions and issuance	Raise redemption coverage; review ownership crowding
Energy stress	Oil remains USD 20 above the chosen baseline for two quarters	Re-estimate pass-through, trade balance and subsidies	Test input-intensive margins and inflation-sensitive duration
Bank repricing	Modelled NIM decline exceeds 30bp without deposit catch-up	Asset/liability reset ladder and credit-cost sensitivity	Reassess financial-sector contribution to aggregate EPS
Currency	INR depreciation exceeds 5% alongside equity losses	USD return, net corporate FX exposure and hedge liquidity	Adjust hedge sizing within mandate constraints
Trading capacity	Stress liquidation horizon exceeds liability horizon	Participation limits, stressed volume and impact	Resize illiquid exposures before forced selling

The structural thesis is falsified at the portfolio level when discounted incremental cash flows fail to compensate for the cost of equity and stress losses exceed the mandate's capital or liquidity budget. Operational evidence should include realised project expenditure, post-commissioning returns, diluted earnings and distributable cash. Monitoring macro growth without those bridges leaves the model unable to distinguish value creation from capital intensity.



Source and calculation notes

S01. Ministry of Statistics and Programme Implementation / PIB, Q1 FY27 GDP release, 31 August 2026, including revised 2022–23-base national accounts. Primary statistical release; figures translated from the official Hindi text. <https://www.pib.gov.in/PressReleaseDetailm.aspx?PRID=2305232&lang=2®=48> Supplementary explanations: PIB, India's GDP Performance, 1 September 2026, <https://www.pib.gov.in/FactsheetDetails.aspx?Id=150991&lang=1®=3> ; MoSPI, Understanding the New Series of GDP, February 2026, https://www.mospi.gov.in/uploads/announcements/announcements_1772117257791_84ae898f-7be2-4b7d-a135-565e1a809513_FAQ_GDP_26022026_1902.pdf ; BEA, The Expenditures Approach to Measuring GDP, 3 June 2025, <https://www.bea.gov/news/blog/2025-06-03/expenditures-approach-measuring-gdp> (general accounting explanation; national presentation conventions differ).

S02. RBI Bulletin authors, Private Corporate Investment Growth in 2025–26 and Outlook for 2026–27, 25 September 2026. Primary author study; comparable envisaged pipeline, not realised total private capex. https://rbi.org.in/scripts/BS_ViewBulletin.aspx?Id=24448

S03. EY India Economic Pulse, FY27 April–July fiscal analysis, accessed 11 October 2026. Published analytical source for the exact 11.4% and 20.8% collection growth inputs; the underlying CGA line-item series was not independently reconstructed for this paper. https://www.ey.com/en_in/services/tax/india-economic-pulse

S04. Government of India, FY27 FRBM statements, central government fiscal and debt targets, pages 9–13 of the retrieved PDF. Primary budget document. <https://www.indiabudget.gov.in/doc/frbm1.pdf>

S05. IndexPE Nifty 50 valuation and derived earnings analysis, 9 October 2026. Primary publisher of its calculations; republished daily market inputs. Exact NSE daily observations and five-year series were not independently reconstructed. <https://indexpe.in/nifty-50>

S06. NSE Indices, Nifty 50 factsheet, 30 September 2026, page 1. Official dated cross-check: 19.36 P/E, 2.78 P/B, 1.22% dividend yield. https://www.niftyindices.com/Factsheet/ind_nifty50.pdf

S07. IndexPE, P/E versus forward returns, as of 9 October 2026. Primary publisher of the historical analysis; TRI returns and overlapping daily windows. Raw sample counts and independent replication are unavailable in this paper. <https://indexpe.in/pe-vs-forward-returns>

S08. NSE Indices, methodology change press release, 23 February 2021, effective 31 March 2021; consolidated earnings and rolling dividend definitions. Primary methodology document. https://www.niftyindices.com/Press_Release/ind_prs23022021_1.pdf

S09. AMFI–Crissil Factbook 2026, March 2021–March 2026 observations, retrieved PDF pages 4, 15 and 37. Its specified SIP definition excludes funds of funds investing overseas. Primary industry study. https://www.amfiindia.com/uploads/AMFI_Factbook2026_14b34a6447.pdf

S10. AMFI SIP industry page, August 2026 collections, accessed 11 October 2026. Primary collection observation. This paper does not label the August observation as October flow data. <https://www.amfiindia.com/mutual-fund>

S11. MSCI India USD factsheet, 30 September 2026, page 2. Primary index composition source; sector classifications do not quantify pure-play AI hardware. <https://www.msci.com/documents/10199/72e88ba4-e198-4e3f-9d43-e95958476403>

S12. Ministry of Petroleum and Natural Gas, Rajya Sabha Unstarred Question 4124, answered 30 March 2026. Primary official source for approximate 88% consumption-basis import dependence and cited oil demand. https://sansad.in/getFile/annex/270/AU4124_gs8kNF.pdf?source=pqars

All sensitivity coefficients, execution rates, model portfolio weights, terminal multiples, discount rates and stress paths identified as illustrative are analytical assumptions. Calculations use rounded supplied inputs and may differ slightly from unrounded source series. The five-year total-return scenarios assume constant EPS/dividend growth, smooth annual multiple adjustment, end-year dividends and reinvestment at the scenario price-growth rate. They exclude expenses, taxes, transaction costs and deviations from those assumptions. Source methodology and cut-off differences remain material.

Compliance footnote: This document constitutes institutional-grade equity strategy research for sophisticated investors and does not represent formal SEBI-registered advisory services.

Prepared by Exyst Ventures for private client circulation. This research is not a personalised suitability assessment, an offer or an instruction to transact. Market investments involve loss of capital; historical frequencies and scenario returns do not guarantee future outcomes. Source information can change and requires validation before execution. The compliance wording does not itself determine regulatory classification or Exyst Ventures' registration status. No legal, tax or registration status is asserted by this paper.

Exyst Venture Private Wealth Desk | 4-Bha-5, Jawahar Nagar, Jaipur – 302004, India | +91 830675 6660 | exystventures@gmail.com



Factor Comparison Matrix

Factor	Current Statistical Reality	5-Year Structural Trend	Institutional Portfolio Implication
Private capex	FY27 pipeline ₹3.2 lakh crore vs FY26 ₹2.6 lakh crore; +23.08% [S02]	Two-year intentions; five-year execution trend unestablished	Test completion, incremental ROIC and cash conversion
Fiscal capacity	Apr–Jul 2026 gross tax +11.4%; corporate tax +20.8% [S03]	Short interval cannot isolate structural formalisation	Normalise refunds, timing and cyclical profits
Debt anchor	FY27 central debt 55.6% of GDP; FY31 objective $50 \pm 1\%$ [S04]	Prospective glide path, not realised history	Model growth-interest spread and term premium
Index valuation	Supplied October P/E 19.27; median 21.95; discount 12.21% [S05]	Provider five-year multiple $27.40 \rightarrow 19.27$	Require adequate returns without re-rating
Earnings	Provider-derived five-year EPS CAGR +12.34% [S05]	Endpoint growth can mask recent downgrades	Audit revisions, dilution and cash earnings
Book and yield	Supplied P/B 2.77; yield 1.52%; official Sep yield 1.22% [S05, S06]	Joint historical rarity not established	Link P/B to ROE; reconcile dividend basis
Conditional returns	P/E 17.3–20.2: five-year median 13.17%; displayed historical positive frequency 100% [S07]	Completed overlapping windows; 2021 methodology break [S08]	Sample evidence, not forward certainty; test tails
MF asset stock	March 2026 AUM ₹73.73 lakh crore [S09]	₹31.43 $\rightarrow$ ₹73.73 lakh crore; CAGR 18.60%	Separate subscriptions from market gains
SIP flows	August 2026 gross collections ₹32,297 crore [S10]	SIP assets ₹4.25 $\rightarrow$ ₹14.83 lakh crore, 2021–2026 [S09]	Stress net inflows against redemptions
Foreign absorption	No calibrated absorption coefficient or beta estimate	Broader ownership; regime effect requires testing	Measure impact and stress trading depth
AI allocation	Financials plus Consumers 48.23% of MSCI India [S11]	Sector snapshot; five-year AI exposure unquantified	Separate domestic growth from AI rent capture
Oil dependence	Supplied ~80%; official recent measure ~88% [S12]	Import vulnerability remains material	Stress oil, FX, subsidies and margins jointly
Banking lag	Modelled –50bp asset / –20bp deposit repricing $\rightarrow$ –32bp NIM	Recurrent mechanism; no uniform lag measured	Underwrite reset ladders and credit cohorts
Joint tails	Model: EPS –15%, P/E 15 $\rightarrow$ INR –33.83%; 10% INR depreciation $\rightarrow$ USD –39.85%	Structural change cannot bound drawdowns	Size to loss, collateral and liquidation budgets