Why Firms Fail in India

Part 2: The Cycle Lens

Maneesh Dangi

Mosaic Asset Management

April 2026

A sequel to 'Why Firms Fail: Insights from My Research with IIM Udaipur's CFR' (Mosaic, March 2025).

1. Why a Part 2

I am writing this in April 2026, with the Strait of Hormuz intermittently closed, oil and LNG routing around the chokepoint at war-risk premiums, and the Indian credit cycle looking benign at the headline level. Part 2 is conceived in this tension — the macro background is hostile, the credit tape is calm, and investors I speak with are reasonably asking what to do about it. Part 1 (read it on our website) gave you the micro. I laid out a twelve-vector framework for assessing firm-level risk, walked through Deepanshu's Impact-Evidence-Cascading evolution on DHFL, and named the three patterns that recur across most Indian bankruptcies — mismanagement (very broad term, I know), competition, and governance. That is the ex-ante toolkit. It tells you how to look at one firm on one day.

But firms in India do not fail uniformly across time. They fail in waves. 2015–18 was a wave. 2018–20 was a wave. 1997–2001 was a wave. Between these, the default rate runs below 1% — benign enough that credit looks easy. Inside these, it crosses 10% — and nothing about the twelve vectors, applied mechanically, would have told you to hide. The framework is necessary; it is not sufficient. You need a second lens.

Part 2 is that second lens. The claim is simple. Firms fail because something hits them from outside — a shock. The shocks come in four flavours: demand, supply, policy, and finance. Leverage and promoter integrity decide whether a firm that gets hit survives or dies. Financiers' behaviour — forbearance, evergreening, eventual recognition — decides whether defaults arrive smoothly or in a wall. The cycle rhymes roughly every ten years. Part 1's HERO factor is not random across firms; it clusters by cycle. The 2012–15 HERO was commodities and policy. The 2018–20 HERO was ALM and governance. The HERO of the current cycle is forming now, and the back half of this note argues what it looks like.

A note on intellectual honesty before we begin. There is no Shumway-style causal hazard model on comprehensive Indian firm data. The CRISIL default study gives us thirty-six years of actuarial rates. Subramanian's 2016–17 Economic Survey gave us the Twin Balance Sheet diagnosis. Chari, Jain and Kulkarni (2019) and Kulkarni, Ritadhi, Vij and Waldock (2022) gave us hard evidence of evergreening. The rest is practitioner judgment and case-study pattern recognition. Where this note is numerical, it is sourced. Where it is inferential, I say so. Maybe purely speculative. Don't read this as academic literature. Just one where a practitioner is recording his views.

2. The Four Shocks

Every large default I have seen in India — and many I have only read about — traces back to one or more of four external shocks. Firms rarely fail on a calm day. Something moves the ground first.

Demand shock

Someone, somewhere, sells what you sell for less. In India, that someone is overwhelmingly China. Textiles, specialty chemicals, steel, auto components, solar cells, toys, LEDs, and now auto parts for

EVs — the list that China impacts is very long. Of course, the disruptor isn't always foreign — if you were an Indian telecom operator over the past decade, the answer was very local: Jio. A demand shock is not necessarily a fall in volumes. Prices collapse first, margins go before volumes do, and by the time the volume line on the chart turns down, the balance sheet is already gone. I watched this in domestic steel 2014–16. I am watching it now in specialty chemicals and in small-ticket textiles. Auto-components and renewables may be next. The tell is always the same: operating margin compresses for three or four consecutive quarters, and management blames 'transient destocking.' It is usually a lot more than that. You may want to call this a price shock, instead of a demand shock.

Supply shock

An input you priced at X becomes 2X, and you cannot pass it through. As I write this in April 2026, the Strait of Hormuz is intermittently closed — the IRGC re-asserted control on 18 April following the US blockade of Iranian ports — and the roughly fifth of global seaborne oil and fifth of global LNG that transit the strait is routing around it at war-risk premiums several times their pre-crisis levels. The IEA's March 2026 Oil Market Report calls this the largest supply shock in the history of the global oil market. Brent moved from the low-sixties to well above a hundred dollars a barrel over the quarter. Fertilizer, aluminium and helium are also disrupted. Indian refiners, gas-based power, petrochem, and fertilizer urea producers are directly exposed. This is the live supply shock of the current cycle and we do not yet know its shape. Supply shocks are deceptively sector-specific, which means they do not show up in aggregate default rates for a year — but within the affected sector, the wipeout is brutal. Check coverage ratio sensitivity to your largest input cost. If a twenty-to-twenty-five percent move in coal or crude or palm oil takes interest coverage below 1.5x, the firm does not actually have a margin of safety — it has a margin of luck. You may think it reads more like a cost shock instead of a supply shock.

Policy shock

The Indian state giveth, the Indian state taketh away — and it does so without warning. The 2G spectrum cancellation (Supreme Court, 2012, 122 licences struck down) effectively ended the first telecom incumbent group. The AGR judgment (Supreme Court, October 2019) upheld DoT's expansive definition of adjusted gross revenue and crystallised roughly ₹1.47 lakh crore of principal dues across fifteen telcos, with interest and penalties stacked on top. The shock concentrated on the surviving private players — Airtel, Vi, and, before it exited, Tata Tele — and for Vi the interest-and-penalty layer alone turned the definition dispute into an existential solvency event. The 2014 coal block cancellation voided nearly all of 218 allocations. The Karnataka iron ore mining ban, the Goa mining suspension, right-of-way denials on road projects, environmental-clearance withdrawals, RBI's November 2018 action against a non-bank, the sudden cancellation of a limestone lease — each one is an idiosyncratic event; collectively they form a pattern. As I wrote in Part 1, be sceptical of sectors that need government licences, and be doubly sceptical of receivables where the payer is a state government. It is not that the money does not come. It comes late, and it comes late enough to kill a levered balance sheet. I once had exposure to an education company setting up computer labs for state governments — good business, real demand, but the states would not simply pay. Receivables ballooned, default

followed. Eventually money came in, but after a lot of struggle. The lesson is that solvency eventually is not the same question as solvency in the interim.

Finance shock

Heard of sudden stops? Capital that was available yesterday at 9% is unavailable at any price today. 2008's global sudden stop hit Indian corporates who had borrowed offshore in ECBs and FCCBs, but the Indian default wave that should have followed was muted — RBI's aggressive liquidity response, the fiscal stimulus, and the rapid 2009–10 recovery let most firms refinance through. The size of a shock is not what determines the wave; the scope of the response is. Which is why the 2018 shock — smaller in absolute terms, narrower in scope, confined to NBFCs and the mutual fund holdings that funded them — produced a larger Indian wave relative to its size. The origin was local, no global cavalry was coming, and the policy response had to come from inside the system that was itself under stress. I then had argued that RBI should have done a little more. But you know how we are as fund managers. We always want system put for every risk we take!

IL&FS's default in September 2018 froze the NBFC commercial-paper market within two weeks; DHFL followed in 2019. Both failed through the same trigger — a funding squeeze on short-tenor CP liabilities — but the underlying cause in each was different. IL&FS had weak assets: bulky long-tenor infrastructure lending whose recoverability was suspect even before the squeeze. DHFL had fraud, or was it? A substantial portion of the loan book was later revealed to be fictitious, and the funding squeeze forced the disclosure of what was already rotten. The common lesson is not that ALM mismatch kills NBFCs — it does, but only when combined with something already wrong on the asset side. Clean ALM with weak or fraudulent assets is what actually blew up. Clean ALM with clean assets survived 2018. This is also why Part 1's Financial Cycle vector scored so highly on the DHFL table, and why it will score highly on every NBFC table I ever build.

In practice, shocks do not arrive one at a time. 2012–15 was demand (global commodity collapse) plus policy (2G, coal blocks, iron ore). 2018–20 was finance (IL&FS-led funding shock) plus demand (real estate absorption). Most firms that died had two shocks against them, not one. A single shock is survivable for a clean balance sheet. Two, stacked, rarely are.

3. The Four Amplifiers

A shock is the trigger; an amplifier is what turns a bruise into a fracture. The four that follow determine whether a shocked firm survives or dies, and unlike the shocks themselves, all four are endogenous — which is another way of saying they are underwritable. The move here is deliberately from macro to micro. Part 1 dwelt on these at length; what follows is the short version, because no discussion of how firms fail is complete without them. So a little bit of repeat here.

Leverage

A firm that is 2x Net Debt / EBITDA absorbs a 30% EBITDA hit and emerges levered but solvent. The same firm at 5x absorbs the same hit and is in default territory. The shock is identical. The outcome is not. This is why the 2003–08 capex cycle — financed at peak commodity assumptions and peak growth assumptions — produced the 2012–18 default wave. The leverage was put on at the top. The shock, when it came, produced expected results.

One refinement on Part 1's framework. Firm leverage is the headline number; promoter leverage is the hidden one, and in India it is frequently the more dangerous of the two. When a promoter has pledged his equity stake to fund share buying, an unrelated venture, or simply a lifestyle, the firm's own balance sheet understates group stress. Promoter leverage migrates to firm leverage, always, and usually late. The mechanisms are known: a side letter to the promoter's personal lender that the firm is quietly on the hook for, a buyback executed at a price that transfers cash upward, a shortfall undertaking to a group company that was never disclosed to the board. A firm at 2x Net Debt / EBITDA whose promoter is levered at the holdco may already be a 5x firm in distress — you just haven't found the mechanism yet.

Which brings me to a heresy. Many fund managers insist on lending only to operating companies, where the cash flows are visible, and refuse to touch loans against promoter shares. The empirical record is murkier than that preference sounds: failure modes for the two are substantially the same, and promoter loans at least carry a risk premium for the sin. It is not that promoter lending is safe. It is that operating-company lending in India is not as safe as the decision rule assumes.

Promoter integrity

Part 1 argued that mismanagement is the single most prevalent cause of Indian failures — more prevalent than the business cycle slowdown story that dominates developed-market bankruptcies. Part 2 pushes further: promoter integrity is not a sub-clause under leverage. It is a co-equal amplifier, and on any honest reading of Indian credit history, it is the more decisive of the two.

The Indian evidence is thinner than I'd like, but it points the same way. Chahal and Ahmad (2022) show that politically-connected Indian firms exhibit materially higher investment inefficiency than comparable peers, and argue this is part of what produced the Twin Balance Sheet problem in the first place. Their 2026 follow-up, using a two-decade panel of S&P BSE 500 firms with a novel donor dataset, finds that firms donating to the BJP and the Congress obtain more long-term debt, primarily through government-owned banks. Aggarwal (2023) documents that politically-connected Indian firms see their cost of debt fall around state elections, and only from PSU lenders — consistent with a quid pro quo. So, key message: political connection distorts access to credit and degrades the quality of the resulting book.

I carry a simpler practitioner rule. Lending to a politically powerful promoter is a non-control trade. Whether the loan is repaid sits almost entirely in the borrower's discretion — the value of collateral, covenants and personal guarantees all fall sharply when the counterparty has the "political strength"

cover to make enforcement slow, public and expensive. That is the rule I use on the desk. The research is what I quote to people who want to override it.

The governance question extends well beyond PEPs. Consider the JSPL-versus-Bhushan contrast I drew in Part 1: both steel firms defaulted into the same down-cycle, but lenders lost almost nothing on JSPL and lost most of their Bhushan exposure. The cycle was identical. What separated the two was cash leakage, related-party transactions, and the willingness of one promoter to treat the firm's liabilities as someone else's problem.

Basic message is, India's loss-given-default of roughly 70% — against an OECD benchmark closer to 30% — is not primarily a story about IBC efficiency. It is a governance story.

Three practitioner tells, in rough order of importance. First, family splits. When a business family divides, at least one faction inherits the weaker half of the group and is forced to make its first independent capital allocation under pressure. The weaker faction goes looking for a big-bang acquisition to prove itself, and finds a seller happy to oblige. ADAG is just one such example; it's actually all over the place. Second, personal guarantees and post-dated cheques posted late in the cycle — PDCs from the promoter's personal account being a specifically Indian late-cycle signal, where the promoter is pledging personal liquidity to keep the firm's facilities live. A promoter who still believes he can pull it off will stake personal collateral, and it looks like integrity. More often it is the last stage of denial — which is why Part 1's CCD and Kingfisher observations matter. Third, unrelated diversification. A promoter in a stable cash-generative business who announces an adjacency in aviation, infrastructure, or financial services is telling you one of two things: the original business has run out of room, or he has mistaken domain competence for general competence. Both end in the same place.

None of these tells requires forensic accounting. They require sitting with the promoter for three hours, talking to two ex-employees, and reading the last five years of related-party notes to accounts. The quantitative tools from Part 1 — Beneish M-Score above −2.22, Altman Z declining through 1.81 — are the confirmation. But to be sure — the conversation is the real detection. I have not yet met a promoter who can lie consistently through multiple meetings. The answer is always to keep asking questions, patiently, again and again.

Broken unit economics

The firm was never going to work at scale. Contribution margin after fully-loaded variable cost is negative or near zero; growth makes the hole larger, not smaller. Byju's is a composite example — broken unit economics in the core tuition business, compounded by debt-funded acquisitions at the top of the cycle and governance breakdowns — but the unit-economics piece is the one that made scale compound the problem rather than solve it. The pattern runs through a generation of consumer-tech, edtech, D2C and challenger-FMCG firms that raised equity against revenue growth rather than unit profit. Unit economics is an amplifier, not a standalone failure mode — the firm survives while equity keeps funding the hole and defaults the moment equity markets close.

Venture debt is a runway trade sitting on top of a broken-unit-economics bet. The borrower is a loss-making firm with equity investors, a cash balance, and a finite runway to the next round. Debt extends the runway by six to twelve months in exchange for a warrant and a coupon. The credit thesis is not ‘this firm will generate EBITDA to service the debt.’ It is ‘this firm will raise the next equity round before the debt amortises, and the warrants will deliver the real return.’ Call it what it is. The two questions that matter at underwriting are therefore narrow ones: will unit economics improve from here, and will the named equity investors be willing and able to lead the next round.

Working capital

The classic Indian mid-market killer, and the one most consistently under-appreciated by credit analysts trained on EBITDA and leverage ratios. A firm with 60-day receivables and 120-day inventory, funded by 90-day payables and a working capital line, is 3–4 supplier decisions away from distress. The distress signals are mundane: debtor days creeping from 75 to 95 to 120, channel stuffing at quarter-end (revenue booked, cash never arrives), inventory write-offs blamed on ‘shade variation’ or ‘obsolescence.’ The book looks profitable until the working capital line is not renewed. Then it is not. Working capital is an amplifier because it converts a small shock — a single delayed payment from a large customer, a bank tightening its WC line at cycle turn — into default. Check debtor days, inventory days, and payable days against both the firm’s own three-year history and its closest peer. Divergence here precedes a default by nine to eighteen months. These exact timelines are made up. But you get the point!

For NBFCs, working capital is the whole game — it just shows up under a different name: asset-liability management. Rule out a serious ALM mismatch and rule out obvious fraud, and the probability of an NBFC going down drops sharply. The sector is more regulated, and therefore more transparent, than almost any other borrower category a credit fund encounters — but it is also structurally more levered than any operating business, and its failure modes do not map cleanly onto the four amplifiers above.

The practitioner rule I carry: never finance a wholesale NBFC without ALM comfort, and treat liability concentration as the single most important metric, ahead of GNPA. Wholesale NBFCs — those lending large tickets to developers, infra SPVs, or corporates — combine bulky assets and bulky liabilities, and the mismatch is structural. Retail NBFCs with diversified liability stacks (term loans across twelve banks, retail NCDs, securitised pools, small-ticket CP) can survive what wholesale NBFCs cannot.

4. The Mechanism — Why Defaults Come in Waves

We just took a tour of four multipliers. Now back to how a shock travels through different firms. A shock hits firms unevenly and continuously. Defaults, in contrast, arrive in walls — 2015–18, 2018–20 — with long and calm stretches in between. If firm distress is the underlying water level, what holds it back until the dam bursts? The answer is lender behaviour. Understanding this mechanism is the single

most practically useful idea in this note, because it tells you when the next wave is due — not whether, but when.

The sequence is always the same. Phase one: a firm suffers a shock and misses a payment. The financier restructures — extends tenor, cuts coupon, adds a moratorium — because recognising the bad loan means a provision, a hit to capital, and often a career event for the lending officer. Phase two: the restructured loan is serviced out of fresh debt from the same or a syndicate bank. The firm is now a zombie — economically insolvent, accounting-solvent, consuming credit that could have funded a viable firm. Phase three: forbearance becomes policy. Between 2008 and 2015, the RBI's regulatory forbearance allowed a whole generation of infra, power and metals exposures to be carried at par while operationally the underlying cash flows were nowhere near debt service. Phase four: something forces recognition. In 2015, it was the Asset Quality Review. In 2018, it was the IL&FS-led funding shock. Once recognition fires, the wall materialises within 6–12 months.

This is not speculation. Academic work supports it. Chari, Jain and Kulkarni (2019) show that between 2008 and 2016, stressed Indian banks systematically extended larger loans to low-solvency and zombie-credit firms even as those firms' capex fell — evergreening identified at a firm-bank pair level. Kulkarni, Ritadhi, Vij and Waldock (2019) document the zombie firm population created by RBI's forbearance regime and show that the eventual losses were larger precisely because the dead firms consumed credit that should have been cut off years earlier. Acharya, both in his academic work on European zombie lending and in his speeches as RBI Deputy Governor on Indian NPAs, drew the parallel to Japan's lost decade. The mechanism travels.

Two operational implications for investors in my funds or other credit funds. First, default rates lie during the forbearance phase. The period 2010–14 shows a deceptively low default rate in the CRISIL series; the underlying stress that exploded in 2015–18 was already there, hidden in restructured-standard assets. If you underwrote in 2013 on a benign default-rate chart, you bought into a wave that was already loaded. Second, the timing of the wall is set by the forbearance exit, not the shock. The 2012 commodity collapse did not produce a default wave in 2012. It produced one in 2015, when the AQR forced recognition. The 2018 IL&FS-led funding shock did not produce the full real estate wave in 2018; it produced it in 2019–20, as successive rounds of refinancing failed.

How long the phase lasts depends on what forces the exit. A regulatory intervention (the 2015 AQR) produced a seven-year gap between the 2008 shock and recognition; a single-firm failure that freezes a funding channel (2018 IL&FS) produced a two-year gap. Your job is to identify, at any given moment, what would force recognition in the current cycle — because that timing, not the underlying distress, determines when the next wall arrives.

5. Four Down Cycles

The argument of Part 2 is that Indian corporate defaults rhyme. Here are the four cycles of the last three decades, each described by its triggering shock, its HERO factor (in Part 1's sense), its concentrated sectors, and what the CRISIL long-dated series tells us about magnitude.

1997–2001: Asian contagion and domestic overcapacity

The 1997–2001 cycle is harder to document than what came later — the Indian corporate-default literature really starts post-2008 — so what follows is reconstructed from contemporaneous RBI reports and practitioner accounts. The shape is clear; the firm-level detail isn't. The shock was twofold — the 1997 Asian crisis (capital flight, rupee pressure, dollar funding costs doubling for ECB borrowers) and the tail of the mid-1990s domestic capex boom in cement, steel, and textiles that had built capacity ahead of demand. HERO factor: overcapacity. Commodity prices collapsed; Indian steel producers who had built at peak assumptions could not service. The CDR mechanism was formalised in 2001 in direct response. Sectoral concentration: cement, steel, textiles. Regulatory response that shaped the next cycle: SARFAESI (2002), DRTs strengthened, and the rupee began its managed float. Learning of the cycle — never fund exuberance, either at macro level or at sector level. Rapid ramp-ups in capacity almost always lead to major sector credit events. Today, I worry about the renewable sector given that lens.

2008–2009: The global sudden stop

A short, sharp finance shock rather than a wave. Lehman's collapse froze global credit for six months; Indian corporates who had stacked ECBs and FCCBs in the 2006–08 euphoria faced simultaneous refinancing stress. HERO factor: rollover risk. The interesting feature of this cycle is that the Indian default wave was muted — RBI's aggressive liquidity response, the fiscal stimulus, and the rapid 2009–10 recovery let most firms refinance their way through. This is the cycle where the foundations of the next wave were laid: the RBI's regulatory forbearance regime started here, and the corporate debt overhang that was not cleaned in 2009 festered until the 2015 AQR forced recognition. In CRISIL's series, 2008–09 registers as an uptick, not a spike — which is itself the lesson. In hindsight, 2008–09 and 2012–18 are best read as one cycle with a deferred wave, bridged by RBI's forbearance regime — a live demonstration of §4's mechanism. As a lender, I learnt a few lessons on ALM. 2018 would re-teach those lessons again. Onlookers may ignore how catastrophic a major asset-liability mismatch tends to be. You have to live through a crisis to see how a benign mismatch in it turns into a delinquency event for a firm.

2012–2018: The Twin Balance Sheet wave

The largest one I saw, experienced and learnt the most lessons from. This is the cycle Subramanian recognised in the 2016–17 Economic Survey as the Twin Balance Sheet problem — overlevered corporates on one side, undercapitalised public-sector banks on the other, locked together in a mutually-reinforcing stress equilibrium. The shocks were stacked: a demand shock (the 2014–15 global commodity collapse), a policy shock (2G cancellation 2012, coal block cancellation 2014, iron ore

mining bans in Karnataka and Goa, land and environmental clearance delays), and eventually a finance shock as the 2015 AQR forced recognition. HERO factors varied by firm: commodities for JSPL and Bhushan, policy for the telecoms and power-sector projects, governance for Bhushan and ADAG.

Sectoral concentration was textbook: the five sectors Part 1 cited — steel, power, telecom, infrastructure, textiles — accounted for the majority of defaulted exposure. IBC was enacted in 2016 directly in response. The first large IBC resolution — Tata Steel's acquisition of Bhushan Steel for ₹35,200 crore (a 37% haircut on ₹56,079 crore of admitted financial creditor claims) — closed in 2018 and set the template for the dozen large cases that followed. Loss-given-default across this wave averaged around 70%, and the post-IBC numbers have not meaningfully improved that figure. This is the cycle that taught me the lesson I carry into every deal: in India, loss-given-default on a stressed exposure will exceed 50% even when the borrower was well-managed going in. What determines recovery is not the cover ratio but the tangibility of the collateral, the ease of enforcing on it, and — above all — the speed of enforcement. My respect for loan-against-shares structures went up sharply through this cycle: I watched lenders holding a residual charge on pledged equity recover near-fully, while senior secured lenders on the same borrower's operating assets struggled to get anything out. Of course, Op-Co lenders will never agree to this!

2018–2020: The NBFC crisis

Triggered by a finance shock — IL&FS's default in September 2018 froze the NBFC CP market within two weeks — and propagated through the asset-side weakness and governance failures already present in the wholesale NBFC book. HERO factor: ALM mismatch amplified by governance. DHFL followed in 2019. The wholesale-NBFC business model was exposed: bulky short-tenor liabilities against bulky long-tenor real estate and infra assets. Second-order casualties were the real estate developers who had relied on NBFC wholesale funding through 2015–18 — Supertech, Amrapali, Jaypee Infratech, and a long tail of Tier-2 developers defaulted or entered IBC through 2019–20. Small and mid-ticket residential construction seized for eighteen months. The regulatory response — scale-based NBFC regulation, tighter ALM norms, the prompt corrective action framework for NBFCs — took until 2021–22 to take final shape. Default rates spiked in the NBFC segment specifically; the broader corporate default rate was already trending down post-IBC. This was the cycle where the value of ALM stopped being a textbook idea and became a visceral one, and where I saw the butterfly effect at close range — IL&FS in September, and by October I was selling genuinely good-quality paper into a bid-less market just to meet redemptions. The funding freeze was universal; credit quality of the asset was not the operative variable, liquidity of the market was. The other lesson was about my own side of the table. Lender standards were at their loosest in the twelve months before the crisis, exuberance peaked in 2017, and credit fund inflows were at their maximum just as the cycle was about to turn. Investors — myself included — have an uncanny ability to flock in at exactly the wrong moment. Fund-flow behaviour, watched on oneself, is the most honest marker of a cycle inversion I have found. If you are an investor, you must beware the confident fund managers!

6. What Survives — The Mirror

The literature on Indian defaults is ten times richer than the literature on what survives them. A few traits recur across firms that made it through at least two of the four cycles above.

One, leverage discipline through the easy years. The survivors did not lever when capital was cheap. Asian Paints and Pidilite ran sustained capex across the 2000s and 2010s without materially levering the balance sheet — both have spent most of the last two decades net debt-free or near it, and Pidilite remains so today. A long tail of listed compounders followed the same pattern. That discipline is what let them absorb stress rather than be defined by it.

The counter-example sits in the same generation. The Birlas levered the group hard to fund Hindalco's 2007 acquisition of Novelis, and quickly realised that the market doesn't care who you are. The lesson took: Hindalco and the wider Aditya Birla group have not taken on comparable leverage since. A parallel lesson played out at the promoter level a cycle later — founders who had pledged shares to raise personal liquidity watched a market drawdown turn a stock-price move into a margin call on the company itself. Many of those who survived the scare unwound or sharply cut their loan-against-shares exposure, and have run cleaner promoter balance sheets since.

Two, a dominant position in a domestic demand category where the product is sticky and the consumer is sticky. Distribution depth, brand equity, and category adjacency let the survivor absorb price shocks without losing share. The defaulters were concentrated in commodity-adjacent or licence-adjacent businesses where the firm was a price-taker on both the input and the output. Survivors were price-makers on at least one side.

Three, boring promoters. The survivor set is conspicuously light on charisma. Firms whose founders gave long interviews, acquired aviation assets, launched cricket franchises, and ran for political office are over-represented among the defaulters — Mallya, Siddhartha, Wadhawan, Ambani (ADAG branch), and a longer tail. Firms whose promoters are visible only at the annual AGM and speak mostly about capacity utilisation and working capital days survive. Part 1's conversation-based integrity test has a cheaper first filter: google the promoter, and be sceptical if the top ten results are not about the firm. Same rules apply to fund managers also. Wonder what you will think of me. 😊

Four, and this one is painful: an auditor who has been there for a full cycle. Firms that survive tend to have stable, mainstream auditors with multi-cycle tenure on the account. Firms that default tend, in my observation, to show auditor changes in the year or two before the default, or mid-tier auditors on books that really warranted a Big Four. I don't have a clean dataset on this, but I'm not the first to notice the pattern. The NFRA was set up in 2018 precisely because of a sequence of audit failures at listed Indian firms — IL&FS, DHFL, CG Power, Coffee Day, Yes Bank — and SEBI's scrutiny since has been disproportionately focused on exactly the combination I'm describing: auditor change close to the event, or an auditor out of proportion to the size and complexity of the book. The auditor is a trailing indicator of governance. Sort of.

7. Where We Are in the Current Cycle — The Investor Playbook

The punchline. We are in the deceptive phase of the cycle — headline default rates are low, spreads are tight, a zombie cohort from the COVID response is being carried quietly on bank books, and the macro setup still reads as favourable for credit. This is when underwriting mistakes look like wins, and when every credit manager looks skilled.

One, underwrite the book assuming this phase ends, not that it persists. Every one of the four cycles began in a tape that looked like this one. My working hypothesis — offered as pattern-match rather than a statistical finding, and meant to be tested against the reader's own portfolio — is that applying Part 1's HERO framework to a current-cycle book will surface stress in four predictable places: asset-liability mismatch in finance-sector exposures, promoter leverage in mid-market industrials, unit economics in anything consumer-tech or SaaS, and working capital in export-oriented mid-market manufacturers. Treat that as a prior to stress-test, not a conclusion to rely on.

The best of credit is almost certainly behind us

A stronger claim to close this section on. The stress pockets I think I can see today are the ones described above. Four more are forming that I want to flag because they are not in the default data yet, and by the time they are, it will be late to reposition.

First, residential real estate. Inventory-to-sales data shows roughly 6–7 quarters of unsold stock across the top eight cities — absorption is keeping pace at the system level. If AI's progress through 2026–28 is what I expect — genuine displacement in white-collar back-office, IT services, BPO, analytics, junior consulting and junior legal — the buyer for ₹1–3 crore apartments in Bengaluru, Hyderabad, Pune and Gurgaon is exposed. My working hypothesis is that the marginal buyer in those four cities is disproportionately drawn from the cohort most exposed to AI displacement. I don't have the buyer-profile data to prove it, but if the hypothesis is right, developers levered against this demand face a problem that isn't in the absorption numbers yet. Watch pre-sales velocity in the top-ten listed developers quarterly, not just inventory.

Second, unsecured retail, as a second-order effect. The sector is visibly recovering from RBI's November 2023 risk-weight tightening; growth has re-accelerated off a lower base after the partial reversal in early 2025. The base case is benign. But if the white-collar job thesis above is right, household debt service ratios deteriorate with a six to twelve month lag, and unsecured retail is where that shows up first — personal loans, credit cards, consumer durable EMIs. The reason I call this second-order is that the underwriting base is strengthening at exactly the moment the demand-side shock arrives. That is a recipe for a sharper-than-expected turn, not a gentler one.

Third, auto. The sector sits at the intersection of two structural shifts simultaneously: the ICE-to-EV transition, and the emergence of China as the dominant global auto exporter. Chinese EV production is now over two-thirds of the global total. Indian auto has not been directly hit yet, partly because of high import duties, but it is exposed through two channels: auto-component manufacturers selling into global

OEMs whose own share is collapsing, and domestic two-wheeler and passenger-vehicle brands whose product cycles are ICE-heavy at a moment when Chinese EV cost parity is arriving in emerging markets. The EV-native Indian cohort — Ola Electric, etc. — faces the opposite exposure: a well-funded Chinese entrant could compress margins in the only segment where Indian auto is genuinely investing.

Fourth, the second-order effects of the Hormuz shock on small business. Section 2's discussion focused on refiners, gas-based power and petrochem — the obvious direct exposures. But a restaurant whose LPG supply has become unreliable, a small textile unit whose dyeing process runs on piped gas, a logistics operator whose energy cost has moved 20% in a quarter — these are the firms where the shock lands without any policy cushion.

All of this is why I think the best of Indian credit is almost certainly behind us. Default rates at current levels are a function of the benign phase continuing, not of underlying resilience, and the four vectors above are not priced into current spreads.

Even though the best of Indian credit, at current spreads, is almost certainly behind us, it does not mean 2015–18 is coming again — the balance sheets, the banking system, and the regulatory architecture are all in better shape than they were a cycle ago. Indian corporate balance sheets are genuinely cleaner. India Inc's net debt-to-equity ratio has improved for most of the last decade following the deleveraging wave; large-cap net debt to EBITDA runs materially below the levels that preceded the Twin Balance Sheet crisis. Interest coverage at the aggregate level has moved the other way — up, not down. The banking system, after the AQR-driven clean-up of 2015–20 and the migration of consumer balance sheets onto bank books, is better capitalised than at any point in the last two decades. What it means is that the stress in the next two to three years will be sectoral and firm-specific, not synchronised, and the investors who navigate it will be the ones who underwrote to that reality rather than to either of the two comfortable priors: 'another 2015–18' or 'the benign phase continues.'

8. What We Don't Know

This note is as honest as I can make it, which means naming the gaps. The biggest unknown sits in front of us in real time: the 2026 Iran supply shock does not yet have a shape. Whether the Hormuz disruption resolves in weeks or grinds on through the year, whether the US naval blockade becomes the durable new regime or a bargaining posture, whether India secures exemptions of the kind Iran announced in late March for Indian-flagged vessels — none of this is knowable from here. Every shock discussed in the note was written about after it was resolved. This one is being underwritten live, which is a different exercise. The honest posture is to mark down certainty, not to pretend it is higher than it is.

The deeper not-knowing is structural. There is no comprehensive hazard-model study of Indian firm defaults on par with Shumway (2001) or Campbell-Hilscher-Szilagyi (2008) for US data. There is no clean empirical separation of fundamental distress from strategic default in the Indian context. There is no published study I have found that disaggregates recovery rates by the cause of default — promoter-

fraud versus cyclical versus policy. The rating-agency default studies are actuarial, not causal. The academic work — Chari-Jain-Kulkarni, Kulkarni et al., Chahal-Ahmad, Aggarwal, Sengupta-Vardhan — is high-quality but fragmented. Until someone closes these gaps, investor notes like this one are stitched together from practitioner judgment, case studies, and the handful of rigorous studies that exist. I would rather tell you that than pretend otherwise.

Sources

Dangi, M. (March 2025). *Why Firms Fail: Insights from My Research with IIM Udaipur's Centre for Financial Research.* Mosaic Asset Management. (Part 1 of this series.)

Government of India, Ministry of Finance. *Economic Survey 2016–17*, Volume 2, Chapter on the Festering Twin Balance Sheet Problem. Principal author: Arvind Subramanian.

Reserve Bank of India. *Financial Stability Reports*, half-yearly, 2010–present.

Reserve Bank of India. *Report on Trend and Progress of Banking in India*, annual.

CRISIL Ratings. *Annual Default and Rating Transition Study*, FY2025 (and prior vintages from FY1988 onward).

India Ratings (Ind-Ra). *Corporate Transition and Default Study*, FY2025.

US Energy Information Administration. *Crude oil and petroleum product prices increased sharply in the first quarter of 2026.* Today in Energy, April 2026.

International Energy Agency. *Oil Market Report*, March 2026.

S&P Global Mobility. Automaker global market share series (as reported by CNBC, February 2026).

International Energy Agency. *Global EV Outlook 2025.*

Automobility. *State of China's Auto Market*, September 2025.

Chari, A., Jain, L., and Kulkarni, N. (2019). *The Origins of India's NPA Crisis.* Working paper (University of North Carolina / CAFRAL, Reserve Bank of India). Later expanded as NBER Working Paper 28435 (2021).

Kulkarni, N., Ritadhi, S. K., Vij, S., and Waldock, K. (2019 working paper; 2025 published). *Unearthing Zombies.* Management Science. Earlier Georgetown McDonough School of Business Research Paper No. 3495660.

Mannil, N., Nishesh, N., and Tantri, P. (2024). *When emergency medicine becomes a staple diet: Evidence from Indian banking crisis.* Journal of Banking & Finance, Vol. 161.

Chahal, R. J. K., and Ahmad, W. (2022). *Political connections, investment inefficiency, and the Indian banking crisis.* Quarterly Review of Economics and Finance, 85.

Chahal, R. J. K., and Ahmad, W. (2026). *Political Connections and Corporate Debt Structure: The Moderating Effect of India's Government-owned Banks.* Abacus.

Aggarwal, D. (2023). *Political Lending, Cost of Debt, and Cash Policy.* SSRN working paper (November 2023).

Acharya, V., Crosignani, M., Eisert, T., and Eufinger, C. (2024). *Zombie Lending: Theoretical, International, and Historical Perspectives.* Annual Review of Financial Economics, Vol. 16.

Chopra, Y., Subramanian, K., and Tantri, P. (2021). *Bank Cleanups, Capitalization, and Lending: Evidence from India.* Review of Financial Studies, 34(9), 4132–4176.

Mohan, R., and Ray, P. (2022). *The Roller Coaster Ride of Non-Performing Assets in Indian Banking.* CSEP Working Paper 22, Centre for Social and Economic Progress, New Delhi.

Sengupta, R., and Vardhan, H. (2017). *Non-performing assets in Indian Banks: This time it is different.* Economic and Political Weekly, 52(12), 85–95.

Insolvency and Bankruptcy Board of India. Commissioned research reports, including IIM Bangalore (May 2025) on *Behavioural Impact of the Insolvency and Bankruptcy Code*, and IIM Ahmedabad (March 2026) on *Effectiveness of the Resolution Process: Firm Outcomes in the Post-IBC Period.*

Chopra, A., Bansal, A., and Wadhwa, A. (2020). *Evidence of Predicting Early Signs of Corporate Bankruptcy Using Financial Ratios in the Indian Landscape.* arXiv:2008.04782.

Shetty, S. H., and Vincent, T. N. (2021). *The Role of Board Independence and Ownership Structure in Improving the Efficacy of Corporate Financial Distress Prediction Model: Evidence from India.* Journal of Risk and Financial Management, 14(7), 333.

Banaji, J. (February 2023). *Indian Big Business.* Phenomenal World.

Altman, E. I. (1968). *Financial Ratios, Discriminant Analysis and the Prediction of Corporate Bankruptcy.* Journal of Finance, 23(4), 589–609.

Beneish, M. D. (1999). *The Detection of Earnings Manipulation.* Financial Analysts Journal, 55(5), 24–36.

Shumway, T. (2001). *Forecasting Bankruptcy More Accurately: A Simple Hazard Model.* Journal of Business, 74(1), 101–124.

Campbell, J. Y., Hilscher, J., and Szilagyi, J. (2008). *In Search of Distress Risk.* Journal of Finance, 63(6), 2899–2939.

This note reflects the views of Maneesh Dangi and Mosaic Asset Management as of April 2026. It is intended for discussion and should not be read as investment advice. Figures drawn from third-party sources are cited above.