VOLUME 3

DS
BUILDING TRUST

FINSECURE LANDSCAPE REPORT

Where India needs to build Trust

A map India has to underwrite



SEPTEMBER 2026

CONTENT

Trust is not value in this report. It is infrastructure, as concrete as an identity rail or a payments rail, and far less finished than either. Each part asks the same question of a different segement: how much of this book can the present apparatus actually see?

FOREWORD	**Why this report?**	Rajit Bhattacharya, Co-Founder & CEO
PART ONE	**The last decade**	What alternate data and predictive AI delivered
PART TWO	**Why and where credit has to grow**	The 11-times expansion as a trust problem
PART THREE	**The challenges**	3 shapes of the gap, days, no record, out of reach
PART FOUR	**Where trust has to be re-engineered**	Risk, KYC, monitoring and income
CLOSING	**The people and the platform**	Trust Dialogues and about Data Sutram

Every generation of Indian lending has been built on a new way of knowing whether to trust a stranger.

For most of the last century that was the branch manager, who knew the family. Then it was the credit bureau, which knew the repayment record. In the decade just ended it became the data exhaust of a country that went digital faster than any in history, read by models that could finally make sense of it.

I have spent that decade inside the third of those transitions. When we started Data Sutram in 2018, the argument that a lender could trust a person on the strength of how they lived was still an argument. It is not any more. The alternate-data view of a borrower has moved from experiment to standing policy at some of the largest lenders in the country. Part I sets out what the decade delivered.

It also taught three lessons that matter more than any of the technology, and they are why this report exists. Every instrument we built was pointed at the population we could already see, and the share of first-time borrowers entering the system fell while we built them. When a system cannot assess a person, the market routes around the problem rather than solving it. And expansion without better assessment does not produce inclusion; it produces a cycle, as the microfinance book has just shown.

Those lessons matter now because the next transition is larger than the last. India intends to be a thirty-trillion-dollar economy by 2047. Worked through, that requires roughly ten times today's credit, growing faster than the economy every year for two decades.

The people already borrowing cannot supply that. It has to come from the sixty crore who are eligible and outside, and the technology that would let us see them has arrived at exactly the moment it is needed.

So this report is an attempt to bring the learnings of the digital era into a framework for the AI era. It asks why credit has to grow and where. It takes each segment of Indian credit, in the grid the industry already reads, and asks how much of it the present apparatus can actually see. It names the seven bets that have to be made and who has to make them, and sets out how risk has to be rethought for a book ten times larger, in a world in which every customer has a software agent. It closes on fraud, because the same technology that lets us see the next sixty crore lets other people impersonate them.

Trust is not a value here. It is infrastructure, as concrete as an identity rail or a payments rail, and far less finished than either. This is a map of where it still has to be built. It draws on public data, on our own platform, and on a season of conversations with senior practitioners across Indian lending, whose thinking informs the argument throughout without being quoted. The argument itself is ours, and so are its errors.

Rajit Bhattacharya
Co-Founder and CEO, Data Sutram

PART 01

The last decade

What alternate data and predictive AI Actually delivered, and what the industry data says happened as a result.

WHERE TRUST HAS TO BE BUILT

The last decade built trust in three things: that a person is who they say they are, through Aadhaar and eKYC; that money will move, through UPI; and that a borrower's past can be read, through the bureau and then alternate data and predictive models. Each is now infrastructure. What it did not build is trust in anyone without a record, trust that renews after disbursement, or trust that can be established faster than a document can be read. Those three gaps are the rest of this report.

The decade from 2015 to 2025 is the one in which Indian lending stopped being a physical business. The industry's own account of it is mostly a story about speed and cost. That misses the more important half.

The rails came first, and were adopted late

Before the global financial crisis, using the bureau was optional. Several lenders believed they underwrote better than the bureau did, and paid for it in 2008 and 2009. The crisis made bureau-based lending universal. That is the decade's first lesson and it is rarely stated: the infrastructure was built in the mid-2000s and adopted only when a crisis forced trust in it.

From roughly 2015 the public rails were laid. Aadhaar and eKYC, conceptualised in 2013 when no bank would touch it, has since processed more than twenty-five billion verifications. NACH and eNACH. UPI from 2017, which settled payments and, as a side effect, created the richest behavioural data source the lending industry has ever had. And Account Aggregator, which by July 2026 carried 538 million cumulative consents across 326 million linked accounts, with 1,020 financial information users and 176 providers connected.

2,500 Cr+
Aadhaar / eKYC verifications processed since inception

53.8 Cr
Cumulative Account Aggregator consents by July 2026, across 326M linked accounts

2017
UPI's launch year now the richest behavioural data source lending has

The net effect, as the head of one large lender put it to us: in 2010 digital lending was an option that banks exercised and physical lending was the norm. By 2026 that has flipped. Digital is the main way. Physical is the exception, and the last holdout, the mortgage with its wet signatures, is expected to sunset within three to five years.

Alternate data went from experiment to policy

A credit bureau file answers one question with authority: how has this person handled the formal credit they already hold? For a borrower who has not built that history, lenders needed an answer to a different question: what does this person's current pattern of life say about how they will handle credit? Over the decade, that second view moved from pilot to a standing input in policy at scale. On our own platform, the same score is now consumed by conservative bank books and high-growth consumer-durable books alike, which is the strongest evidence we have that it has become infrastructure rather than a niche tool.

150%

Enquiry volume growth for the **Trust Score** of a borrower (Apr - Aug 2026 vs the prior full year), against 6.1% industry retail-enquiry growth. Demand for alternate assessment is outrunning credit itself.

Predictive AI matured, and was sold on the wrong argument

The predictive layer is now genuinely mature. Loans below ₹5 crore are straight-through at several large institutions with no human in the loop. Around 80% of call-centre volume is handled by bots, an outcome that took roughly eighteen to twenty-four months from first deployment and that nobody predicted; at one of the largest NBFCs, 85% of pre-delinquency collections calls are now made by voice agents. A full AI decisioning stack runs at roughly one-seventh to one-eighth of the cost of a human underwriter, file for file.

That cost figure is real and it is close to irrelevant, and the fact that the industry led with it is the decade's most consequential mistake. Even a model that is ninety per cent cheaper than a human does not matter, because what breaks a lender's back is never the human cost. It is delinquency in the profile. Sold as a cost story, AI in credit was bought by operations and measured on headcount. Sold as a prediction story, it is bought by risk and measured on loss. Only the second compounds. Lenders that made the prediction bet early report credit costs on the affected books falling by more than half over two to three years, with market share dipping first and then recovering as the models bedded in.

~80%

Call-centre volume now handled by bots in 18–24 months from first deployment

1/8th

Cost of a full AI decisioning stack versus a human underwriter, file for file

>50%

Fall in credit costs on affected books over 2–3 years for early prediction-bet lenders

What the industry data says happened

Set the decade's technology against its outcome and the picture is uncomfortable. Total retail lending reached ₹170.2 lakh crore in March 2026, growing 16.6% a year. Retail 90-plus delinquency improved from 1.6% to 1.3%. By every conventional measure the decade worked.

And yet balances grew faster than accounts in every product the bureau reports, 16.6% against 6.2% overall, which means lenders extended more credit per customer rather than reaching more customers. The new-to-credit share of originations fell from 16.6% to 14.8%; the below-prime share fell from 27.6% to 26.5%. The industry's word for this is premiumisation. A less flattering word is deepening. The system became very good at lending more to the people it already understood, and narrowed at the edge where anyone new would enter.

₹170.2L Cr

Total retail lending, March 2026, growing 16.6% a year

1.6% →1.3%

Retail 90-plus delinquency across the decade it improved

16.6% →14.8%

New-to-credit share of originations it fell

THE INHERITANCE

That is the inheritance the AI decade starts with. The rails exist. The second view of a borrower exists and is being bought. The models work. And the share of Indians entering formal credit for the first time went down.

PART 02

Why and where credit has to grow?

A national target with a consequence nobody states: credit has to grow faster than the economy, every year, for twenty-one years. And it has to grow in the products and the places where trust is thinnest.

WHERE TRUST HAS TO BE BUILT

The eleven times expansion is a trust problem before it is a capital problem. Credit grows only where a lender can establish enough confidence to price a loan and the growth has to come from exactly that segments where that confidence is thinnest.

In July 2024, NITI Aayog tabled an approach paper at the ninth meeting of its Governing Council setting out an ambition that has since been repeated in almost every speech about India's economic future: a thirty-trillion-dollar economy by 2047. Most commentary argues about whether it is achievable. We are interested in what it implies for lending, if it is.

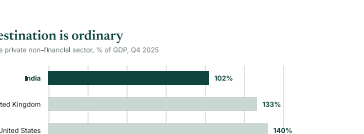
India's nominal output for FY2025-26 was $3.92 trillion. Reaching $30 trillion in twenty-one years requires growth of about 10.2% a year in dollar terms. That is the well-worn part of the debate. The part that is not discussed is the credit side. Credit to India's private non-financial sector stood at 102.3% of GDP at the end of 2025. The United Kingdom is at 133%, the United States at 140%, the average across all economies the BIS reports at 153%.

If India converges toward the UK's level, a $30 trillion economy carries a private credit stock of about $39 trillion. If it converges toward 150%, the figure is $45 trillion. Today the stock is about $4 trillion. The expansion required is therefore not the seven-and-a-half times implied by GDP alone. It is between ten and eleven times: 9.6x at the UK level, 11.2x at 150%, which is 11.4% to 12.2% compounded in dollars for twenty-one years.

EXHIBIT 1

The destination is ordinary

Credit to the private non-financial sector, % of GDP, Q4 2025



India today sits below every advanced economy. Convergence toward the UK implies a private-credit stock near $39tn; toward 150%, about $45tn — roughly eleven times today's $4tn.

Source: Bank for International Settlements, total credit to the private non-financial sector, adjusted for breaks, Q4 2025.

The arithmetic is not ours alone. Senior lenders we have discussed it with work it through on the same public inputs, reach the same $45 trillion, and arrive at the same eleven times in twenty years. The number is not in dispute among the people who would have to deliver it. It is not yet treated as an operating constraint.

THE HONEST CAVEAT

The $30 trillion target is contested. Reaching it requires around 9.25% annual growth in dollar terms; India has recorded real growth of 8.5% or better in only four years since the 1980s; the rupee fell roughly 7% in the year to September 2026, which makes a dollar target harder. None of that changes this chapter. If the target is met, credit grows tenfold. If it is missed but India develops normally, credit still grows several times over and the same population has to be underwritten. The trust problem is identical in both worlds; only the timetable moves.

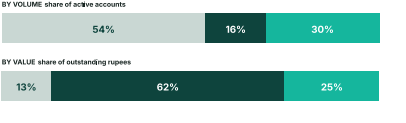
If India is to carry $45 Trillion of credit, the growth has to come from the large ticket.

The retail book has two shapes, and they do not overlap. Count the loans and India is an unsecured, small-ticket lender: personal loans, cards, consumer durables and two-wheelers are 54% of all active retail accounts. Count the rupees and it is a secured and business lender: those same four products are 13% of outstanding value, while housing, vehicles and sole-proprietor business lending, 16% of accounts, carry 62% of the value.

EXHIBIT 2

Two shapes of the same book

Share of retail credit by product group, value against volume, March 2026



Count the loans and India is an unsecured, small-ticket lender. Count the rupees and it is a secured and business lender. The two shapes barely overlap a decade of fast unsecured origination widened the count and barely moved the value.
Source: CRIF High Mark, How India Lends, March 2026, portfolio snapshot. Shares rounded.

This is the chart that explains why the digital decade did not move the arithmetic. Unsecured retail is where the decade happened, and it is 13% of the value.

THE MAP ON ONE PAGE

Where credit has to grow

By value, growth has to come from secured and business credit. the only books that can scale tenfold without the Reserve Bank stopping them. By people, from the sixty crore eligible adults outside the system. By geography, from the north and east and beyond the top hundred cities. Trust has been built in the metros and the south, for large tickets. It has to be built everywhere else.

Segment	Book	Growth	Trust gap	Where it has to grow
Home loans	₹44.4L cr	+9.4%	Built - slow	Beyond the top eight cities, which hold 52% of value; below the ₹75L ticket that now takes 41%; KA, TS and UP growing fastest
Business, sole-proprietor	₹51.6L cr	+19.7%	Absent	The 7.92 cr establishments of which 41% have ever borrowed formally; the unincorporated majority outside GST and ITR
Gold loans	₹18.6L cr	+50.4%	Substituted	Not more gold: assessed products for the 9 cr pledge customers, a third of them in Tamil Nadu; UP and TS growing fastest
Personal loans	₹16.5L cr	+12.9%	Thin	Above the ₹1L ticket where 89% of loans and 19% of value sit; BR, UP and RJ growing fastest, BR and WB most stressed
Two-wheeler and CD	₹2.9L cr	+17%	Thin	Already 42 to 51% beyond the top hundred cities; the job is the second loan, not the first; BR, MP and RJ most stressed
Rural and microfinance	₹3.3L cr	−13%	Absent	Back to the 1.8 cr borrowers who left; the eastern states where smartphone penetration and income measurement are weakest
Credit cards	₹3.4L cr	0.0%	Built	Nowhere under current underwriting; 41% of new cards from the top eight cities and issuance falling
Young and new-to-credit	n.a.	n.a.	Absent	Cross-cutting: 3 to 4 cr turn 21 each year; 46% of new card customers under 30 are already semi-urban or rural

Sources: CRIF High Mark, March 2026; MFIN, March 2026; MoSPI ASUSE 2025; CIBIL and SIDBI MSME Pulse, July 2026; TransUnion CIBIL, May 2026. Ratings are defined at the start of Part II.

Built apparatus can assess most of the population

Substituted lends without assessing, via collateral

Thin works for part, fails at the growth edge

Absent most of the population cannot be assessed

A tenfold expansion is survivable only if assessment improves faster than volume does.

A tenfold credit expansion is also an excellent description of how a country manufactures a debt crisis. Credit growing faster than income, for two decades, into populations with no repayment history, is the setup for every household-leverage bust in modern financial history. The objection is not hypothetical. It is visible in the microfinance book now, and Part III-C sets out the numbers. The mechanism in earlier cycles was not bad customers. It was good customers lent far more than they could repay, until the book that was built to include them became the trap that excluded them.

Two numbers you will see quoted differently

Bureau data puts the credit-active base near 29 crore; several practitioners cite roughly 40 crore. The higher figure counts everyone with a trade line ever recorded; the lower counts those currently active. We use the current-active figure. Retail aggregates also differ between bureaus because of coverage: CRIF reports ₹170.2 lakh crore, TransUnion CIBIL ₹161 lakh crore. Each figure in this report is cited with its source and the two are not mixed within an argument.

PART 03

The challenges

Three obstacles, not eight products. Each part groups the segments that share the same reason trust cannot be built, and asks one question: how much of this book can the present apparatus actually see?

The three parts are organised by obstacle rather than product, because that is how the challenges cluster in practitioners' own accounts.

A · Trust that takes days: secured lending and business documentation, where the borrower is already trusted and the process that proves it is slow.

B · Trust with no record: small-ticket unsecured, the entry products, the young, cards at their ceiling, and gold, where trust is blocked by the absence of a history.

C · Trust out of reach: rural households, microfinance and the micro-enterprise, where trust is blocked by cost and channel.

Four ratings are used across the map

Built	The apparatus can assess most of the addressable population; where growth is limited, it is by the time assessment takes, not its accuracy.
Substituted	The market has found a way to lend without assessing the person, typically through collateral.
Thin	The apparatus works for part of the population and fails at the edge where growth would come from.
Absent	Most of the addressable population cannot be assessed by any instrument in current use.

Trust that takes days

Home loans, loan against property, vehicle finance and the documented end of business credit. The borrower is already trusted. The process that proves it takes a week.

Secured lending is where India's credit system behaves best and grows least. Housing delinquency is low and improving across every lender type and ticket band. The borrower here is not the problem. The lender already trusts him, and the repayment record proves it. What takes time is establishing that trust: a legal check on the title, a technical valuation, a field verification, a personal discussion, and a full underwrite. Each step is a document a person has to read.

Product	Book, Mar-26	Growth	Active loans	PAR 31 to 180	Where it sits
Home loans	₹44.4L cr	+9.4%	2.36 cr	2.1%	52% of value top-8 cities; ₹75L+ now 41% of originations
Auto loans	₹9.3L cr	+13.9%	1.72 cr	2.7%	Balanced urban and rural; ₹5 to 10L tickets dominate
Sole-proprietor business	₹51.6L cr	+19.7%	6.74 cr	3.3%	CV, CEL, BL, LAP; fastest-growing large book after gold

Source: CRF High Mark, How India Lends, March 2026.

At one of the largest secured lenders the cycle from sales contact to money in the account averages eight days, and the delinquency that results is among the lowest in retail. The trust is real; it is embedded in a process, and the process is slow. Housing balances grew 9.4% in FY26 while active loans grew 2.5%; the ₹75 lakh-plus band now accounts for 40.7% of origination value, up from 33.6% two years earlier; 51.8% of value originates in the top eight cities. A process that costs a week can only be spent on customers who justify it.

What practitioners tell us

Chief risk officers draw the boundary the same way. Time has been solved for unsecured products; loan against property, mortgage and auto remain blocked by property registration, RTO processes and land-record digitisation, which no lender can fix. Property records in Vietnam are available with a click, with consent and with governance, the way a bureau score is in India, and here they are not. The problem is not the technology; it is bringing that infrastructure into being, which is as much a political-economy problem as a technical one. Of 8.56 crore income-tax returns filed in FY2024-25, only 11.4 lakh were filed by companies and 16.6 lakh by partnership firms the overwhelming majority of Indian businesses are unincorporated and invisible to every underwriting process built around audited financials.

Where AI has to enter

Document intelligence on the secured file	Title deeds in a regional language that took a specialist ten to twelve hours to read against government records are read in under thirty minutes at leading lenders, and underwriter throughput has roughly tripled. The most mature AI application in secured lending, and the least discussed.
Collapsing the process, not the standard	Each step that manufactures trust by hand KYC, income estimation, field investigation, title, valuation, the credit memo done by a machine in minutes with the same confidence, and a human taking only the final call. The destination is any file below fifty lakh rupees underwritten with no human at all.
Trust established before the application	A borrower whose footprint has been read continuously arrives pre-qualified, so the process has nothing left to establish. That is how eight days become ten minutes, and the only route to it for mortgages.

Where government and the regulator have to enter

Digitise the registry	The eight days are a records problem: land records, property registration, encumbrance and RTO data a lender still fetches and verifies by hand. That belongs to state and central government, and no lender can build around it.
A single window for the business	GST, Income-tax, MCA and Udyam describe the same enterprise and do not talk to each other. A consented pull across them of the kind the Unified Lending Interface promises would do for the business what the bureau did for the individual.
A registry of the unincorporated	Beyond private limited companies there is no central register of who owns what, so the family behind a proprietorship is invisible. Creating one is a statutory act, not a technology one.
The standard, not the speed	The regulator's role is to hold the bar where it is while the process under it collapses: same delinquency, same disclosure, same consent, in minutes.

Trust with no record

Small-ticket unsecured lending, the entry products, the young, cards at their ceiling, and gold. Trust is blocked by the absence of a history.

Consumer durables, 10.2 crore accounts averaging under ₹10,000 outstanding, are the entry point for most first-time borrowers into formal credit. Two-wheelers carry the highest early delinquency in retail. Cards show the end of the road: balances flat, issuance down 15% for a second year, a product refined over three decades on the population that was always visible and topping out at roughly one in eight eligible adults.

Product	Book, Mar-26	Growth	Active loans	PAR 31 to 180	Where it sits
Personal loans	₹16.5L cr	+12.9%	12.2 cr	2.4%	<₹1L = 89% of volume, 19% of value; NBFCs 91% of volume
Two-wheeler loans	₹1.9L cr	+15.1%	3.72 cr	4.1%	51% of value BT100; NBFCs 70%; highest PAR in retail
Consumer durable loans	₹1.0L cr	+20.8%	10.2 cr	1.7%	42% BT100; entry product for first-time borrowers
Credit cards	₹3.4L cr	0.0%	11.0 cr	3.0%	41% of new cards top-8 cities; issuance −15%
Gold loans	₹18.6L cr	+50.4%	8.99 cr	1.2%	Southern states; NBFCs 49% of volume; ticket-driven

Source: CRIF High Mark, How India Lends, March 2026. Card PAR shown is 31 to 90.

The record that is not there

The circularity is stated plainly by everyone who runs a marketplace: to have a good credit score you need credit, and to get credit you need a good credit score. Roughly 15 crore Indians sit on exactly one trade line. The second blind spot surfaces in the kind of call every lender has had: a customer earning ₹50,000 with ₹20,000 of EMIs is offered four lakh and asks for seven, because his wife earns 35,000, a fact the engine had no field for. Repayment comes from income, but resilience comes from wealth, and household wealth is exactly what an individual assessment misses.

The customer who has changed

Nobody in the industry describes the young as a lending problem; they describe a generation the industry was not built for. This is a credit-native generation. It has only known UPI, wallets and pay-later; it has never queued at a branch and never expects to. Credit, for it, is a feature one expects inside the transaction at the checkout, in the ride, on the marketplace, decided in real time. It attaches no stigma to borrowing: nobody wants a car loan, they want a car. It is not loyal; it will move to whoever gives it a better rate, faster.

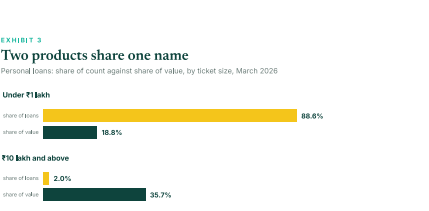
The door and the bounce

Pay-later, consumer durable and two-wheeler loans now run with no human in the loop, yet a retail book of a hundred thousand borrowers nobody knows gets no monitoring until a bounce. And the door matters as much as the decision.

EXHIBIT 3

Two products share one name

Personal loans: share of count against share of value, by ticket size, March 2026



Sub-₹1-lakh loans are 88% of personal-loan originations by count but a fifth of the value, and carry the highest stress.

Source: CIRF High Mark, How India Lends, March 2026. NBFCs originate 81.3% of volume and 38.8% of value.

Where AI has to enter · trust with no record

The family, not the individual	For a first-time borrower the strongest record that exists is not theirs it is their household's. Mapping an applicant to the family network, and reading the parents' and the household's credit behaviour beside a thin file of their own, is the missing field in every rule engine, and for this generation the most valuable one.
Every product in ten minutes	Quick finance today means two or three product categories. For a generation that expects a real-time decision, every category, including the secured ones, has to reach the speed pay-later reached, and AI is the only way there without lowering the standard.
Credit embedded in the transaction	This generation does not expect to go to a lender. Every lending category has to rediscover itself inside the transaction and the experience to serve it. A lender that is not in the transaction is not in the consideration set.
A first assessment that needs no history	The most mature and most under-deployed application in the map. Our own data finds score separation largest for thin-file applicants and the smallest tickets, narrowing as the file matures.
Door-level risk, and monitoring as the second product	The existing score applied to the acquisition channel, so price, cut-off and monitoring follow the door; and the same borrower watched repeatedly, so a filter becomes an early-warning layer. No new model. Almost nobody does either.

Trust out of reach

The rural household, the microfinance borrower and the micro-enterprise: the largest populations on the map, and the ones where the apparatus is largely absent.

This is the segment the industry most recently tried to widen, and the result is the most instructive thing in this report. Between June 2024 and March 2026 the number of microfinance borrowers fell from 8.7 crore to 6.9 crore, while average loan size rose 23% and portfolio at risk beyond 180 days tripled.

Segment	Book	Growth	Borrowers	Stress	Where it sits
Microfinance (GLP)	₹3.3L cr	−13.2%	6.9 cr	16.3% PAR 180+	Down from 8.7 cr; average ticket +23% to ₹61,500
Micro-enterprise	—	—	7.92 cr est.	—	41% ever formally credited; 11.4 lakh file as companies
Two-wheeler (rural share)	₹1.9L cr	+15.1%	3.72 cr	4.1% PAR	51% of value BT100; worst PAR in BR, MP, RJ

Sources: MFIN Micrometer, March 2026; MoSPI ASUSE 2025; CIBIL and SIDBI MSME Pulse, July 2026; CRIF High Mark, March 2026.

THE LESSON OF THE MICROFINANCE RETREAT

A system that cannot assess new borowers does not become more careful about everybody under stress; it retreats to the familiar and concentrates there. It lent more, to fewer people, the ones it already knew

The household that cannot be measured

The people who know rural lending best appear, at first, to disagree. Those who build models say data is not the constraint: fill rates for a rural customer are as high as for an urban one. Those who run microfinance books say income estimation remains a very inexact science. Both are right. The first is describing the digital footprint. The second is describing what it does not contain: a household whose income is agriculture plus the woman's work plus a kirana counter, measured by a field officer's conversation checked against a static bureau, with harvest migration breaking every address-based assumption.

Income in more than one dimension

The deeper problem is that income for this population is being read with an instrument built for the salaried class: one employer, one credit, one slip. The rural household and the informal worker have several sources and none documented. Income here has to be read in several dimensions at once: mapped to the family, because the household pools it; to migration and travel, because where a person moves says as much about earning as any declaration; to the gig platform, because for a growing share that is the only earnings record there is; and as a pattern of small flows rather than one large one.

THE LARGEST SINGLE OPPORTUNITY IN THIS PART

Distribution has to move from the field to a voice-native assistant that does origination and collections in the borrower's own language, not the app, because many rural women still do not have smartphones. It is voice, the medium rural India already uses.

The business that is a family

A small business that defaults often vanishes as an entity while other businesses owned by the same family, still creditworthy, continue to thrive. A mechanism that maps that network is the most obvious use of AI in MSME lending, and nobody has one. The experiment has been run: with forty data points per gig worker, an NBFC could lend to 70% of applicants against 25% without them. Under the data-protection law that data belongs to the worker.

The last friction: identity, and who it excludes

The friction that finally decides who is reachable is not income or collections. It is identity. Ask the heads of large lenders what they would rebuild and they do not name underwriting: bureau data, alternate data, models, e-mandates and e-signatures all flow, and the one place friction remains is KYC. Proof of identity is largely solved by Aadhaar; proof of address is not, and a graduate who has moved from a northern hometown to work in Bangalore has no address proof there. The people address-based identity excludes are exactly the people in this part: the migrant worker, the seasonal labourer, the rural woman without a smartphone, the proprietor whose business has no registered address. KYC built for a settled, documented, urban population is a reach problem before it is a compliance one.

Where AI has to enter · trust out of reach

1 · Voice-native origination and collections	A conversational assistant that conducts the application in the borrower's language, assesses what it hears, and carries the same conversation through collections. At the country's largest microfinance lender, 85% of pre-delinquency calls are already voice agents; the front of the journey is next.
2 · Income from several dimensions at once	Area-level consumption and mobility as a base rate; the family network; migration and travel patterns; gig-platform earnings; and the undocumented business's GST, UPI, bank and lease trails. A field officer's estimate is checked against all of it rather than standing alone.
3 · Identity as triangulation	KYC rebuilt as signals held together the rural address pattern, the number's tenure, the device, the household, the consent chain rather than a document presented. This is the population most exposed to fraud, and a single credential is exactly what gets borrowed, sold or stolen.

PART 04

Where trust has to be re-engineered

The three obstacles in Part III share four roots. The risk framework was built for a world with fewer kinds of risk. KYC is still a document. The borrower is read once, at origination, and next when the payment bounces. And income is still a salary slip. Each has to be rebuilt, not improved.

WHAT HAS TO BE REBUILT

A · The risk framework, because it carries the categories of 1995 and the world has added climate, geopolitics and a counterparty that is software.

B · KYC, because every other form of trust inherits it, and a document that anyone can copy is no longer a proof.

C · Monitoring, because trust established once at origination is not trust; it is a snapshot of a person who is moving.

D · Income, because the instrument used to read it describes a minority of working Indians.

Part III described where trust is slow, absent or out of reach. Underneath each of those is a piece of the apparatus that was designed for a smaller, more documented, more stable lending book than the one India is now building. The four sections that follow take them one at a time and ask what re-engineering each would actually mean.

PART FOUR · A

Evolving the risk framework

Enterprise risk was built by accretion: credit, then fraud, then cyber, then operational, each arriving as a new department. Three categories now sit outside it, and each is already moving the book.

That is how institutions end up comprehensive about the last crisis and silent about the next. The risk heads we have spoken to say the same thing in different words: there is no past history to model the risks now arriving, so risk management has to become forward-looking, and the frameworks are not. Three categories in particular have to be brought inside.



Climate: a regional shock that moves along a value chain

Climate risk is usually filed under agriculture, which is the wrong drawer. A drought does not damage a farmer. It damages the farmer, then the input dealer, the transporter, the cold store, the exporter, the district's trading community, and the salaried employees of all of them. The exposure is regional and it moves along a value chain, and almost no Indian lender holds a stress scenario shaped like that. In October 2023 the South Lhonak glacial lake in Sikkim burst its moraine and sent a wall of water down the Teesta, destroying a 1,200 megawatt hydropower dam, bridges, roads and settlements in a single night. None of these show up in a borrower's file, and every one of them shows up in a district's repayment six to nine months later.

The counter-case deserves equal weight. Between one and one-and-a-half thousand monitorable indicators exist for any product in any geography; most can be bought and most are noise. A monitoring capability without a demonstrated link between signal and loss is a machine for crying wolf. The re-engineering is not "add climate data". It is a small number of regional stress scenarios shaped like a value chain, each tied to an indicator whose link to loss has been proven on the institution's own book, and each with an action attached.



Geopolitics: repricing faster than the model refreshes

Supply chains, energy prices, shipping routes and export demand now reprice on political events faster than a credit model refreshed quarterly can track. The West Asia conflicts of the last two years are the working example. The book's exposure is not to a country; it is to a corridor, a commodity and a sector, and very few lenders hold portfolio concentration by export corridor or input dependency as a live risk parameter.

The next geopolitical shock to the Indian book may be technological rather than military. The salaried class, the most trusted population in Indian lending, is concentrated in a few sectors and a few cities, and a large share of it works in the services that AI is now automating: software services, business process outsourcing, customer support, back-office finance. A personal-loan or mortgage book concentrated in the technology corridors of Bengaluru, Pune, Hyderabad and Gurugram carries a correlated exposure that no bureau score records. A framework that treats a salary slip as the safest evidence of repayment capacity has to start asking how safe the salary is.



Cyber: fraud as a service, and the counterparty that is software

7.1%

India's suspected digital fraud rate, against a 3.8% global average close to double

₹22k Cr

Reported cyber-fraud losses in 2025 across 28.15 lakh complaints

3/4

Share of that value in investment fraud alone

The people who run fraud defence describe a change in the nature of the adversary rather than its scale. Fraud has become an industry: there are people who commit fraud, and people who sell fraud as a service. It is organised crime with no hierarchy every operator an entrepreneur, no single lord, and when law enforcement arrests one, the next has never met him. And the newest layer is agentic: software agents already open accounts, fill applications and move money through chains of mule accounts faster than a rule can fire, and the payment networks are now issuing tokens to AI agents so they can transact on a customer's behalf. The best of the fraud, in the words of the people fighting it, is yet to come.

Detection is not defence: the safety valve

The structural disadvantage is that the fraudster has no false-positive problem. He can try a thousand times and needs one success. The institution cannot build infrastructure that inconveniences bona fide customers, so every control it adds has a cost the attacker never pays. That asymmetry is why a fraud strategy built on detection alone loses. The re-engineering the best fraud heads describe is a safety valve a water tank's overflow pipe, an old house's fuse that blows before the wiring does. What it requires is that the institution knows the moment something has been breached and reacts immediately: block the journey, hold the transaction, freeze the maintenance change, for twenty-four hours if need be, and analyse afterwards. Act and react first; the AI and the analytics follow.

Three things sit behind the valve. A strategy for each kind of fraud first-party, third-party and collusion. Authentication differentiated by what the customer is doing: viewing an account the easiest, a payment harder, and a maintenance change the hardest of all, because it is what precedes the theft. And the customer's own valve, which can only be awareness. Underneath all three, cyber and fraud risk have to be treated as a category with their own capital and loss line, and a trust layer has to exist for agents themselves.

THIS REPORT PROPOSES

A 5-star fraud protected rating, issued for every product category and every lender category, and published the way a car crash rating is published. Five years ago nobody buying a car asked about crash ratings. Financial products are at the same team, and have no equivalent.

The standard has to be by product and by lender category, and it has to test outcomes, not process: how quickly a breach is detected, whether a safety valve fires automatically, whether authentication is differentiated by transaction type, whether fraud loss is measured separately from credit loss, and what the institution's actual fraud rate per lakh transactions is against its category. Building the rating is a job for the industry with the regulator, and it is one this report's author is willing to help start.

PART FOUR · B

Rethinking KYC

Every loan and every transaction in India begins with the same question, and the answer is still a document. Documents can be copied, and now they can be faked. KYC has to become a triangulation of identity, not a check of a document.

The document has stopped being a proof

Because every hotel, telecom counter, courier desk, housing society and small lender has demanded a photocopy of Aadhaar and PAN for a decade, copies of almost every Indian's core identity documents exist in thousands of places, and a great many have leaked. A document that anyone can obtain is not a proof of anything; it is a starting kit. Onto that base the last two years have added generative tools: synthetic faces that pass a liveness check, cloned voices that pass a verification call, forged salary slips and bank statements. The result is a system in which the institution can no longer believe the document and has responded the only way a document-based system can — by asking for more documents, more re-verification and more time.

How the rest of the world moved

Advanced markets have quietly abandoned the idea that identity is a number. The social security number in the United States is treated as already compromised. Authentication has moved to the device and to behaviour. Singapore's national digital identity is a single consented credential any institution can verify without ever holding the underlying document; the European Union is issuing a digital identity wallet on the same principle. Identity has become multidimensional and continuous a triangulation of signals that agree with each other, rather than a single credential presented once. India built the world's largest credential-based identity system. It has the harder and more valuable transition to make on top of it.

THE DESTINATION

A journey with no document and no OTP in it, the person recognised rather than asked to prove themselves, the asset verified from records that agree, and a business recognised from the trail it leaves rather than the licence it holds.

Monitoring the borrower

A borrower is assessed twice: when the loan is given and when it goes bad. Everything in between, which is where the loss is decided, is unobserved.

The first trigger any retail lender gets is the bounce. It is the only trigger. Meanwhile the factors that change a borrower's capacity have multiplied beyond the normal cycles: a drought in the district, a layoff in the sector, a new loan taken elsewhere, a change in daily movement that says the shop has closed. Each leaves a trace in the borrower's footprint months before the bounce, and nothing in the standard stack is reading it. The honest reason is economics: monitoring is not a growth lever, so it never clears an investment committee against origination — and the cost of not doing it is invisible until the cycle turns, as it has in microfinance.

~13×

Delinquency, thin-file bottom vs top Trust-Score quintile, first six months of a bureau file

~15×

Bottom vs top quintile at just three months after booking — before the six-month view most lenders wait for

~7×

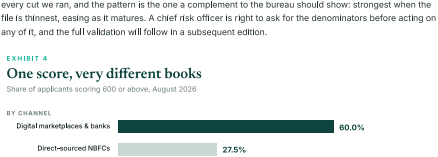
Separation among the smallest-ticket loans

Those are relative comparisons across score bands; this edition does not disclose the base rates, sample sizes or windows behind them, and a ratio without a base rate can flatter. The direction is consistent across every cut we ran, and the pattern is the one a complement to the bureau should show: strongest when the file is thinnest, easing as it matures. A chief risk officer is right to ask for the denominators before acting on any of it, and the full validation will follow in a subsequent edition.

EXHIBIT 4

One score, very different books

Share of applicants scoring 600 or above, August 2026



The same Trust Score, read across acquisition channels, describes different books — not better or worse ones. A dealer counter reaches customers far earlier in their credit journey, so risk enters through the door before any credit decision is made.

Source: Data Sutram Trust Score query data, August 2026; volume-weighted across funnels. Distributions at query, not approvals or outcomes.

What continuous monitoring has to mean

The re-engineering is a handover of primacy rather than a replacement. Early, before a formal history exists, the second view carries the decision. As the file matures, the bureau takes the lead, because direct repayment evidence should. And through the life of the loan, the second view continues as a monitoring layer between reporting cycles, refreshed on the live borrower rather than pulled once at origination. The unit of time changes from the origination event to the life of the exposure. The first signal of stress becomes something other than a missed payment a change of device, number or address pattern; a new enquiry elsewhere; a district scenario firing. Every signal has to carry a proven link to loss, or the credit team stops reading the red flags. And each signal needs an action attached.

A SECOND PRODUCT HIDING IN THE FIRST

The same borrower watched repeatedly is also the same borrower re-qualified repeatedly. Applied to the nine crore gold-loan accounts, onboarded, paying and scored for nothing beyond their collateral, monitoring becomes the largest acquisition engine in Indian credit, and the one area where the case for AI needs no new data at all.

Re-engineering income

Income is the number every credit decision rests on, and it is still derived from a bank statement or a salary slip. For most of the next sixty crore, neither exists. Income has to be triangulated, not read.

For the salaried borrower income is easy: one employer, one credit a month, one slip. That is the instrument every rule engine uses, and it describes a minority of working Indians. The rest earn in ways the instrument cannot see: a harvest, a daily wage, a kirana counter, a delivery platform, a remittance, a seasonal migration very often several at once. The consequence is not only exclusion. It is the debt trap the microfinance book has just produced, in which good customers were lent more than they could repay because nobody could measure what they could repay.

The gig worker, and democratising consumer data

The system that decides creditworthiness in India was built before the gig worker existed. The experiment that shows what is possible has been run: in a pilot in Bengaluru with around forty data points per gig worker, an NBFC could lend to 70% of applicants where it would have lent to 25% on the data it otherwise had. The same pilot found the product had to change with the income gig workers repay best weekly, not monthly. The data exists, on every platform the gig worker uses, and under the data-protection law it belongs to them. What does not exist is the mechanism to democratise it: a consented way for a worker to pull her platform earnings into a locker and share them with a lender of her choice, the way an account statement is shared through Account Aggregator today.

WHAT HAS TO CHANGE

For the lender: income as a triangulated construct in the rule engine, with fields for the household, the area base rate, the platform record and the conversation.

For the regulator and government: consented portability of platform earnings data, so a gig worker's income record is hers to share, and repayment products shaped to how people earn.

For the provider: the area-level base rates, the household map, the conversational assessment and the cash-flow read of the undocumented business, delivered as one income view rather than five data feeds.

India spent the last two decades building the rails that made credit fast. The next two depend on building the ones that make it safe to extend to people we do not yet know. The map in this report is of where those rails still have to be laid.

The people whose thinking shaped this report

This report is a landscape, and a landscape is only as good as the people who have walked it. Through the summer of 2026 we sat down with nine leaders of Indian financial services for our Trust Dialogues series and asked each the same questions: where does trust break today, where does it have to be built, and what would they tear apart and rebuild. Their answers are the raw material of Parts III and IV. The argument, the arithmetic and the mistakes are ours; the ideas that made it possible are theirs, and we are grateful.

Sudipta Roy
Managing Director & CEO, L&T Finance

Bhavesh Gupta
Former Chief Executive Officer, PayTm

Prithvi Chandrasekhar
CEO, COnsumer Finance, InCred

Amit Jain
Chief Risk Officer, Hero FinCorp

Sudipto Nag
Group Head & CCO, Axis Bank

Soumyajit Ghosh
COO, True Balance

Manish Agrawal
Senior Executive Vice President, HDFC Bank

Abhinav Bansal
Managing Director & Senior Partner, BCG

Vivek Chopra
Chief Operating Officer, Tata Capital

Rajesh Bansal
Former CEO, Reserve Bank Innovation Hub

Santosh Agarwal
Chief Executive Officer, PaisaBazaar

Ravi Sandepudi
Head of Platform Products, Socure

Our thanks also to CRIF High Mark and TransUnion CIBIL, whose published data gives the industry a common set of numbers to argue from, and to the team at Data Sutram who built the platform whose data appears in these pages and who worked through the summer to bring this edition to the Global Fintech Festival.

ABOUT DATA SUTRAM

An enterprise AI platform that solves for trust

Trust

An agentic platform that does the back-office work for banks.

Specialised agents that access the full context of a case and take the right decision across fraud, credit and recovery.

Alerts: Docs | Financial | Authenticate | Voice | Assesment

Auth

A data platform that processes customer risk from external sources and proprietary models.

DPDP-compliant by design. The Trust Score is its flagship, with seven risk bundles beside it.

FLAGSHIP · TRUST SCORE

A digital score that reflects a customer's stability in the digital world.

Built from proprietary data signals rather than bureau history alone.

BUNDLES:

Identity | Selfie-Fraud | Employment | Business | Financial | Traceability | Vehicle

Echo

A voice agent platform that enables conversational AI.

Human-like AI agents with near-zero latency and multi-language capability, running predefined workflows.

Voice Agents: Pre-EMI Reminders | Bounce Follow-Up | NPA Notice | KYC & Adress Verification | Personal Discussion

Speech to Text: Sentiment | Intent | Keywords | Compliance Signals

DS
BUILDING TRUST

About DS

Data Sutram also now known as DS is an enterprise AI platform that solves for trust. The platform brings together agentic workflows, intelligence from external sources, and proprietary models and algorithms to build solutions on trust across payments, lending, insurance and wealth, and is used by more than forty financial institutions in India. It has three products. Trust, Auth and Echo.

© 2026 Data Sutram/DS (Extrapolate Advisors Pvt. Ltd.). All rights reserved. The analysis, insights and underlying research in this report are the proprietary intellectual property of Data Sutram/DS. No part may be reproduced without permission. "Data Sutram", "DS" and related product names are trademarks of Extrapolate Advisors Pvt. Ltd.; all other marks belong to their respective owners. Information is current as at the reference periods stated. Provided for information only; not financial, legal or investment advice.

For more information visit us at: www.datasutram.ai