India FinTech Debt: AI Cash-Flow Underwriting with Explainability and Override Controls

A Lender Framework for Decision Lineage, Model Drift, Human Accountability and Facility Controls

Authored by Chennakeshav Adya
Independent Researcher
September 2026

Abstract

Indian fintech lenders can use bank transactions, Account Aggregator feeds, GST records, invoices, payment histories and bureau information to evaluate cash generation before conventional financial statements become available. Artificial intelligence can organise those signals, identify instability and support faster credit decisions. The financing case remains incomplete when a lender cannot reproduce an approval, explain its material drivers, govern manual exceptions or show how model performance reaches cash available for debt service.

This paper develops an evidence-gated framework for banks, private-credit funds, family offices and institutional capital providers financing Indian fintech lenders. It maps the chain from permitted data acquisition through feature construction, model output, credit policy, human review, disbursement, servicing, collections and funder cash. The framework separates model ranking from probability calibration, affordability, loan sizing and facility leverage. It requires time-stamped lineage, transparent benchmarks, out-of-time validation, cohort and vintage analysis, drift thresholds, reason codes, override logs, fallback procedures and independent challenge.

The paper translates that evidence into a debt structure. Receivables qualify only when the lender can demonstrate decision-time data, approved model and policy versions, enforceable asset ownership, controlled cash and servicing continuity. Dynamic reserves, concentration caps, model-change consent, reporting tests and cash-sweep triggers respond to deterioration in realised losses, data coverage, override behaviour or model calibration. A hypothetical facility illustrates the method. Every balance, rate, threshold and transaction term in that case is a modelling assumption and does not describe an identified lender, borrower, portfolio or investment.

The central conclusion is operational. AI can support fintech debt capacity when the institution proves the lineage of each decision and retains accountable human authority for consequential outcomes. Explainability and overrides then become components of credit architecture. Current Indian regulatory obligations and their application require advice from qualified Indian counsel and the relevant regulated entity.

JEL Classification: C53, G21, G23, G28, O33

Keywords: India fintech debt, cash-flow underwriting, artificial intelligence, model explainability, human override, NBFC, Account Aggregator, model drift, credit policy, facility covenants

1. Underwrite the facility and its cash-flow evidence chain

The capital provider should begin with the financing decision. A revolving warehouse, term loan, forward-flow purchase, securitisation exposure and corporate facility create different repayment claims. Each structure needs a defined borrower, asset perimeter, collection route, priority, reserve, servicing arrangement and remedy. The investment memorandum should state which cash repays the facility and which operating events can interrupt that cash.

Draw the legal and operational chain before reviewing a model. Identify the regulated lender, lending-service providers, digital applications, embedded-finance partners, data providers, credit bureaus, Account Aggregators, payment processors, collection agencies, security trustee and controlled accounts. Record who owns each receivable, who can alter a decision, who communicates with the borrower and who controls the collection account. RBI guidance keeps responsibility with the regulated entity when functions are outsourced.[1][14]

Reconcile the proposed collateral pool to executed contracts, disbursement evidence, servicing records, general-ledger balances and bank cash. The same loan identifier should connect the application, decision snapshot, agreement, receivable, repayment schedule, collection and write-off record. Exceptions need an owner and an eligibility consequence. Duplicate assets, missing agreements and unexplained cash timing weaken the borrowing base regardless of model sophistication.

Define the investment committee's actual decisions. The committee may approve a commitment, opening advance rate, concentration limit, reserve, model-change covenant or cash sweep. Each decision requires its own evidence. Ranking statistics help evaluate discrimination; realised vintage losses support reserve calibration; controlled-account evidence supports cash dominion; override records support confidence in policy execution. A single accuracy measure cannot answer all of these questions.

The resulting map connects data, underwriting and debt service. It allows the capital provider to test where a failure would enter the facility and which control would contain it. Legal, regulatory, tax, accounting and data-protection conclusions remain subject to current professional advice.

Figure 1. From borrower data to private-credit repayment

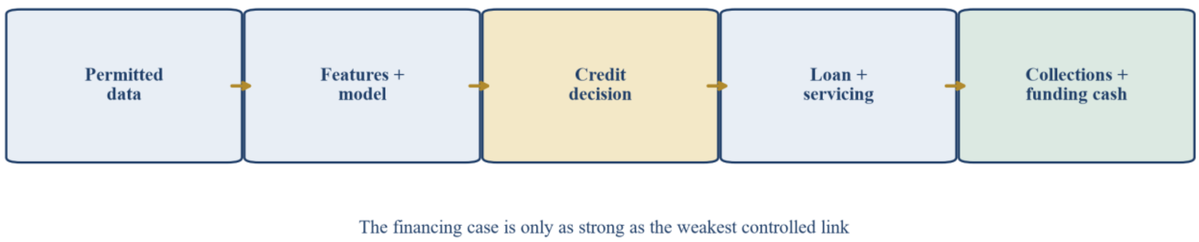

Management framework. Each arrow requires an evidence source, owner, control and financing consequence.

Table 1. Investment-committee questions before committing capital

Decision	Evidence required	Failure signal	Financing response
eligible receivables	executed loan records, borrower identity, policy version and payment status	missing decision history or duplicate asset	exclude until reconciled
advance rate	cumulative net loss, recovery timing, dilution and volatility by vintage	recent cohorts outside historical range	lower advance or increase reserve
concentration	product, geography, channel, employer, data source and risk tier	hidden dependence on one acquisition or data partner	concentration limit and reporting trigger

Decision	Evidence required	Failure signal	Financing response
model reliance	independent validation, challenger, calibration, stability and overrides	no reproducible decision or weak out-of-time result	cap eligibility and require remediation
cash control	bank statements, waterfall, servicing file and reconciliation	commingling or unexplained cash lag	controlled account and cash sweep
change control	policy, model, feature and vendor change logs	material change without approval or back-test	consent right and eligibility pause

The evidence standard depends on instrument, borrower type and regulated perimeter.

2. Place cash-flow underwriting inside the Indian regulatory perimeter

Cash-flow underwriting uses data that may originate outside a conventional credit file. Bank transactions, GST records, invoices, payment receipts, commerce activity and operational accounts can reveal turnover, seasonality, customer concentration, liquidity and payment behaviour. Each field needs a documented purpose and permitted use. Broad access to data creates control obligations that extend beyond statistical performance.

The Reserve Bank of India consolidated requirements for digital lending in the Digital Lending Directions, 2025.[1] The framework addresses regulated entities, lending-service providers, digital lending applications, borrower disclosures, data practices, disbursement, repayment, grievance handling and reporting. The RBI's regulatory handbook highlights need-based collection, explicit consent, audit trails, borrower choice and restrictions on storage by service providers.[2] Transaction counsel should confirm the provisions applicable to the lender, product and structure.

The Account Aggregator framework supports consented movement of financial information between regulated participants. RBI educational material identifies its potential for MSME cash-flow lending and states that the aggregator does not see or store the financial information it transfers.[13] The lender should evidence consent artefacts, requested data types, purpose, duration, revocation handling, failed connections and fallback decisions. A funding party should also understand how a revoked or unavailable connection affects monitoring after origination.

Credit-information reporting remains part of the control environment. RBI directions require frequent updating of borrower records and define responsibilities around submission and correction.[3] The lender should reconcile bureau pulls and submissions to its own loan records. Identifier mismatches, reporting delays and unresolved disputes can affect repeat borrowing, total indebtedness and risk selection.

India's Digital Personal Data Protection Act and notified rules create a wider data-governance perimeter.[4][5][6] Required evidence can include a data inventory, notices, permissions, processor terms, retention rules, security controls, rights handling and incident response. The RBI FREE-AI Committee report adds a financial-sector governance reference covering accountability, fairness, understandability, human oversight, monitoring and assurance.[7] These sources support an integrated diligence plan rather than a checklist isolated from credit economics.

Table 2. Regulatory and governance evidence map

Evidence area	Primary reference	Diligence artefact	Financing relevance
digital lending	RBI Digital Lending Directions 2025	regulated-entity map, LSP agreements, app inventory, KFS and grievance logs	enforceability, conduct and servicing continuity

Evidence area	Primary reference	Diligence artefact	Financing relevance
credit reporting	RBI Credit Information Reporting Directions 2025	CIC submissions, rejection log, correction process and refresh evidence	indebtedness, repeat borrowing and pool accuracy
personal data	DPDP Act and Rules	data inventory, notices, permissions, processor terms and retention	lawful use, remediation cost and interruption risk
AI governance	RBI FREE-AI Committee report	board policy, AI inventory, model files, approval and audit evidence	model change, accountability and reputational risk
outsourced technology	RBI outsourcing and IT governance directions	vendor diligence, service levels, audit rights and exit plan	data-feed, servicing and operational continuity
borrower treatment	KFS, pricing, redress and recovery requirements	offer screens, reason codes, complaints, resolution and recovery scripts	cash collection, disputes and regulatory exposure

This is a diligence map, not a legal conclusion. Qualified advisers should confirm current applicability.

3. Reconstruct cash-flow data before trusting predictive power

Every decision should be reproducible from information available at its original decision time. Build a source register covering provider, legal entity, field definition, currency, time stamp, refresh cycle, coverage, transformation, model use, retention and permitted purpose. Preserve the raw event and the derived feature. A reviewer should be able to move from an approval back to the source value without relying on a current database that may have changed.

Cash-flow features require economic definitions. Revenue credits should exclude internal transfers, loan proceeds, refunds and reversible entries where appropriate. Expense classifications should distinguish operating payments, tax, financing, owner drawings and one-off capital expenditure. Seasonality needs a sufficiently long window. Customer concentration requires reliable payer identity. GST and invoice data should reconcile to collections rather than serve as unchallenged proxies for cash.

Time leakage can create false predictive power. A later bureau update, post-disbursement collection tag, corrected transaction category or subsequently issued invoice may reveal the outcome. Freeze the source snapshot and transformation code used at approval. Conduct a feature-availability test for every variable. Leakage should invalidate the relevant validation result and trigger a controlled rebuild.

Coverage is an underwriting variable. Some applicants complete an Account Aggregator connection; others provide statements or remain outside the connected-data path. GST coverage differs by business profile. Embedded partners may deliver different fields. Report missingness and eligibility by product, channel, business age, geography and risk tier. A model should have an approved path for absent, partial and contradictory information.

Data suppliers and software vendors create continuity risk. Contracts should address audit rights, schema changes, incident notices, portability, subcontractors, service levels and exit assistance. The lender should test an outage and a corrupted-feed scenario. Funding eligibility should narrow when a critical feed fails and the approved fallback cannot reproduce the original policy standard.

Figure 2. Cash-flow data provenance and decision lineage

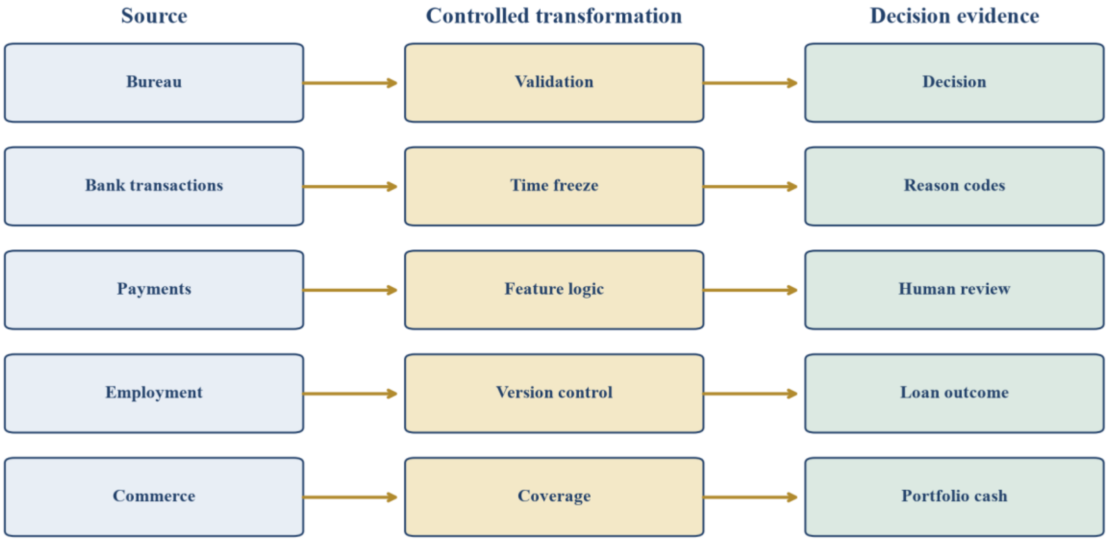

Proposed control architecture. Raw sources remain traceable through features, versions, decisions and outcomes.

Table 3. Minimum cash-flow data register

Field group	Required record	Core test	Stop condition
provenance	provider, system, contract, permission and permitted purpose	can the source and right to use be evidenced?	undocumented source or prohibited use
timing	event time, ingestion time, decision snapshot and refresh	was the value available before the decision?	leakage or retrospective substitution
transformation	raw field, code, aggregation window and feature version	can the feature be reproduced?	undocumented or non-deterministic transformation
coverage	eligible population, observed population and missingness by cohort	who is excluded when data are absent?	material unexplained selection effect
quality	validity, completeness, duplication, reconciliation and drift	does the feed remain within approved tolerances?	critical breach without fallback
retention	storage, access, encryption, deletion and archive evidence	are data retained and removed under the approved rule?	uncontrolled copies or unverifiable deletion

The register should preserve source-level detail and decision-time availability.

4. Separate probability estimates from credit policy

An underwriting model estimates an outcome under a stated target, population and data set. Credit policy converts that estimate into approval, amount, tenor, price, conditions and documentation. The distinction is central to financing. Losses can rise when the model remains stable and the lender changes a cut-off, increases limits, relaxes affordability rules or shifts acquisition towards a weaker channel.

Define each target precisely. Default probability, first-payment default, fraud, cash-flow volatility, prepayment and recovery answer different questions. State the observation window, performance window, treatment of incomplete outcomes and rule for restructurings or settlements. Recent originations need sufficient seasoning before they can support long-horizon conclusions.

Maintain a transparent benchmark. A scorecard, logistic model or rules-based cash test provides a challenger to a more complex system. Compare discrimination, calibration, stability, operational cost and economic value. Calibration deserves particular attention because facility reserves depend on the relationship between predicted and realised loss. A model can rank borrowers well while understating the absolute probability of default.

Validation should follow time. Train on earlier periods, tune on a separate sample and test on later originations. Repeat results across product, channel, geography, business age, data-coverage path and risk band. Economic regime changes, tax cycles and partner mix can alter relationships. Random splitting across the same period can conceal that instability.

Keep model and policy versions linked to each loan. The decision file should capture the data snapshot, feature version, model score, calibration, policy rule, limit, tenor, price, explanation and human action. This record allows the capital provider to attribute deterioration to changing borrowers, changing data, changing models or changing commercial policy.

Table 4. Model-validation evidence for a capital provider

Test	Question	Evidence	Financing use
discrimination	does the model rank higher-risk borrowers above lower-risk borrowers?	Gini, AUC, KS and confidence ranges by cohort	eligibility and monitoring
calibration	do predicted losses match observed losses?	calibration curve, expected-to-actual ratio and tail error	reserve and advance rate
stability	does performance persist over time and changing mix?	population and characteristic stability, rolling metrics	trigger and reporting frequency
challenger	does complexity add repeatable value?	out-of-time comparison with simpler model	reliance and fallback
sensitivity	do small input changes create disproportionate decisions?	perturbation and boundary tests	manual review zone
reproducibility	can a historical result be recreated from preserved evidence?	code, data snapshot, feature version and decision log	audit right and eligibility

No single metric proves suitability. Each test informs a financing or control decision.

5. Read cash-flow vintages as repayment evidence

Facility repayment depends on realised collections after credit losses, refunds, servicing costs and funding expenses. Model metrics should therefore connect to loan-level cash outcomes. Build monthly or quarterly vintages using origination date and track scheduled principal, collected principal, interest, delinquency, cure, restructuring, recovery, write-off and prepayment.

Compare vintages by product, term, ticket size, risk band, channel, business age, geography, model version, policy version and data path. A headline loss curve can improve while a fast-growing partner channel deteriorates. Report both absolute balances and rates. Small cohorts can show unstable percentages; large cohorts can dominate averages and conceal a local failure.

Cash-flow underwriting needs cash-flow diagnostics. Measure recurring inflow, volatility, lowest monthly balance, debt-service burden, customer concentration, seasonality and the frequency of negative liquidity intervals. Test whether each feature remains related to repayment in later periods. A relationship that weakens after a payments-platform redesign or GST schema change should trigger investigation.

Use expected-to-actual analysis after calibration. Compare observed defaults and losses with model estimates by band and cohort. Set minimum seasoning, sample and materiality rules. A ratio above one can indicate underestimation; a ratio below one can reflect conservative calibration or changing selection. The committee should review the economic cause before changing reserves.

Vintage evidence should enter the facility model. Advance rates, loss reserves and concentration limits can respond to mature loss, developing delinquency and coverage. Growth should not automatically create availability. Receivables from an unseasoned product or materially changed model can enter a separate pool with tighter limits until evidence matures.

Figure 3. Hypothetical cumulative net loss by origination vintage

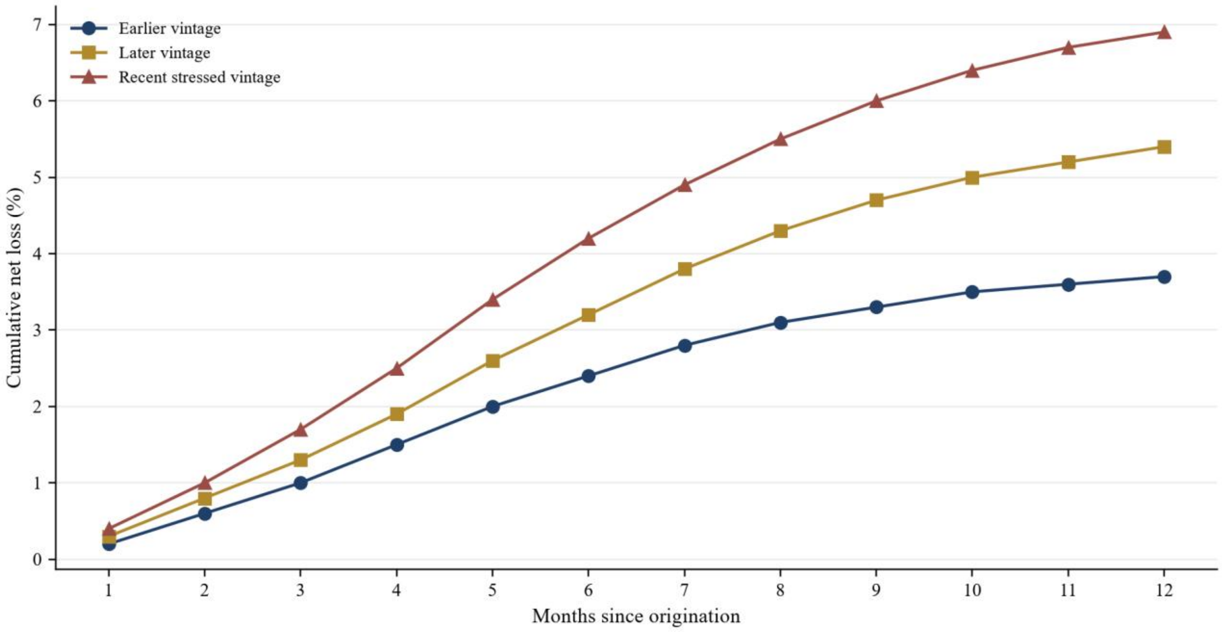

Percentages are modelling assumptions and do not describe an actual lender or portfolio.

6. Test inclusion and fairness as credit-control issues

Cash-flow data can widen access for businesses with limited formal borrowing history. It can also reproduce exclusions through geography, language, operating model, digital access, transaction type or partner selection. Fairness testing should examine data access, approval, pricing, limit, tenor, error, explanation, override and subsequent outcome. The purpose is to identify material differences and their drivers within the applicable legal framework.

Begin with coverage. Determine which businesses can connect accounts, provide GST records or produce machine-readable statements. Compare missingness and failed connections across relevant cohorts. An apparently neutral feature may encode access to technology, formality or a particular platform. The lender should document its economic rationale, incremental value and fallback path.

Evaluate decision and error measures. Approval rates, predicted risk, realised loss, false-positive and false-negative rates, calibration and reason-code distribution provide different views. No single measure establishes equitable treatment. The institution should define its objective, threshold, sample rule, review process and remediation in consultation with qualified advisers.

Human review also needs testing. Compare referrals, approvals, declines, pricing and outcomes by reviewer and reason code. Overrides can correct missing context; they can also introduce inconsistency and favour applicants who can provide additional documentation. Outcome tracking should show whether override groups perform within the approved range after accounting for selection.

The capital provider should receive a concise record of findings, limitations and remediation. Material unresolved issues can affect collections, complaints, partner continuity and funding confidence. Facility documents may require notification of significant findings while respecting personal-data restrictions and appropriate confidentiality.

Figure 4. Hypothetical cohort monitoring dashboard

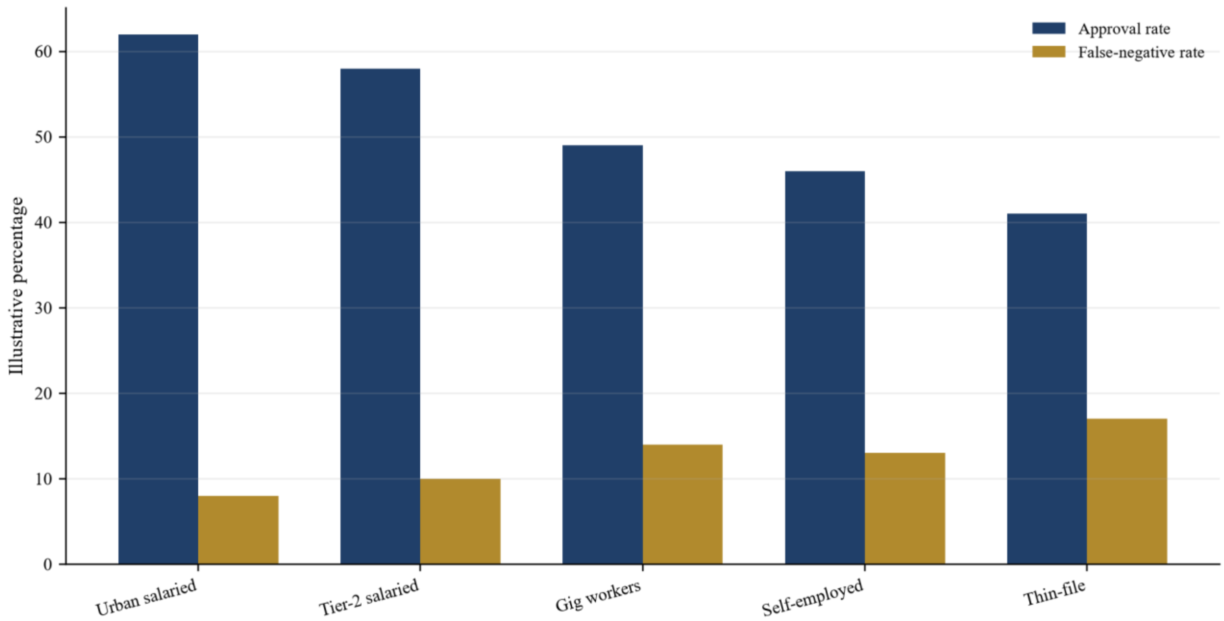

Approval and false-negative rates are modelling assumptions. Differences require investigation, context and legal review.

Table 5. Fairness and borrower-outcome test set

Test area	Measure	Diagnostic question	Potential action
data access	source coverage and missingness	does one cohort lose access to a critical feature?	alternative path or narrower feature use
decision	approval, limit, price and tenor	are outcomes materially different after relevant risk controls?	policy review and controlled retest
error	false-positive and false-negative rates	who bears incorrect approvals or declines?	threshold or review-zone adjustment
explanation	reason-code frequency and stability	can the lender give a consistent material reason?	simplify model or improve reason mapping
human review	referral, override and turnaround	does review correct or amplify disparity?	reviewer training and dual control
borrower outcome	delinquency, complaints, cures and hardship	do adverse outcomes concentrate after approval?	product, servicing or affordability change

Metric selection and lawful testing design require the regulated entity's approved governance process.

7. Make explanations usable for decisions and challenges

Explainability serves several audiences. Borrowers need clear and actionable reasons. Credit officers need the main drivers, data quality and policy rule. Validators need model behaviour, sensitivity, stability and limitations. Capital providers need evidence that a funded receivable followed an approved and reproducible process. Each audience requires a different depth and format.

Separate global and local explanations. Global evidence describes overall model behaviour, feature importance, interactions and stability. Local evidence describes the material factors for one decision. Use both with caution. Post-hoc methods can approximate a complex model and may change under small perturbations. Compare explanation results with a simpler challenger and known policy logic.

The decision packet should preserve input snapshot, missing-data path, feature values, model and policy versions, score, calibrated risk, rule outcome, reason codes, explanation, referral status, reviewer action and approval terms. The record should be immutable or strongly controlled. A capital provider can sample packets and reproduce eligibility without receiving more personal data than needed.

Reason codes need borrower-facing language and internal precision. A broad phrase such as insufficient cash flow provides limited control value. A structured code can identify volatility, debt burden, incomplete coverage, concentration, liquidity floor or inconsistent records. The lender should test whether the code reflects the actual decision logic and whether borrowers receive the explanation required by applicable rules.

Generative AI can summarise source documents or draft an explanation for review. It should remain outside deterministic approval logic unless separately governed, validated and approved. Prompts, versions, grounding data and reviewer action need records. Any output used in a consequential decision should be reproducible and attributable.

8. Design override authority around consequence

Human intervention should have a defined purpose. Review zones may cover model uncertainty, conflicting data, suspected fraud, missing records, material exposure, policy exceptions or new borrower types. Each zone needs permitted actions, evidence requirements, delegated limits and escalation. A reviewer should know whether they can add evidence, change a classification, alter terms or approve outside policy.

Capture overrides in structured form. Record the original recommendation, changed outcome, reason code, supporting evidence, reviewer, authority level, time stamp and any condition. Free text can supplement the record and should not replace the controlled fields. Link the override to repayment outcomes so the lender can test whether the intervention improves decisions.

Monitor direction and frequency. Approval overrides and decline overrides create different risks. A sudden rise in approvals can indicate commercial pressure or model drift. Concentration by reviewer, channel or partner deserves investigation. Track repeat overrides of the same rule; they may show that policy needs revision or that the review process is bypassing a necessary control.

Set quantitative and qualitative thresholds. A high rate alone does not prove weak governance because portfolio complexity can change. Thresholds should use a baseline, confidence range, materiality and minimum sample. Breaches can require enhanced review, independent validation, reduced eligibility or a pause in the affected cohort.

The board or delegated credit committee should approve the override framework. Independent validation should analyse outcomes. Internal audit should test authority, evidence and record integrity. The capital provider should receive aggregated reporting and targeted access rights appropriate to the financing structure.

9. Convert model uncertainty into debt capacity

Debt capacity follows cash available for debt service after losses, servicing cost, operating requirements, taxes, reserves and permitted distributions. A model influences that cash through borrower selection, loan size, price, tenor and collections. The financing model should trace each link instead of treating model accuracy as a direct source of leverage.

Translate uncertainty into scenarios. Vary approval mix, realised default, recovery timing, prepayment, data coverage, override rate, funding cost and collection lag. Correlated stresses matter. A weak channel may experience higher defaults and slower recoveries while the facility also becomes more expensive. The scenario should show liquidity, availability and covenant headroom through time.

Model benefit requires a conservative baseline. Compare the champion with a simpler approved challenger using contemporaneous cohorts where practical. Deduct implementation cost, human-review cost, vendor fees, remediation and transition effects. Treat unseasoned approval lift as prospective evidence until realised cash confirms it.

Apply an evidence haircut to benefits that depend on fragile data, incomplete validation or uncertain transferability. A capital provider can recognise stronger evidence through broader eligibility or a smaller reserve. Weak evidence can lead to a separate pool, concentration cap or exclusion. The financing response should be stated in advance.

Keep assumptions visible. Management forecasts should identify which inputs are observed, calculated, contractual or estimated. Sensitivities should vary one driver and combined scenarios should vary correlated drivers. Every output remains a scenario rather than a prediction of a particular borrower or portfolio.

Table 6. Model-risk findings translated into facility terms

Finding	Cash or risk effect	Potential facility control	Release evidence
weak recent-vintage calibration	loss reserve may be understated	lower advance rate and higher dynamic reserve	two matured cohorts within approved tolerance
data-feed concentration	origination and decision continuity risk	provider concentration cap and tested fallback	successful continuity drill and challenger evidence
missing decision lineage	eligibility and audit uncertainty	exclude affected receivables	reproducible sampled decisions
unstable group outcome	conduct, remediation and loss risk	enhanced reporting and restricted model change	independent review and approved remediation
high override concentration	inconsistent policy execution	manual-review cap and reviewer monitoring	stable override performance by cohort
unexplained cash lag	liquidity and commingling risk	controlled account, sweep and reconciliation trigger	reconciled collections within stated timing

Terms are illustrative and must be negotiated for the actual instrument and legal structure.

Figure 5. Hypothetical debt-service coverage under loss and funding-cost stress

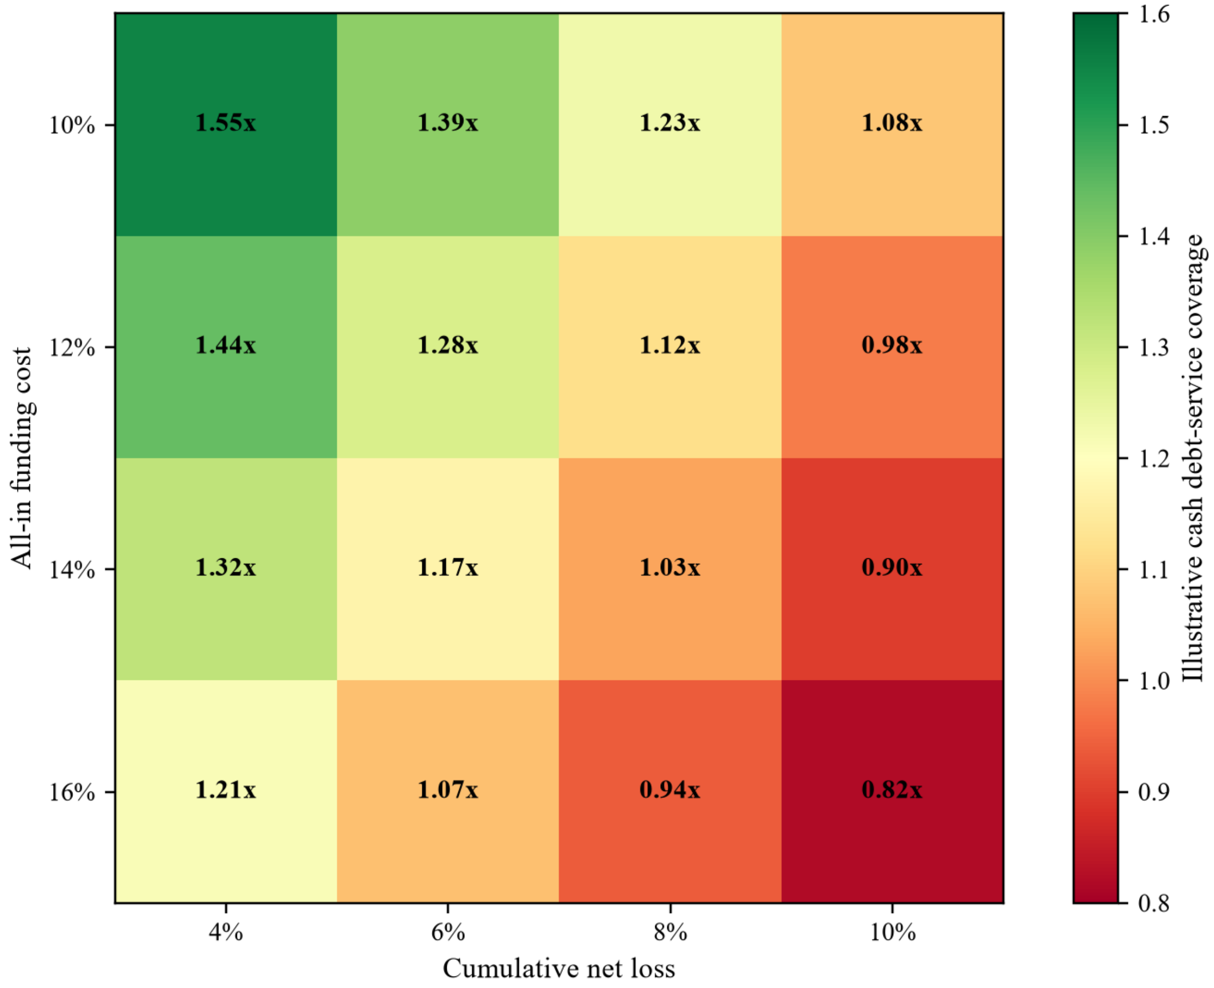

Ratios are modelling assumptions for an illustrative facility. They are not forecasts or proposed terms.

10. Build a facility control package around model evidence

Eligibility criteria should link legal asset status and underwriting evidence. A funded receivable can require verified borrower identity, executed agreement, disbursement, current payment status, approved product, permitted jurisdiction, decision-time data snapshot, model version, policy version and any override record. Missing critical evidence should exclude the asset until repaired.

Advance rates and reserves should respond to observed cash risk. Inputs can include cumulative net loss, recovery timing, delinquency development, prepayment, servicing cost, cash volatility, data coverage and concentration. A dynamic reserve can increase when recent vintages develop outside tolerance. The formula should avoid double counting between exclusions, haircuts and reserves.

Concentration limits can address product, partner, geography, borrower type, model version, data path and policy version. A new embedded partner or material model update can receive a temporary limit. The limit can step up after minimum seasoning, successful validation and operational evidence. This links growth to verified performance.

Model-change controls should define materiality. Changes to target, features, algorithm, calibration, cut-off, product, vendor or data source can require notice, validation and consent before new assets qualify. Routine maintenance can follow a documented lower-risk process. Emergency changes need bounded authority and retrospective review.

Triggers should have objective definitions, data sources, calculation owners and remedies. Examples include loss development, expected-to-actual ratio, data coverage, unexplained overrides, reporting failure and cash-reconciliation breaks. Remedies can include added reserve, reduced advance, cohort exclusion, cash sweep or controlled amortisation. The legal documents should reflect qualified advice and negotiated terms.

11. Monitor cash, cohorts, models and overrides together

Monthly funding reports should reconcile collateral, collections and covenant calculations to controlled records. Provide opening eligible balance, additions, removals, collections, losses, recoveries, reserves, utilisation, availability and exceptions. The report should identify the model, policy and data path attached to each funded cohort.

Portfolio monitoring should cover origination, approval, ticket, tenor, yield, delinquency, cure, recovery, prepayment, fraud, complaints and net loss. Model monitoring should cover discrimination, calibration, stability, feature availability, missingness, reason codes and override outcomes. Data monitoring should cover schema, timeliness, completeness, duplication, reconciliation and incidents.

Link these measures through common identifiers and reporting dates. A rise in delinquency can be traced to a partner, policy version, data outage or model segment. A change in approval rate can be traced to applicant mix, missing data or revised cut-offs. The capital provider gains an operational diagnosis rather than a disconnected set of dashboards.

Define escalation and reopening rules. A material breach should identify the accountable owner, immediate containment, affected population, cash effect, funding response and remediation evidence. The lender should preserve an incident record and show when the issue is closed. Repeated breaches can indicate a design problem requiring facility amendment or portfolio run-off.

Independent validation should report limitations, unresolved findings and remediation dates. Internal audit can test governance, model inventory, access, change records and evidence integrity. External specialists may review higher-risk systems or data suppliers. Reporting to the capital provider can preserve confidentiality through aggregated metrics and controlled access to sampled evidence.

Table 7. Monthly funding and model-risk report

Domain	Monthly measure	Escalation example	Accountable owner
collateral	eligible balance, exclusions, advance and deficiency	unexplained ineligible growth	finance and collateral agent
cash	collections, timing, reconciliation and waterfall	delayed or unreconciled cash	treasury and servicer
credit	vintage loss, roll rate, cure and recovery	recent cohort outside tolerance	chief risk officer
model	calibration, stability, challenger and override	material drift or challenger underperformance	model owner and validation
data	availability, coverage, quality and incidents	critical feed breach or schema change	data owner and technology
borrower	complaints, reason codes, hardship and redress	recurring issue concentrated by cohort	compliance and operations

Measures should be defined in the finance documents and operating procedures.

12. Work through a hypothetical Indian fintech-debt facility

Consider a hypothetical Indian non-bank financial company that provides short-duration working-capital loans to small merchants and service businesses. It originates through its own application and selected embedded-finance partners. The underwriting system combines bureau data with consented bank transactions, GST records, invoices, payment receipts and prior repayment behaviour. The company seeks a secured revolving facility from an institutional lender. Every figure in this section is a modelling assumption used only to demonstrate the framework.

The proposed commitment is INR 4.0 billion. The opening receivables pool is INR 3.1 billion and the requested advance rate is 78 percent. Average original tenor is nine months. The lender reports a contractual portfolio yield of 22 percent, servicing and collection cost of 3.5 percent and an all-in facility cost of 12.5 percent. Mature-vintage cumulative net loss ranges from 4.8 to 6.6 percent. Recent cohorts acquired through two embedded partners are seasoning above the earlier delinquency curve.

The cash-flow model estimates twelve-month default probability and a separate liquidity score. Independent validation confirms useful ranking on the frozen holdout, while calibration weakens for businesses with less than six months of connected bank history. The expected-to-actual default ratio rises from 1.01 for the validation population to 1.24 for that cohort under the assumed calculation. GST-feed completion falls from 69 percent in model development to 51 percent in the latest quarter. A portfolio-wide advance rate would conceal both changes.

The credit team reviews overrides. Twenty-two percent of applications in the newest partner channel received manual action; eight percent changed the model recommendation. Most changes cite verified invoices or seasonal contracts. A material minority uses free-text reasons that cannot be grouped or back-tested. The lender therefore introduces controlled reason codes, evidence attachments, delegated limits and monthly outcome analysis. Overrides without the required record cease to qualify for funding until remediated.

The agreed opening structure uses a 70 percent advance rate, a dynamic loss reserve and a 12 percent concentration cap for the affected channel. Receivables need a decision-time data snapshot, an approved model and policy version, a recorded explanation and any override evidence. A trigger applies if the three-month expected-to-actual ratio exceeds 1.15 after minimum seasoning, if a critical data feed falls below its approved coverage floor, or if unexplained overrides exceed the negotiated tolerance. Material changes to features, models, cut-offs, products or vendors require notice and defined consent before new assets qualify.

The downside case assumes cumulative net loss of 8.5 percent, a two-month collection slowdown, lower origination and a facility cost of 14.5 percent. Availability contracts through the dynamic reserve and concentration limits. If the assumed cash-coverage floor is breached, the waterfall shifts to cash sweep and new receivable funding stops. The remedial case restores eligibility only after sampled decisions reproduce, the affected cohort returns inside its threshold and independent validation accepts the change.

This example shows how model governance enters debt capacity. Ranking power supports screening. Calibration, cash conversion, data coverage, override quality and controlled collections determine whether receivables support leverage. The facility rewards evidence that survives operational stress.

13. Establish implementation and independent assurance

Implementation should use controlled gates. The first gate confirms legal structure, regulated perimeter, asset ownership, security, cash accounts and servicing continuity. The second gate reconciles the portfolio tape, ledger, collections and mature vintages. The third gate reviews data provenance, models, policy, explanations, overrides, fairness and fallback. The fourth gate translates findings into facility terms and reporting.

Run a complete reporting cycle before first funding. Load a collateral file, calculate eligibility, apply concentrations and reserves, reconcile collections, produce covenants and sample decision packets. Simulate a missing Account Aggregator feed, a model change and an override spike. Record each exception, decision owner and closure requirement.

Assign decision rights. The board or delegated committee approves risk appetite and AI policy. Credit owns lending decisions and policy. Model development builds the system. Independent validation challenges it. Data and technology own lineage and resilience. Compliance and legal oversee conduct and rights. Finance reconciles cash and reporting. Internal audit tests the complete system.

Create a quarterly assurance pack combining facility, portfolio, model, data, override, borrower and vendor evidence. Material changes return to the relevant gate. A new product, geography, partner or data source can require fresh validation and a revised financing decision. The pack should distinguish observed results, calculated metrics, management assumptions and unresolved findings.

The operating architecture should preserve separation of duties. Production access, feature approval, model release, policy deployment and facility reporting should not sit with one individual or one uncontrolled technical account. Privileged activity needs logging and periodic review. Emergency access should expire automatically and leave a reviewable record. These controls support the integrity of the decision history used by validators, auditors and funders.

Version management should cover more than model code. Store the data schema, feature definitions, reference data, calibration, cut-offs, policy rules, explanation method, reason-code taxonomy and user interface presented to the reviewer. A decision can change when any of these components changes. A release register should identify the approved package, deployment date, validation evidence, affected population and rollback plan.

Operational capacity requires testing during rapid growth. Review queues, manual evidence checks, complaint handling and data-quality investigations may expand faster than headcount. The lender should measure review ageing, escalation time, rework, abandoned applications and control exceptions. Funding growth can be capped when the operating team cannot complete required reviews within the approved service level.

Assurance findings need economic prioritisation. Classify each issue by affected receivable balance, potential cash impact, borrower consequence, likelihood, detectability and time to remediate. Link the classification to funding treatment. A documentation issue with complete underlying evidence can follow a timed cure; a missing decision snapshot or uncontrolled model release may require immediate exclusion. Closure should require evidence reviewed by an independent owner.

The capital provider should test information rights before closing. Monthly files, sampled decision packets, validation reports, incident notices and audit findings need agreed formats, timing and confidentiality protections. A dry run can reveal data fields that cannot be delivered or metrics whose definitions differ across finance, credit and technology. Resolving these gaps before funding reduces later disputes and makes covenant action executable.

Figure 6. Proposed ninety-day financing and model-assurance roadmap

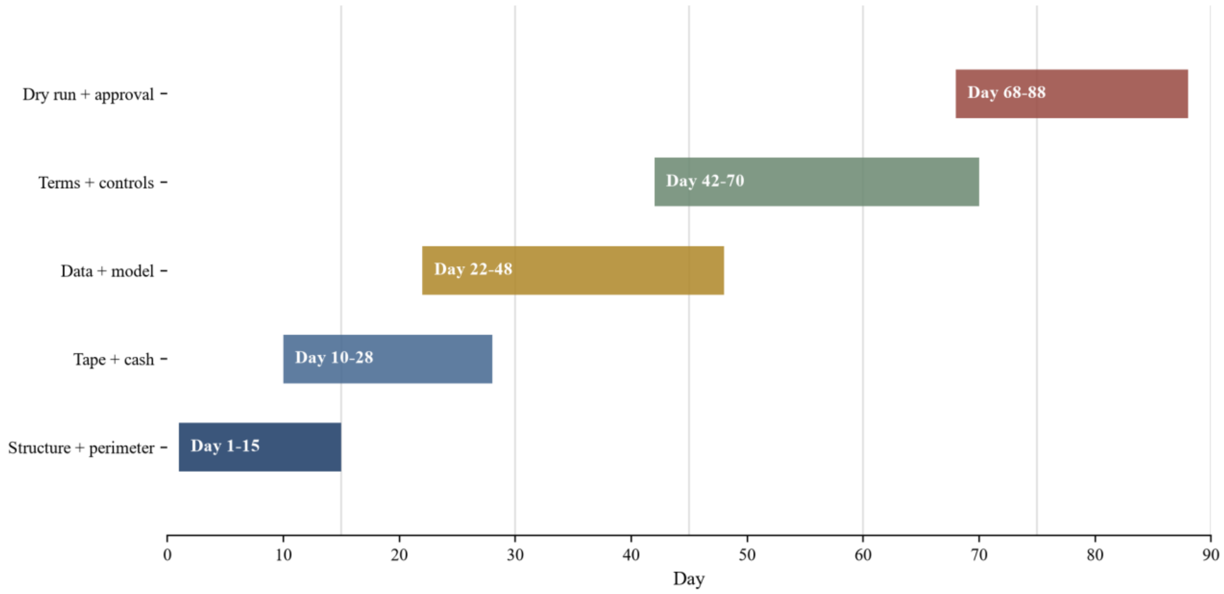

Sequence should be adapted to instrument, portfolio maturity and regulatory requirements.

14. Limitations and further research

This framework does not determine whether a particular data field may be collected or used, whether a lending product complies with applicable rules, or whether a security interest is enforceable. Those conclusions depend on current facts and qualified professional advice. Regulatory obligations, commencement dates and supervisory expectations can change.

Observed model performance does not establish causality. Marketing, pricing, fraud controls, servicing, economic conditions and borrower selection can change repayment. Connected-data coverage can reflect self-selection. Outcomes are unavailable for declined applicants, which complicates reject inference and fairness analysis.

Explainability methods have limits. A post-hoc explanation may approximate a model without reproducing its internal reasoning. Feature importance can vary by sample and method. A stable explanation does not prove correct data, fair treatment or suitable policy. Validation should connect explanation tests to the specific decision and audience.

The hypothetical facility contains no market forecast or recommended term. Advance rates, reserves, triggers, prices, loss assumptions and thresholds should come from the actual portfolio, legal structure, mandate and negotiated documents. Further research could compare the stability of Account Aggregator, GST, invoice and payment features across Indian MSME cohorts; evaluate privacy-preserving validation; and measure whether structured overrides improve realised cash outcomes.

15. Conclusion

AI cash-flow underwriting can support Indian fintech debt when the lender proves a continuous evidence chain from permitted source data to realised collections. The capital provider should be able to reproduce a sampled decision, identify its material drivers, verify the policy rule, inspect any override and trace the funded receivable to controlled cash.

Explainability supports borrower communication, credit challenge, validation and funding diligence. Override governance supports accountability when data are incomplete or the model is uncertain. Both controls become economically relevant through eligibility, reserves, concentration limits, model-change rights, reporting triggers and remedies.

The financing decision should follow observed cash, mature-vintage performance, calibration, data reliability and operational control. Strong evidence can support broader eligibility and more efficient capital. Weak or irreproducible evidence should narrow funding until remediation is verified.

Frequently asked questions

Q: What is AI cash-flow underwriting? A: It uses models to evaluate transaction, GST, invoice, payment and related financial data as evidence of repayment capacity. The lender still needs an approved credit policy, documented data rights, decision-time lineage, validation and accountable human governance.

Q: Which cash-flow data can support an Indian fintech credit decision? A: Potential sources include consented bank transactions, Account Aggregator feeds, GST records, invoices, payment receipts, bureau data and prior repayment history. Permitted use, reliability and relevance require case-specific assessment.

Q: What should an explanation contain? A: A useful explanation identifies the material decision drivers, applicable policy rule, data limitations and available review path in language suited to its audience. Internal records should preserve the source snapshot, versions, score, policy outcome and reason codes.

Q: When can a credit officer override the model? A: The lender should define review zones, permitted actions, evidence requirements, delegated authority and escalation. Every override should record the original recommendation, changed outcome, reason, evidence, reviewer and time stamp.

Q: How should override quality be tested? A: Compare frequency, direction, reviewer concentration, reason-code use and realised outcomes. Investigate repeated changes to the same rule, unexplained free text and approval overrides that develop outside the approved loss range.

Q: How does model drift affect a debt facility? A: Drift can change calibration, approvals and realised losses. Facility responses can include added reserves, reduced advance rates, cohort limits, enhanced reporting, new-asset suspension or controlled amortisation under negotiated terms.

Q: Can a strong model justify a higher advance rate by itself? A: Debt capacity also depends on asset ownership, cash control, servicing, mature losses, recoveries, concentration, data continuity and liquidity. The advance rate should reflect the combined evidence and downside case.

Q: What should happen when a critical data feed fails? A: The lender should use a tested fallback, narrow limits or route cases to an approved manual path. New receivables should qualify only when the fallback meets the facility's evidence standard.

Appendix A. Due-Diligence Data Request

A1. Portfolio and cash

- Loan-level origination, payment, delinquency, cure, recovery, write-off, prepayment and cash-collection records.
- Product, channel, geography, acquisition, risk band, model version, policy version and servicing fields.
- General-ledger, bank-account, trust, escrow and investor-report reconciliations.
- Complaints, bureau corrections, hardship, fraud and recovery-dispute records linked through controlled identifiers.

A2. Data model and governance

- Data inventory, source contracts, permissions, schema versions, quality reports and incident history.

- Feature definitions, training snapshots, code, validation, challenger, calibration and stability results.
- Credit-policy versions, cut-offs, pricing, limits, affordability rules, overrides and approval minutes.
- AI inventory, board policy, model owner, independent validation, audit findings, remediation and fallback tests.

Appendix B. Facility Control Checklist

B1. Before first funding

- Confirm regulated perimeter, asset ownership, security, cash accounts and servicing continuity.
- Reproduce sampled decisions using decision-time data, feature logic, model and policy versions.
- Reconcile eligible balances and collections to source records and bank cash.
- Approve advance rate, reserves, concentrations, model-change controls, triggers and reporting definitions.

B2. During the facility

- Reconcile collateral and cash every reporting period.
- Review mature and developing vintages, calibration, stability, data coverage and overrides.
- Escalate material complaints, incidents, vendor failures, regulatory findings and model changes.
- Test fallback, servicing transfer and enforcement readiness under the negotiated documents.

References

[1] Reserve Bank of India. Reserve Bank of India Digital Lending Directions, 2025 and related official material. https://www.rbi.org.in/

[2] Reserve Bank of India. Handbook on Regulations at a Glance, February 2025. https://website.rbi.org.in/documents/d/rbi/handbookg27022025d0f3f53f5d3c4310a6bb2f8ac2175d3a

[3] Reserve Bank of India. Master Direction - Reserve Bank of India Credit Information Reporting Directions, 2025. https://www.rbi.org.in/Scripts/BS_ViewMasDirections.aspx?id=12764

[4] Government of India. Digital Personal Data Protection Act, 2023. https://www.indiacode.nic.in/indiacode/handle/123456789/22037

[5] Ministry of Electronics and Information Technology. Digital Personal Data Protection Rules, 2025. https://www.meity.gov.in/documents/act-and-policies/digital-personal-data-protection-rules-2025-gDOxUjMtQWa

[6] Ministry of Electronics and Information Technology. Enforcement Timeline for the Digital Personal Data Protection Act, 14 November 2025. https://www.meity.gov.in/documents/act-and-policies/digital-personal-data-protection-rules-2025-gDOxUjMtQWa

[7] Reserve Bank of India. FREE-AI Committee Report - Framework for Responsible and Ethical Enablement of Artificial Intelligence, 13 August 2025. https://m.rbi.org.in/Scripts/PublicationReportDetails.aspx?ID=1306&UrlPage=

[8] Financial Stability Institute. In Data We Trust? Emerging Policy and Supervisory Approaches to AI Data Use in Financial Services, FSI Insights 73, 26 March 2026. https://www.bis.org/publications/fsi-insight-73-data-we-trust-emerging-policy-and-supervisory-approaches-ai-data-use-financial-services

[9] Reserve Bank of India. Annual Report 2024-25, Regulation, Supervision and Financial Stability. https://www.rbi.org.in/scripts/AnnualReportPublications.aspx?Id=1436

[10] Reserve Bank of India. Annual Report 2023-24, Public Tech Platform for Frictionless Credit and Account Aggregator developments. https://www.rbi.org.in/scripts/AnnualReportPublications.aspx?Id=1406

[11] Financial Stability Institute. Regulating AI in the Financial Sector: Recent Developments and Main Challenges, FSI Insights 63, 12 December 2024. https://www.bis.org/publications/fsi-insight-63-regulating-ai-financial-sector-recent-developments-and-main-challenges

[12] Bank for International Settlements. Artificial Intelligence and the Economy: Implications for Central Banks, Annual Economic Report 2024. https://www.bis.org/publications/aer-2024/artificial-intelligence-economy-implications-central-banks

[13] Reserve Bank of India. Financial Awareness Messages, Account Aggregator framework and cash-flow-based lending. https://www.rbi.org.in/commonperson/images/FAME202426022024.pdf

[14] Reserve Bank of India. Guidelines on Digital Lending, 2 September 2022. https://www.rbi.org.in/Scripts/NotificationUser.aspx?Id=12382&Mode=0

[15] Reserve Bank of India. Key Facts Statement for Loans and Advances, 15 April 2024. https://www.rbi.org.in/Scripts/NotificationUser.aspx?Id=12648&Mode=0

[16] Reserve Bank of India. Default Loss Guarantee in Digital Lending, 8 June 2023, updated through official FAQs. https://www.rbi.org.in/Scripts/NotificationUser.aspx?Id=12535&Mode=0

[17] Reserve Bank of India. Master Direction on Outsourcing of Information Technology Services, 10 April 2023. https://www.rbi.org.in/Scripts/NotificationUser.aspx?Id=12486&Mode=0

[18] Reserve Bank of India. Master Direction on Information Technology Governance, Risk, Controls and Assurance Practices, 7 November 2023. https://www.rbi.org.in/Scripts/NotificationUser.aspx?Id=12562&Mode=0

[19] National Institute of Standards and Technology. Artificial Intelligence Risk Management Framework 1.0. https://doi.org/10.6028/NIST.AI.100-1

[20] Organisation for Economic Co-operation and Development. OECD Principles on Artificial Intelligence. https://oecd.ai/en/ai-principles

[21] Basel Committee on Banking Supervision. Principles for the Management of Credit Risk. https://www.bis.org/bcbs/publ/d591.htm

[22] International Organization for Standardization. ISO IEC 23894 Artificial Intelligence Guidance on Risk Management. https://www.iso.org/standard/77304.html

About the Author

Chennakeshav (CK) is a corporate finance and investment banking executive with 25+ years of global experience in deal origination, structuring and execution across M&A, growth capital and corporate strategy. He has led value-creation mandates for founders, corporates and funds — bridging the boardroom view to hands-on execution and closure.

His career spans Morgan Stanley, HSBC, Lloyds Banking Group, EWEC, ADQ portfolio companies and Emirates Growth Fund, across TMT, real estate, fintech, deeptech, cleantech, infrastructure and energy. He has partnered with C-suite leaders, private equity and venture funds, sovereign wealth funds and family offices to finance complex fund raises and scale-up ventures, and has led M&A due diligence, post-merger integration and business-transformation initiatives to create value.

At Matchpoint Partners, he is Managing Partner and leads the firm's corporate finance, M&A and capital-raising practice. He holds an MBA from London Business School, an engineering degree from VTU and a Master of Laws (LLM, in progress) from UCL London.

An active start-up mentor, CK mentors at Techstars, DIFC FinTech Hive, Startup Grind, Founder Institute and IN5, serves as Entrepreneur Mentor in Residence (EMiR) at London Business School, and judges the Entrepreneurship World Cup.

https://www.linkedin.com/in/ckadya/

https://www.matchpoint-partners.com/team/ck-adya.html

This paper is part of a continuing series on the structure of private and alternative markets. The views expressed are the author's own. The paper is for information only, describes market structure in general terms, and does not constitute investment, legal, tax or regulatory advice or a recommendation in respect of any security, vehicle or counterparty.