India Consumer Private Credit: Alternative-Data Underwriting without Model-Risk Blind Spots

An Evidence-Gated Framework for Data Provenance, Fairness, Vintage Risk and Funding Covenants

Authored by Chennakeshav Adya
Independent Researcher
September 2026

Abstract

Indian consumer lenders increasingly combine bureau information with transaction, bank-account, payment, device, employment, commerce and distribution data. The resulting models may widen access and sharpen risk selection, especially for borrowers with limited formal credit history. They can also create blind spots when a lender cannot prove where a feature came from, whether it was available at the decision date, how it behaves across borrower groups, or what happens when the data provider, model or channel changes.

This paper develops an evidence-gated framework for a private-credit fund, bank, family office or institutional capital provider evaluating an Indian consumer-credit platform, non-bank financial company or financing vehicle. It treats underwriting as a chain that runs from lawful data acquisition through feature construction, model development, credit policy, human review, loan servicing, collections, portfolio performance and cash available for debt service. Each link receives an owner, definition, evidence source, control, failure signal and financing consequence.

The framework separates predictive performance from decision quality. It requires time-based validation, challenger models, out-of-time and out-of-distribution tests, cohort and vintage analysis, calibration, stability monitoring, fairness tests, reason codes, override controls and a documented fallback when a model or data feed fails. The capital structure then translates those findings into eligibility criteria, advance rates, reserves, concentration limits, performance triggers, reporting rights and remedies. Regulatory and consumer-protection conclusions remain within the remit of qualified Indian counsel and the relevant regulated entity.

A hypothetical case illustrates a secured facility to an Indian consumer lender. All loan balances, loss rates, model metrics, thresholds, advance rates, prices, time periods and transaction terms are modelling assumptions. They do not describe an identified lender, borrower, client, portfolio or investment recommendation. The paper concludes that alternative data becomes financeable when the lender can reproduce a decision, explain its material drivers, demonstrate lawful and representative use, show stable cash performance by cohort, and operate a credible human and technical control path when the model is wrong.

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

Keywords: India consumer credit, private credit, alternative data, underwriting, model risk, digital lending, NBFC, fairness, human oversight, portfolio covenants

1. Define the financing decision and the repayment source

The first question is not whether the lender uses artificial intelligence. The capital provider needs to know what obligation it is funding, which cash flows repay it, which entity controls those cash flows, and how underwriting affects the probability and timing of repayment. A warehouse line to a non-bank financial company, a bilateral term facility, a securitisation investment, a forward-flow arrangement and corporate debt expose the investor to different assets, controls and remedies. The diligence plan should begin with the legal and economic structure.

Map the regulated lender, lending-service providers, digital applications, originators, servicers, collection agencies, data providers, account aggregator connections, credit information companies, payment rails, trust or special-purpose vehicle, bank accounts and security package. Record which party approves credit, owns the receivable, holds customer data, communicates with the borrower, collects cash and reports performance. Outsourcing does not remove the regulated entity's responsibility under RBI's digital-lending framework.[1][2] A capital provider should therefore test the regulated entity's control over every critical outsourced step.

Define the repayment waterfall using observed contracts and cash records. Loan collections may pass through operating accounts, escrow accounts, trusts or collection accounts before reaching the funder. Identify set-off rights, commingling risk, servicing fees, taxes, refunds, chargebacks, customer redress, first-loss support, cash reserves and senior claims. Reconcile the contractual waterfall with bank statements and the servicing system. A model can rank borrowers accurately while weak cash control still impairs debt service.

State the decision that each analytical test supports. The investment committee may need to decide whether to lend, how much to advance, which receivables qualify, what reserve to hold, what concentration to allow, what trigger should reduce availability, and when a model or data change requires consent. Those decisions need different evidence. A one-page model-accuracy presentation cannot replace a portfolio tape, vintage analysis, data-lineage record, credit-policy history and cash reconciliation.

The output is a financing decision map. It links the proposed instrument to the receivable pool, underwriting system, servicing process, data rights, covenants, reporting package and remedies. It also identifies professional questions that require Indian legal, regulatory, tax, accounting, data-protection or consumer-protection advice.

Figure 1. From borrower data to private-credit repayment

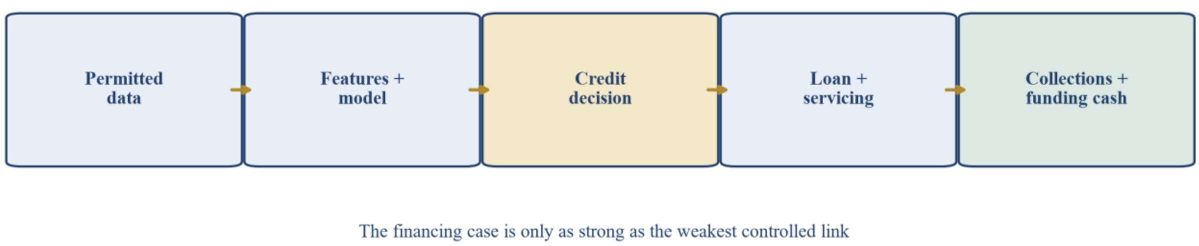

Management framework. Each arrow requires an evidence source, owner, control and financing consequence.

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

Decision	Evidence required	Failure signal	Financing response
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
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 alternative data inside the Indian regulatory perimeter

Alternative data is a broad commercial label. The diligence team should classify every data element by source, purpose, legal basis, regulatory treatment, decision use and retention. Bureau data, bank-account transactions, goods and services tax filings, utility payments, payroll records, e-commerce activity, device signals, location, contacts and behavioural variables do not carry the same relevance or risk. Some sources can be necessary for underwriting; others may be intrusive, weakly related to repayment or difficult to explain.

The Reserve Bank of India consolidated its digital-lending requirements in the Reserve Bank of India Digital Lending Directions, 2025.[1] The framework addresses regulated entities, lending-service providers, digital lending applications, disclosures, data collection, disbursement, repayment, grievance handling, reporting and default-loss guarantees. The RBI handbook summarises requirements for need-based data collection, clear audit trails, prior explicit consent, borrower choice over specific data and restrictions on storage by service providers.[2] The exact application to a transaction requires current legal review.

The RBI Credit Information Reporting Directions, 2025 require credit institutions to submit borrower credit information to credit information companies and keep information updated on a fortnightly basis, subject to the direction's detailed provisions.[3] More current formal credit records can reduce dependence on weak proxies. A lender should show how bureau data are obtained, matched, refreshed, disputed and corrected. The capital provider should test whether reporting delays or identifier errors distort repeat-loan decisions and portfolio monitoring.

India's Digital Personal Data Protection Act, 2023 establishes the statutory framework for processing digital personal data.[4] The Digital Personal Data Protection Rules, 2025 and the notified commencement framework add implementation requirements and staged dates.[5][6] The lender should maintain a current applicability analysis, notices, consent or other lawful grounds as applicable, processor contracts, security safeguards, retention schedules, rights handling and breach response. A financing memorandum should not substitute for counsel's opinion on live obligations.

The RBI FREE-AI Committee's August 2025 report provides a detailed governance reference for artificial intelligence in financial services.[7] Its recommendations include board-approved policy, lifecycle governance, product approval, fairness, bias, understandability, customer protection, monitoring, human oversight and independent assurance. The report records that only a minority of surveyed users maintained some interpretability, audit-log, bias-mitigation and regular-audit practices at the time of the survey. That evidence supports close diligence of implementation rather than an assumption that industry adoption proves control maturity.

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
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
consumer 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. Build data provenance before building predictive power

A reliable model starts with a data contract. For every source, record the provider, system, borrower permission or other lawful basis as applicable, field definition, unit, time stamp, refresh frequency, coverage, missing-value treatment, validation rule, transformation, retention period and permitted use. Preserve the raw value, transformed feature, model version and decision outcome. The record should allow an independent reviewer to reproduce a historical decision using information that existed at that time.

Time integrity is central to credit validation. A feature may look predictive because it contains information created after approval, disbursement or delinquency. Examples include a revised bureau record, post-disbursement account balance, collections tag, later device event or manually corrected employer field. Freeze the decision-time snapshot and test feature availability. Any leakage should invalidate the relevant back-test and trigger a controlled rebuild.

Coverage deserves the same attention as accuracy. Bank-transaction features may exist only for applicants who complete a consented connection. Payroll variables may favour salaried workers. E-commerce histories can represent platform-active customers more strongly than cash-dependent borrowers. Device or language patterns can correlate with geography, income, age or other characteristics. The lender should publish coverage by cohort and explain what model or policy applies when a source is absent.

Data quality tests should include completeness, validity, duplication, range, consistency, timeliness, reconciliation and unexpected distribution shifts. Source-specific thresholds belong in an operational dashboard. A sudden rise in missing transaction categories, shorter account history or a changed bureau match rate can alter approval and loss outcomes before headline model performance shows the problem.

Third-party providers introduce contractual and operational dependencies. Review audit rights, data ownership, subcontractors, service levels, incident obligations, portability, termination assistance and the lender's ability to operate during an outage. The BIS Financial Stability Institute identifies data privacy, quality, security, third-party dependence and concentration as important concerns in AI adoption by financial institutions.[8] A capital provider should translate those dependencies into reporting and fallback requirements.

Figure 2. Alternative-data provenance and decision lineage

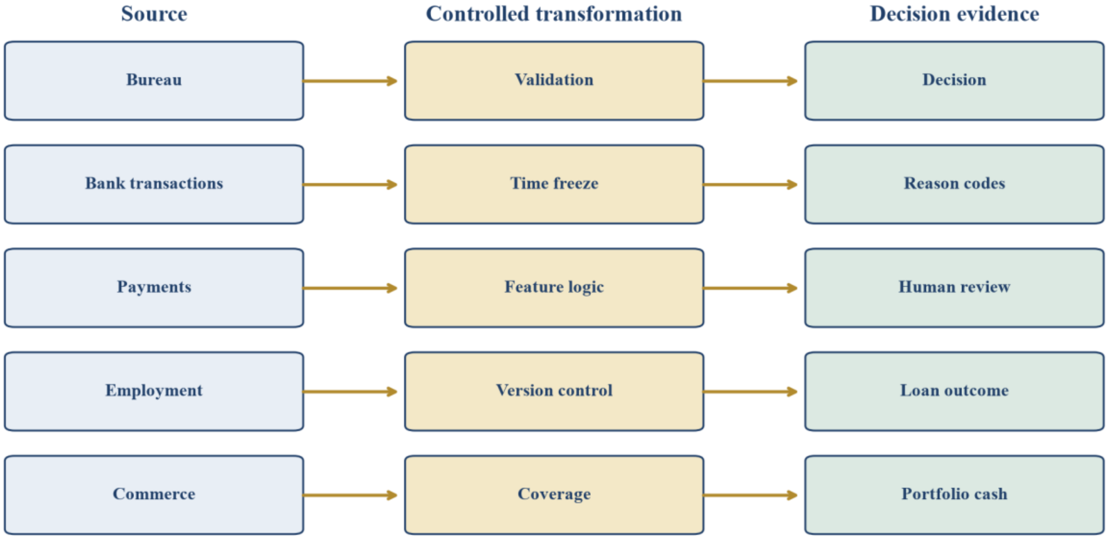

Proposed control architecture. Raw sources remain traceable through features, versions, decisions and outcomes.

Table 3. Minimum alternative-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 model development from credit policy

The statistical model estimates an outcome under defined data and assumptions. Credit policy converts that estimate into an approval, limit, price, tenor, documentation requirement or decline. The distinction matters because losses can arise from a sound ranking model paired with an aggressive cut-off, high loan size, weak affordability test or acquisition incentive. Diligence should reconstruct both layers.

Start with the target variable. Ninety-day delinquency, first-payment default, cumulative net loss, fraud, prepayment and collections cost answer different questions. Define the performance window and censoring treatment. Recent loans have not had time to reveal long-horizon loss, so compare matured vintages and show development curves. Exclude or separately identify restructurings, settlements, refunds, deceased borrowers, operational errors and fraud according to a documented policy.

Use a transparent benchmark before accepting a complex model. Logistic regression, scorecard or simple rule set can provide a challenger. Compare discrimination, calibration, stability and economic value. Area under the receiver-operating curve or Gini measures ranking, while calibration tests whether predicted probabilities correspond to observed outcomes. The financing case needs both. A model can rank well and still understate absolute loss, weakening advance rates and reserves.

Validation should use time-based splits that resemble live deployment. Randomly splitting records from the same economic period can overstate performance when applicant mix, acquisition channels and data vendors are stable across train and test samples. Hold out later vintages, new channels, regions and product changes. Stress missing data and degraded feeds. Test monotonicity and reason stability around cut-offs so small, immaterial input changes do not generate erratic decisions.

Document every policy overlay: minimum income, debt-service threshold, loan-size cap, geographic exclusion, employer rule, fraud flag, manual decline and repeat-borrower treatment. Quantify approvals and losses removed or added by each overlay. The committee can then distinguish value created by the model from value created by policy, pricing, servicing and channel selection.

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 vintages as cash-flow evidence

Private credit is repaid by cash, so model diligence must connect to realised portfolio behaviour. Build monthly or quarterly vintages using origination date and track scheduled principal, interest, fees, delinquency transitions, cures, prepayments, settlements, recoveries, write-offs, servicing costs and net cash. Reconcile the loan tape to the general ledger, bank accounts, regulatory reports and investor statements.

Segment vintages by product, tenor, ticket, risk band, state or region, acquisition channel, repeat status, model version, policy version, data-source coverage and servicing strategy. The objective is to find where the portfolio's economics differ from the headline average. A stable total loss rate can conceal a weakening new channel offset by a shrinking seasoned book. A favourable early-payment curve can coexist with later loss when tenors extend.

Use delinquency roll rates and cumulative net loss curves. First-payment default can identify fraud, affordability or origination-quality problems. Thirty-to-sixty and sixty-to-ninety transitions show servicing and borrower stress. Cure rates, recovery timing and collection expense affect liquidity even when lifetime loss is unchanged. Compare booked yield with cash yield after reversals, refunds, waivers, prepayment and collection cost.

The Reserve Bank has highlighted consumer-credit risk and previously increased risk weights for certain unsecured consumer exposures and bank exposures to non-bank financial companies.[9][10] The specific prudential treatment has changed over time and must be checked at the transaction date. The financing model should therefore use current regulation and observed portfolio evidence, with sensitivity to funding cost, capital requirements and risk appetite.

Vintage maturity controls prevent false comfort. State what percentage of each cohort has completed the relevant performance window. Use development factors only when supported by stable historical relationships, and show sensitivity when the book, model or economy changed. Recent growth should not be valued as mature performance.

Figure 3. Hypothetical cumulative net loss by origination vintage

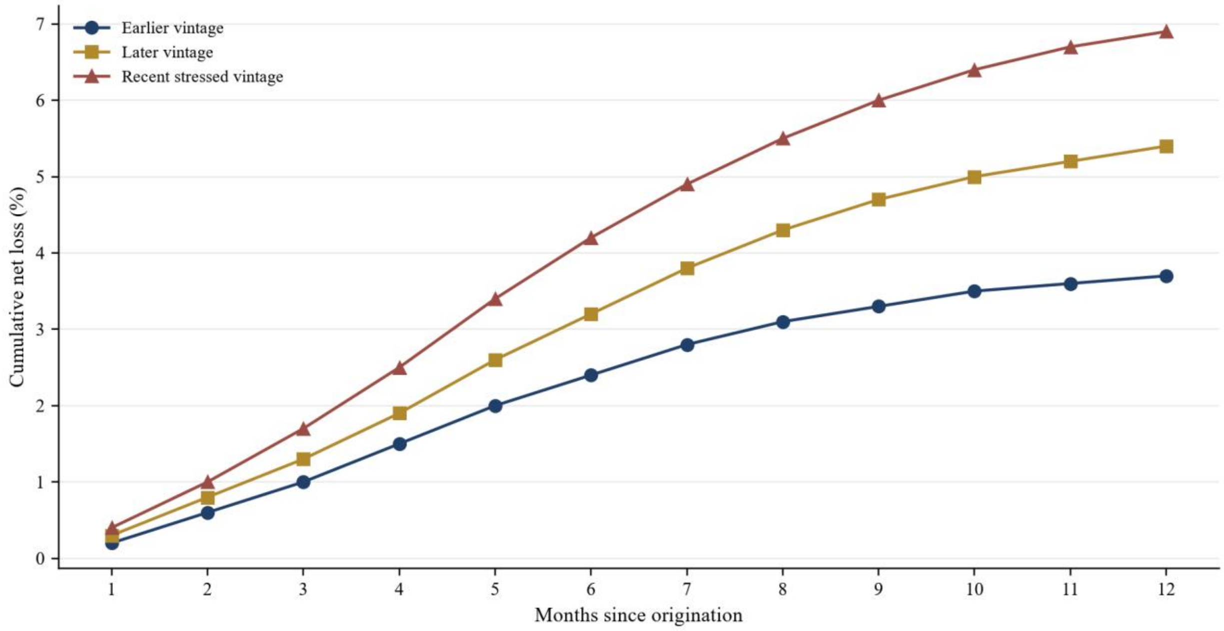

Percentages are modelling assumptions and do not describe an actual lender or portfolio.

6. Test fairness as a portfolio and control issue

Fairness testing asks whether a model or policy creates materially different outcomes across relevant groups and whether those differences can be justified by legitimate credit objectives, data quality and applicable law. The design of protected or sensitive-group testing requires qualified Indian legal and compliance input. The analytical team should avoid collecting sensitive attributes without an approved basis. It may need controlled, privacy-preserving or proxy-based testing under appropriate governance.

Compare application coverage, approval, limit, price, tenor, reason codes, manual review, false-positive and false-negative rates, delinquency, cure and complaints. No single fairness metric is universally correct. Demographic parity, equal opportunity, equalised odds, calibration and error-rate balance can conflict. Choose metrics based on the decision, harm, legal context and control objective, then document the choice.

Intersectional analysis matters because aggregate tests can hide concentrated effects. A model may appear stable by region and employment type separately while treating a smaller regional employment cohort differently. Minimum sample rules, confidence intervals and multiple-testing controls reduce false conclusions. Small cohorts should trigger cautious investigation rather than confident claims.

Alternative data can encode indirect proxies. Device price, operating system, language, transaction location, merchant mix, contact structure and digital behaviour may correlate with socio-economic or geographic factors. Remove variables that lack a credible repayment relationship or create unacceptable risk. Where a feature remains, document its economic rationale, incremental predictive value, sensitivity and explanation.

Fairness monitoring should continue after deployment. Changes in channel mix, marketing, data coverage and economic conditions can alter group outcomes even when model code stays constant. The FREE-AI report places fairness, accountability and understandability within the proposed lifecycle framework and calls for product-approval assessments that include data quality, sensitive-attribute exclusion, back-testing and independent review.[7] The funder can require evidence of this process through reporting and audit rights.

Figure 4. Hypothetical cohort monitoring dashboard

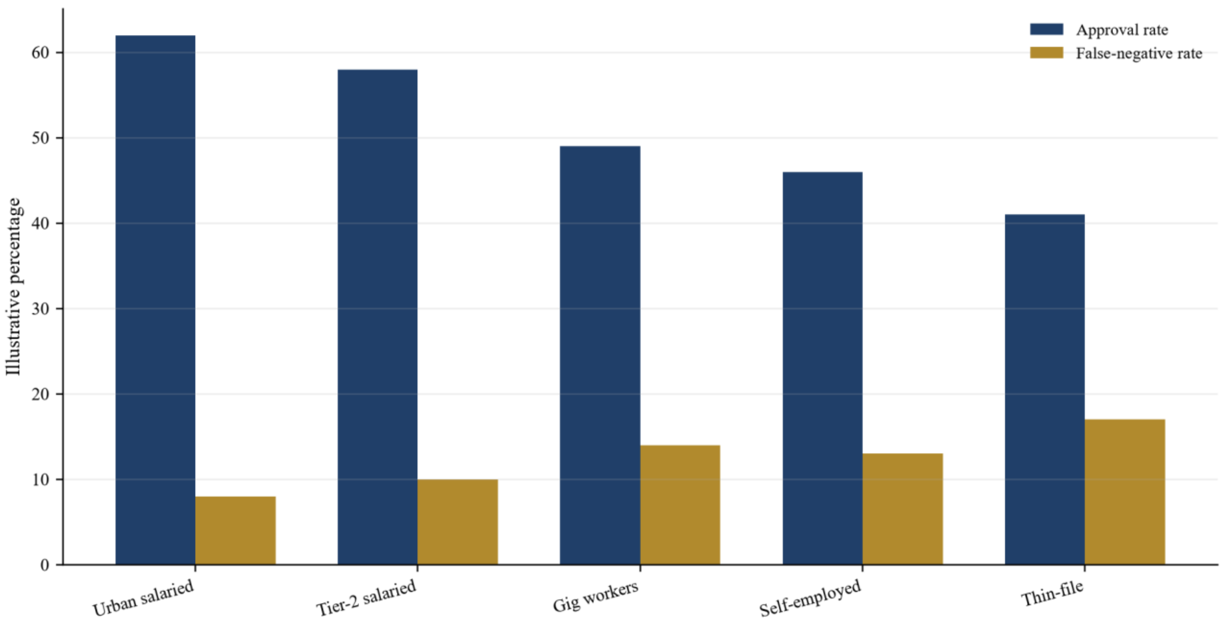

Approval and false-negative rates are modelling assumptions. Differences require investigation, context and legal review.

Table 5. Fairness and consumer-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
customer 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 operational

Explainability serves several audiences. The credit officer needs to understand the drivers and limits of the recommendation. The borrower needs clear information and recourse appropriate to the product and applicable requirements. The validator needs feature behaviour, sensitivity and reason mapping. The board and capital provider need aggregate evidence that the model operates within approved risk appetite.

Global explanations describe the model's overall structure and feature influence. Local explanations identify material factors in one decision. Both can mislead if they are unstable, correlated or detached from the credit policy. Compare explanation methods across representative records and around policy cut-offs. A reason code should correspond to an actionable and factually supported driver, rather than a generic label created for presentation.

Preserve a decision packet for each funded receivable: input snapshot, missing-data path, feature version, model version, score, probability, policy rules, reason codes, human action, approval terms and time stamp. The capital provider does not need unrestricted access to personal data. It needs sufficient contractual rights and privacy-preserving evidence to test the control environment, reproduce sampled decisions and verify portfolio eligibility.

Customer disputes are a valuable control signal. Link complaints, bureau corrections, identity errors, consent withdrawals where applicable, fraud cases, hardship and recovery disputes to the original decision and data sources. Trend recurring reasons by model, vendor, channel and cohort. A high complaint rate may show that a technically valid model operates through a poor product, interface or servicing process.

Generative AI should not invent approval reasons or make autonomous adverse decisions without approved controls. The BIS notes that generative models intensify explainability challenges because their complexity and variability make inputs and outputs harder to trace.[11] If a lender uses a language model to summarise documents or draft explanations, keep deterministic credit decision logic separate, ground outputs in preserved records, restrict permissions and require human review for consequential use.

8. Design human oversight around consequence

Human review should have a defined purpose. Sending every application to a person creates cost without necessarily improving control. Fully automating every case can amplify a bad feed, policy or model.

Define review zones using uncertainty, missing data, conflicting signals, material loan size, suspected fraud, vulnerable-customer indicators, novel cohorts and policy exceptions.

Reviewers need evidence, authority and time. Present the underlying fields, source dates, model range, reason codes, policy limits and comparable cases. Record the reviewer, decision, rationale, changed terms and supporting evidence. Prohibit overrides that exist only as free text. Categorise them so the lender can test whether human action improves or weakens portfolio performance.

Analyse overrides in both directions. Approval overrides can create adverse selection; decline overrides can reduce access and revenue. Compare override rates, performance and customer outcomes by reviewer, branch or digital channel, product and cohort. A reviewer who consistently overrides one group or reason code may reveal training, incentive or process problems. The control owner should investigate before changing the model.

The RBI FREE-AI Committee states that deploying entities remain accountable and recommends human oversight for medium- and high-risk uses, model documentation, validation, monitoring, drift detection and independent product approval.[7] A capital provider can reflect those principles in conditions precedent and ongoing information rights. Board approval should identify who owns the model, who validates it, who can change it, who can suspend it and who communicates with customers.

Fallback is part of oversight. If a bureau, account aggregator, cloud service or model endpoint fails, the lender should know whether to stop originations, use a challenger, narrow products, lower limits or route cases to manual policy. Test the fallback through an operational drill. A written continuity plan without evidence of execution provides weak protection.

9. Convert model uncertainty into credit economics

The lender's unit economics should start with cash received from borrowers and deduct funding cost, expected credit loss, fraud, servicing, collections, customer redress, payment cost, acquisition expense and operating overhead. Booked interest and fee income can differ from cash yield because of prepayment, delinquency, reversals, waivers, refunds and write-offs. Reconcile each bridge with historical records.

Translate model error into loss and liquidity ranges. If predicted default probability is miscalibrated, the resulting loss depends on exposure at default, recovery, timing and collection cost. A small error in a fast-growing, high-ticket or high-advance-rate segment can matter more than a larger error in a small seasoned segment. Weight validation by financing exposure and cash consequence.

Stress correlated changes. A funding-cost increase can coincide with weaker borrower cash flow, lower approval, slower collections and tighter lender liquidity. A data-provider outage can reduce originations and change the risk mix if the fallback approves a different population. A policy response can limit data use or require remediation. The financing case should show how these events affect collections, excess spread, borrowing-base availability and covenant headroom.

Price should not compensate for an unbounded control failure. Higher yield cannot make an unverifiable asset eligible. Set minimum evidence for inclusion, then use advance rate, reserve, subordination and price to absorb measurable risk. Exclude records with missing ownership, duplicate financing, irreconcilable balances, absent decision history or prohibited data use until cured.

The capital provider should maintain an independent base case. Management forecasts may assume continued model lift, benign funding, stable acquisition and improving collections. Rebuild the case using observed mature vintages, conservative development, current funding terms, implementation cost and realistic cash timing. Show which assumptions drive premium or leverage capacity.

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.

10. Build the borrowing base and covenant package

Eligibility criteria should be objective, observable and testable from the loan tape and legal documents. Typical fields include lender ownership, borrower verification, executed agreement, disbursement, product, jurisdiction, payment status, maximum delinquency, remaining term, loan size, concentration, policy version, model version and absence of duplicate financing. The actual criteria depend on transaction counsel, collateral law and instrument structure.

Advance rates should reflect cumulative net loss, recovery timing, dilution, prepayment, servicing cost, cash volatility and data confidence. Use dynamic reserves for segments whose realised performance departs from the approved base. A model-risk reserve can respond to calibration drift, missing decision records, data outages or unapproved changes. Its formula should be transparent and avoid double-counting risk captured elsewhere.

Performance triggers should use measures that can be calculated consistently. Examples include first-payment default, thirty-plus and ninety-plus delinquency, cumulative net loss, expected-to-actual loss, cash collection rate, excess spread and borrowing-base deficiency. Define numerator, denominator, cohort, seasoning, cure, recovery and reporting lag in the finance documents. Ambiguous definitions create disputes precisely when performance weakens.

Model and data covenants can include notice or consent for material feature, algorithm, cut-off, product, vendor or data-source changes. Set quantitative thresholds for changes that matter and allow routine maintenance under documented controls. Require periodic independent validation, audit logs, data-quality reporting, override analysis, fairness evidence and incident reporting, while respecting personal-data restrictions.

Remedies should match consequence. A small reporting delay may require cure. A material eligibility breach may stop new funding. Unauthorised model change, data-law breach, loss of licence, fraud or cash-control failure may justify stronger action subject to negotiated documents. The operating playbook should state who calculates a breach, who verifies it, how the lender responds and how affected borrowers continue to receive appropriate service.

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

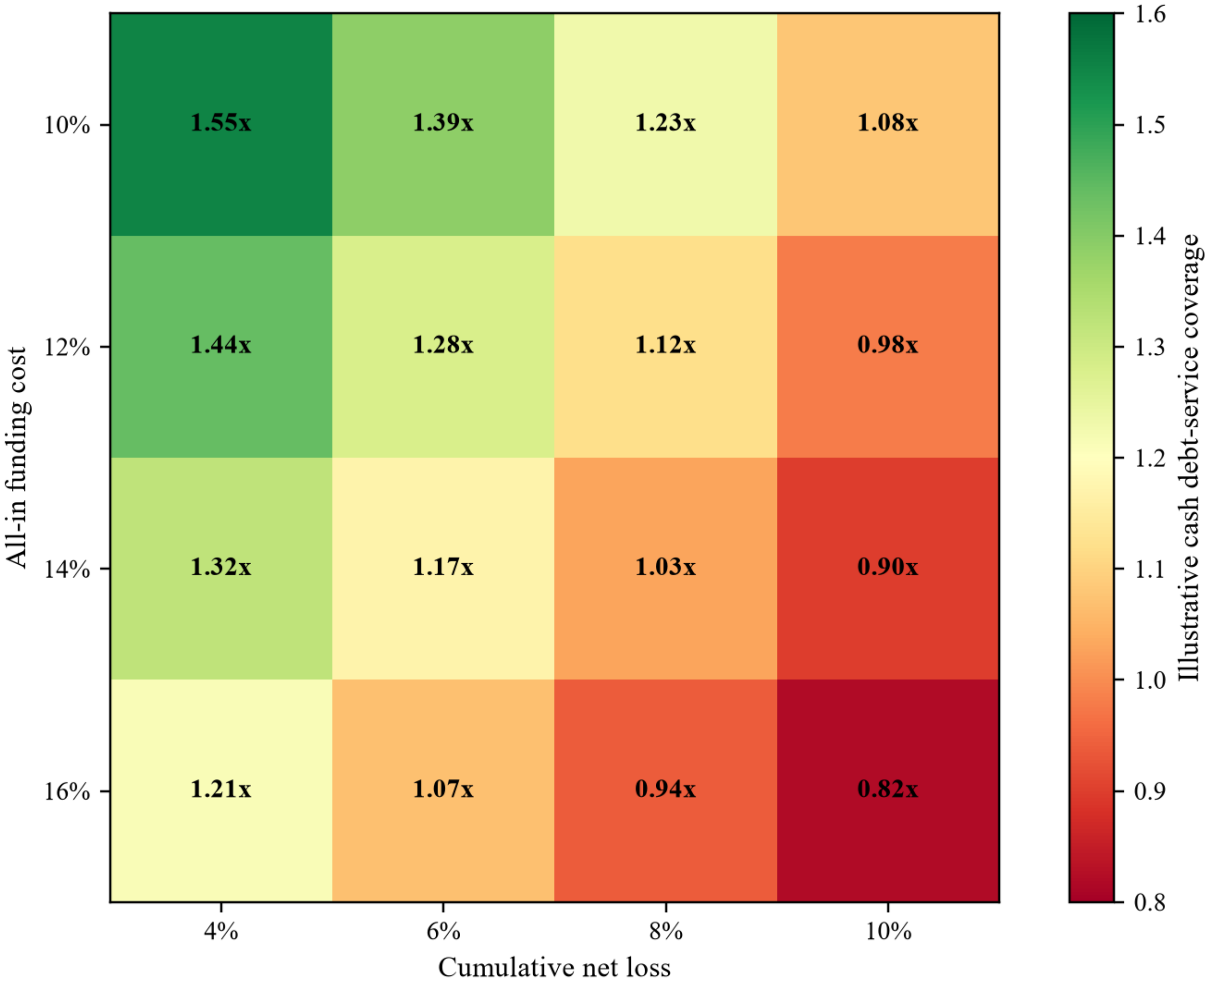

Ratios are modelling assumptions for an illustrative facility. They are not forecasts or proposed terms.

11. Monitor the portfolio and the model together

The monthly funding report should reconcile collateral, collections and covenant calculations with model and data indicators. Separate contractual loan performance from analytical indicators so finance records remain controlled. Link them through stable identifiers, reporting dates and documented transformations.

Portfolio measures include originations, approvals, average ticket, price, tenor, delinquency, cure, recovery, prepayment, fraud, complaints, cash collections and cumulative net loss. Model measures include score distribution, calibration, discrimination, stability, feature availability, missingness, reason codes, overrides and challenger performance. Data measures include provider uptime, ingestion delay, match rate, schema change, validation failure and consent or rights events where applicable.

Set warning, action and stop thresholds. Warning prompts investigation. Action requires a documented mitigation such as lower limits, extra review or reserve. Stop suspends affected originations or funding until evidence supports reopening. The committee should approve thresholds before deterioration occurs. A model owner should not be able to redefine the metric after a breach without independent approval.

Economic change can create out-of-distribution risk. Inflation, employment shocks, rate changes, regional disruption, regulatory change or a new acquisition channel can move applicant behaviour beyond the training sample. Monitor both input distributions and realised outcomes. Scenario tests should show what happens if historical relationships weaken rather than assuming the model will adapt automatically.

Independent validation should report limitations, unresolved findings and remediation dates. Internal audit can test governance and evidence. External specialists may review high-risk systems or areas where expertise is limited. The capital provider should receive material findings and closure evidence, subject to appropriate confidentiality and data safeguards. A clean summary without access to findings provides limited assurance.

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
consumer	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 consumer-credit facility

Consider a hypothetical Indian non-bank financial company originating unsecured instalment loans through its own application and selected distribution partners. It seeks a three-year secured facility to fund eligible receivables. The model uses bureau information, consented bank-transaction data, application fields and repayment history. All figures in this section are assumptions for method demonstration.

The proposed committed facility is INR 3.0 billion. The opening eligible pool is INR 2.4 billion, the headline advance rate is 75 percent, and the lender contributes the remaining funding and reserves. Weighted average original tenor is twelve months. Contractual portfolio yield is 24 percent, assumed servicing and collection cost is 4 percent, and all-in facility cost is 13 percent. The observed mature-vintage cumulative net loss range is 5.2 to 7.1 percent, while two recent vintages are developing above the earlier curve.

The lender presents a model with an out-of-time Gini of 43 percent and an expected-to-actual loss ratio of 0.94 on the validation sample. The same ratio rises to 1.18 for recent digitally acquired thin-file borrowers, meaning observed loss exceeds the model estimate in that cohort under the assumed calculation. Bank-transaction coverage also falls from 72 percent in the development sample to 54 percent in the latest quarter. These facts weaken a uniform advance-rate case.

The diligence team reconstructs sampled decisions and finds that most records reproduce. A material minority uses a vendor-derived feature whose definition changed without a versioned schema. Those receivables are excluded until the lender rebuilds lineage and validation. The funder also separates loans approved by the automated path from manual overrides and sets a concentration cap for recent thin-file digital acquisition.

The negotiated case uses a 70 percent opening advance rate, a dynamic loss reserve, a 10 percent cap on the affected cohort, and monthly expected-to-actual reporting. A trigger applies if the three-month expected-to-actual ratio exceeds 1.15 after minimum seasoning or if critical data coverage falls below the approved floor. Model, cut-off, vendor and policy changes above defined thresholds require notice and, for material changes, consent before new receivables qualify.

The downside case assumes cumulative net loss of 9 percent, funding cost of 15 percent, slower recoveries and a lower approval rate. Availability contracts as the reserve increases. The borrower retains enough equity and liquidity under the assumed case to continue servicing. If the downside breaches the approved cash coverage floor, the facility shifts to cash sweep and stops financing new receivables. These are illustrative mechanics rather than recommended terms.

The case shows why model metrics need a transaction translation. The credit decision follows from cohort loss, data coverage, reproducibility, cash control and the lender's ability to remediate. A headline Gini or approval lift does not determine leverage.

13. Establish a retained implementation and assurance cycle

The facility should move through controlled gates. Gate one confirms legal structure, regulated perimeter, asset ownership, cash control and data rights. Gate two reconciles the loan tape, financial records, collections and mature vintages. Gate three validates data provenance, models, policy, fairness, explanations, overrides and fallback. Gate four documents terms, eligibility, reserves, triggers, reporting and remedies. Gate five tests operations before first funding.

The first operational test should simulate a complete reporting cycle. Load the collateral file, calculate eligibility, apply concentrations and reserves, reconcile collections, produce covenant calculations, trace sampled decisions and issue the investor report. Run a data-feed outage and a model-change scenario. Record exceptions, owners and closure evidence.

After closing, maintain a quarterly review that combines funding, portfolio, model, data, consumer and vendor evidence. Material changes should return to the relevant gate. A new product or distribution channel may require fresh validation and facility approval. A minor code correction can follow a lower-risk path if the policy defines it and independent control confirms the classification.

Assign decision rights. The board or delegated committee approves risk appetite and policy. Credit owns lending decisions. Model development builds the system. Independent validation challenges it. Data and technology control inputs and resilience. Compliance and legal oversee applicable conduct and rights. Finance reconciles cash and reporting. Internal audit tests the framework. The capital provider controls funding eligibility and remedies under the agreed documents.

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

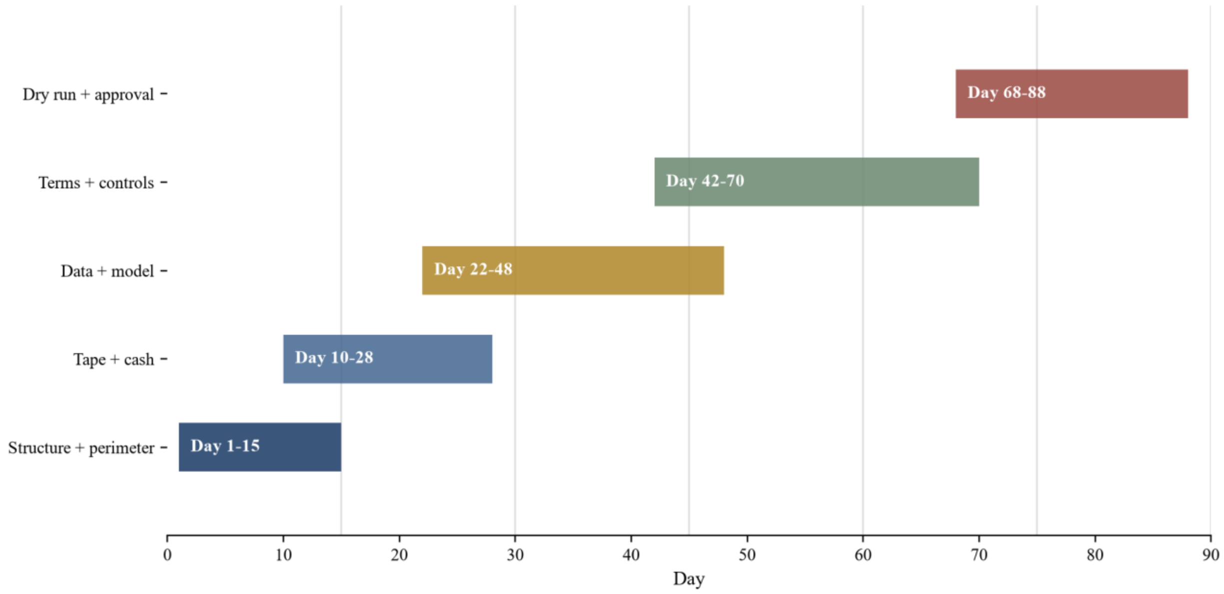

Sequence should be adapted to instrument, portfolio maturity and regulatory requirements.

14. Limitations and further research

This framework does not determine whether a particular data element may lawfully be collected or used, whether a lending product complies with applicable regulation, or whether a proposed security interest is enforceable. Those conclusions require current transaction facts and qualified professional advice. RBI directions, the DPDP commencement schedule and other requirements can change, so diligence should use the rules in force at the relevant date.

Observed model performance does not prove causality. Applicant selection, marketing, policy, pricing, fraud controls, servicing and economic conditions can change outcomes. Alternative-data coverage often reflects self-selection. A model can appear stronger because it is tested only on applicants who completed the data connection. Reject inference and missing labels also complicate validation because repayment outcomes are not observed for declined applicants.

Fairness analysis has legal, statistical and operational limits. Sensitive attributes may be unavailable or restricted. Proxy methods can introduce error. Multiple definitions of fairness can conflict. Small sample sizes widen uncertainty. The institution should document its objective, approved data, methods, limitations and remediation, and avoid claiming that one metric proves equitable treatment.

The hypothetical facility illustrates mechanics and contains no market forecast. Advance rates, reserves, triggers, prices and loss assumptions should come from the specific portfolio, legal structure, fund mandate and negotiated documents. Further empirical research could compare the stability of bank-transaction, bureau and commerce features across Indian consumer cohorts; evaluate privacy-preserving fairness methods; and measure whether human review improves decisions after accounting for selection and incentive effects.

15. Conclusion

Alternative data becomes financeable when it produces a controlled chain of evidence from lawful source to realised cash. The lender should be able to show what data were used, when they were available, how they were transformed, which model and policy applied, why the decision changed the loan terms, who reviewed exceptions and how the resulting cohort performed.

The capital provider should test four connected risks: asset and cash control; data provenance and continuity; model and policy performance; and consumer and governance outcomes. Time-based validation, vintage analysis, calibration, fairness tests, reason codes, human review, fallback and independent challenge give the investment committee a fuller view than headline accuracy.

Facility terms then make the evidence operational. Eligibility, advance rates, reserves, concentration limits, model-change controls, reporting, triggers and remedies should respond to the observed portfolio and its uncertainties. When a lender cannot reproduce a decision or reconcile its cash, the affected asset should not receive funding until the evidence is repaired.

Frequently asked questions

Q: What is alternative data in consumer-credit underwriting? A: It is information beyond conventional application and bureau fields, such as consented bank transactions, payments, payroll, utilities or commerce records. Each source needs a documented purpose, permission, provenance, timing, quality and decision use.

Q: Can a high model Gini justify a higher advance rate? A: No single model metric determines leverage. The funder should also test calibration, recent vintages, cash collections, recoveries, data coverage, concentration, servicing, legal eligibility and model-change controls.

Q: How should a lender validate a thin-file underwriting model? A: Use later-period and cohort holdouts, compare a transparent challenger, test calibration and stability, examine missing-data paths, preserve decision-time inputs and track realised outcomes through mature vintages.

Q: What fairness evidence should a capital provider request? A: Request the approved test design, data basis, cohort coverage, decision and error metrics, reason-code analysis, human-review outcomes, findings, remediation and continuing monitoring, subject to applicable law and privacy safeguards.

Q: When should human review apply? A: Define review zones for uncertainty, missing or conflicting data, material exposure, suspected fraud, policy exceptions, vulnerable-customer indicators and novel cohorts. Record the reviewer, rationale, evidence and outcome.

Q: What happens if an alternative-data feed fails? A: The lender should follow a tested fallback: pause affected originations, use an approved challenger, narrow limits or route cases to manual policy. The facility should state the reporting and eligibility consequence.

Q: Which model changes should require funder consent? A: Material changes to features, algorithms, cut-offs, products, vendors, data sources or decision policy may require consent under negotiated terms. Routine maintenance can follow a documented lower-risk process with validation and notice.

Q: How does model risk enter a borrowing base? A: It can affect asset eligibility, advance rates, dynamic reserves, concentration limits, reporting frequency, performance triggers and the conditions for financing newly originated receivables.

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. Statement on Developmental and Regulatory Policies, 8 August 2024. https://www.rbi.org.in/scripts/BS_PressReleaseDisplay.aspx?prid=58448

[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. Report of the Working Group on Digital Lending including Lending through Online Platforms and Mobile Apps, November 2021. https://www.rbi.org.in/Scripts/PublicationReportDetails.aspx?UrlPage=&ID=1189

[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.