AI Credit Models in Emerging Markets: From Predictive Lift to Financeable Governance

A Lender Framework for Data Rights, Calibration, Drift, Explainability, Overrides, Realised Loss and Capital Eligibility

Authored by Chennakeshav Adya, Independent Researcher

September 2026

Abstract

Artificial intelligence and alternative data can expand the information available for underwriting borrowers who have limited formal credit histories. Transaction records, mobile-money activity, utility payments, platform sales, accounting data and other signals can reveal cash behaviour that conventional bureau files miss. The same models can create risks through unstable data access, selection bias, opaque decision logic, weak default definitions, rapid portfolio growth, untested overrides and performance deterioration when economic conditions change.

This paper develops a lender framework for deciding when an AI credit model in an emerging market is financeable. It begins with the lending decision and cash outcome, maps the model and data supply chain, tests predictive lift against transparent baselines, reconciles score bands to realised delinquency and loss, measures drift, evaluates explanations and overrides, and converts the findings into eligibility rules, reserves, covenants, validation duties and capital gates. Five original figures and five implementation tables provide a practical architecture for diligence, documentation and post-close monitoring.

Current public guidance supports a lifecycle approach. The Central Bank of the UAE requires financial institutions using big data analytics and AI to establish documented governance, validate applications before launch, conduct ongoing calibration and review, maintain traceable records and cover design through discontinuation.[9] The Financial Stability Board's June 2026 consultation proposes sound practices across organisation-wide governance, development and deployment, cyber, technology and third-party risk.[1] IFC's 2026 work on alternative data and AI reports meaningful potential for inclusion while emphasising transparent, trustworthy implementation and stronger data governance.[16] These sources set governance context; they do not establish the performance of any model or loan pool.

Every score, probability, threshold, approval rate, lift measure, default rate, loss rate, drift statistic, capital treatment, cash flow and transaction outcome in this paper is a hypothetical analytical assumption used to demonstrate the framework. It is not a forecast, market quotation, credit opinion, investment recommendation, legal conclusion, accounting conclusion or regulatory conclusion. An actual financing requires verified source records and qualified advice across credit, law, regulation, accounting, tax, consumer protection, data protection, technology, cybersecurity and model risk.

JEL Classification: G21, G28, G32, O16, O33

Keywords: AI credit models, emerging markets, alternative data, model governance, credit underwriting, model drift, explainability, overrides, realised loss, private credit, financial inclusion

1. Begin with the credit decision and the cash outcome

An AI model is valuable to a lender only when its decisions improve risk-adjusted cash outcomes under an approved credit strategy. Accuracy, area under the curve, Gini, precision or another technical metric can describe model behaviour, while financeability depends on approval, pricing, exposure, delinquency, recovery, loss, liquidity and borrower treatment. The diligence team should therefore start from a decision record and trace it to contractual and realised cash.

The minimum record links application, borrower, product, amount, term, price, model version, input snapshot, score, decision threshold, explanation, human override, approval authority, disbursement, repayment schedule, cash received, arrears, cure, restructure, recovery and write-off. Stable identifiers should connect the model platform, loan-management system, bank accounts, credit bureau submissions and general ledger. A sample should be reproducible from original data and code or a controlled scoring service.

The lender should distinguish three questions. Did the model rank risk better than an approved baseline? Did the institution convert that ranking into sensible limits, pricing and policy decisions? Did those decisions produce cash performance consistent with the underwriting case? A positive answer to the first question cannot compensate for weak policy or operations.

The opening diligence output should be a score-to-cash lineage. Every later claim about inclusion, predictive lift, capital efficiency or valuation should reconcile to it.

Figure 1. Model-governance map from data rights to realised cash

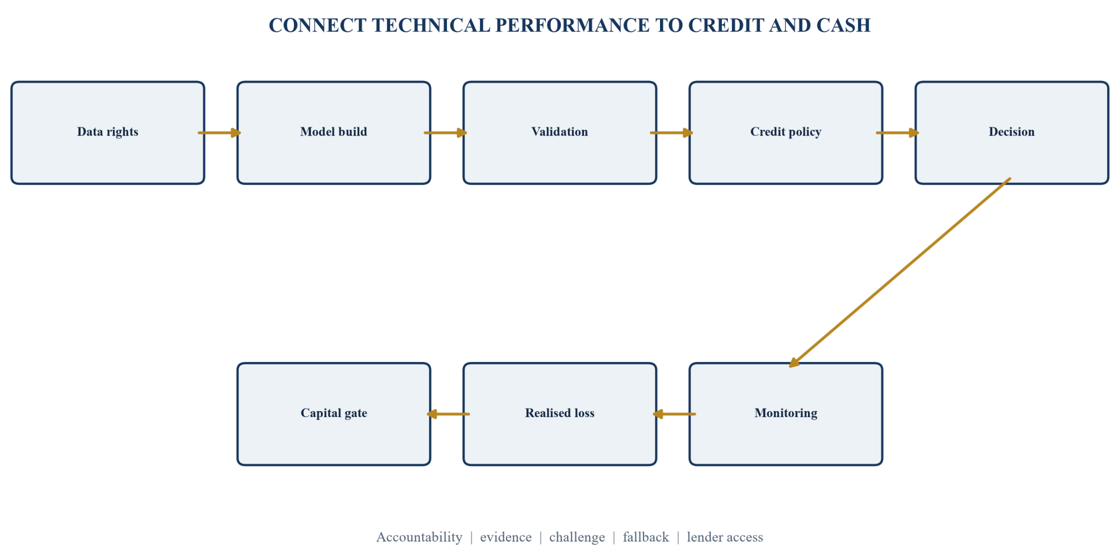

The diagram is a governance framework. Roles, evidence and approvals must be verified for each institution and jurisdiction.

2. Define financeability before measuring model performance

A financeable model has a defined use, accountable owner, lawful data, controlled implementation, independent challenge, stable monitoring, explainable decisions, operational fallback and demonstrated connection to portfolio cash. These attributes should be defined before the lender reviews headline performance. Otherwise, a strong statistic can dominate the discussion while essential controls remain untested.

The model may support screening, approval, limit, price, fraud, collections or early warning. Each use has a different materiality and error cost. A ranking model used as one input to a senior credit officer differs from an automated decline model. A model supporting small, short-tenor advances differs from a model setting a multi-year corporate facility. Financeability criteria should follow the decision's impact and reversibility.

The FSB's proposed 2026 sound practices organise AI risk across governance and the development, deployment, cyber, technology and third-party lifecycle.[1] NIST's AI Risk Management Framework uses govern, map, measure and manage functions and calls for ongoing evaluation, documented limitations and procedures to disengage systems whose performance is inconsistent with intended use.[19] These frameworks can inform the control architecture while local law and prudential requirements remain decisive.

The term financeable does not mean that a model is perfect. It means that residual uncertainty is identified, measured, governed and reflected in advance rate, reserve, covenant, price and lender rights.

3. Map the model, policy and system perimeter

An institution can operate several models behind one score. Identity, fraud, affordability, credit, pricing, limit and collections components may feed a policy engine. Vendor scores can be transformed or combined with rules and human judgement. The diligence perimeter should include every component that changes borrower outcome or lender loss.

The inventory should record purpose, owner, legal entity, jurisdiction, borrower segment, product, decision role, algorithm, model version, training period, input sources, vendor, validation date, materiality, fallback and retirement plan. Shadow models, champion-challenger tests and spreadsheet adjustments should remain visible. A model outside the formal inventory is a governance finding even where its technical performance appears sound.

The Central Bank of the UAE's Model Management Standards identify governance, data management, development, implementation, usage, performance monitoring and independent validation as core components.[11] Its enabling-technology guidance requires an enterprise-wide record and lifecycle coverage from design to discontinuation.[9] These principles support a perimeter that extends beyond code to policy, people, data and operations.

The lender should also map which entity owns the model, which entity originates and services loans, and which entity bears losses. Intercompany scoring or servicing arrangements can affect data access and continuity after a transaction or enforcement.

Table 1. Model inventory and risk-tier evidence map

Inventory item	Evidence	Financeability question	Control response
decision purpose	policy, product terms and workflow	what borrower or exposure outcome can the model change?	materiality tier and approval authority
model mechanics	code, specification, version and transformation logic	can the result be reproduced and challenged?	documentation, validation and access right
data supply	contracts, consent, lineage and quality reports	may inputs be used and will they remain available?	eligibility condition, fallback and vendor covenant
operational use	system logs, thresholds, overrides and exceptions	is production use consistent with approved design?	implementation test and monitoring
outcome linkage	account performance, cash, recovery and ledger	does the score connect to realised loss?	calibration gate, reserve and covenant

Risk tier should reflect decision impact, exposure, autonomy, complexity and substitutability.

4. Establish data rights, provenance and continued availability

Alternative data can include bank transactions, mobile-money flows, utilities, rent, platform sales, accounting records, device signals, geolocation or behavioural information. A lender should verify the legal basis, borrower notice, consent where required, contractual permission, provenance, accuracy, retention, transfer and permitted model use for each field or derived feature.

Data access can be temporary. A platform partner can terminate a feed, change an application programming interface, narrow consent or alter field definitions. A model trained on broad data can degrade when production receives a smaller or slower set. The transaction team should compare development and production coverage and test fallback under loss of each material source.

RBI's digital-lending framework requires need-based collection, prior and explicit borrower consent, clear audit trails and regulated-entity accountability for service providers.[12][13] Kenya's Digital Credit Providers Regulations require customer-information confidentiality and timely, complete and accurate credit information submitted to licensed bureaus.[14] These requirements illustrate why data rights and quality are credit issues rather than technical housekeeping.

The lender should require a data-rights register tied to model features and borrowers. Unsupported fields should be removed or quarantined, the model should be retested, and any performance reduction should enter the credit case.

Table 2. Data-rights and evidence register

Data source	Rights and evidence	Model risk	Lender response
bank or wallet transactions	consent, API terms, account ownership and timestamps	missing periods, reversals and access loss	completeness test, fallback and minimum coverage
platform sales	merchant contract, field dictionary and settlement reconciliation	platform dependence and artificial volume	source reconciliation and concentration limit
bureau records	permissible purpose, enquiry log and dispute process	latency, identity mismatch and incomplete obligations	recency rule and exception reserve
device or behavioural data	lawful basis, necessity, notice and retention	proxy discrimination and weak borrower understanding	feature review, exclusion and fairness test
financial statements	authority, provenance, accounting period and audit status	stale or manipulated inputs	verification hierarchy and manual challenge

Applicable requirements depend on local law, product, borrower and data source.

5. Define the target variable and default consistently

Model performance depends on the outcome it predicts. A target can be missed payment, 30-day delinquency, 90-day delinquency, write-off, fraud, restructuring or net cash loss over a defined horizon. The development definition should match the product, policy and financing objective. A short-horizon delinquency target may rank early payment behaviour while failing to capture lifetime loss.

The target window should allow enough time for the outcome to mature. Recent loans without full performance cannot be labelled good merely because they have not yet defaulted. Cure, restructure, settlement, moratorium, sale and write-off should be treated consistently. Outcomes should reconcile to contractual due dates, cash receipts and accounting records.

The Basel Committee's 2025 Principles for the Management of Credit Risk emphasise sound credit-granting, administration, measurement, monitoring and controls.[7] CBUAE validation guidance calls for review of default definition, development data, model design, assumptions, stability, discriminatory power and calibration.[10] These principles apply directly to the evidence behind the target.

A lender should rerun performance under alternative definitions. If predictive lift disappears when the target changes from early delinquency to net loss, the model may be useful for collections prioritisation and unsuitable for capital eligibility.

6. Construct cohorts that reveal selection and seasoning

Random train-test splits can overstate performance when records from similar time periods, borrowers, merchants or locations appear in both samples. Financeability requires out-of-time and out-of-segment tests. Cohorts should follow origination date, product, geography, channel, borrower type, ticket, term, score band and model version at equal months-on-book.

Selection should be visible. The model observes outcomes only for approved loans unless the institution has another source for declined applicants. Approval policy can therefore create sample-selection bias. A model trained on prior approvals may perform poorly when expanded into previously declined or thin-file groups. Controlled pilots and conservative limits can build evidence.

IFC's 2026 alternative-credit work reports that transaction data are common inputs and that AI is used selectively across models serving borrowers with limited formal histories.[16] The evidence supports opportunity and also reinforces the need to assess context, segment and implementation. A model's success in one country, partner or product does not establish transferability.

The cohort pack should state observation and performance windows, exclusions, missingness and number of accounts and exposure. Small samples should be identified and pooled only with a defensible reason.

7. Translate predictive lift into an economic bridge

Predictive lift describes improvement relative to a baseline. It can mean higher discrimination, better calibration, lower loss at the same approval rate, or higher approval at the same loss. The lender should require a baseline that reflects the actual decision alternative, such as a bureau score, scorecard or policy rule, rather than an artificially weak benchmark.

The bridge should hold one business constraint constant. At the same approval rate, compare realised loss and contribution. At the same expected loss, compare approvals and funded exposure. The analysis should include acquisition, data, model, verification, servicing, fraud, funding and capital costs. A small statistical gain can have limited value after implementation expense and adverse selection.

The hypothetical example below begins with 100 units of baseline expected cash loss. Ranking improvement reduces loss by 18 units, while expansion into thin-file segments adds 7, policy looseness adds 5, data and model operating cost adds 4, and weaker recoveries add 3. The net analytical benefit is 1 unit. Each amount is a hypothetical analytical assumption and should be replaced by verified account-level evidence.

Figure 2. Illustrative bridge from predictive lift to net cash-loss benefit

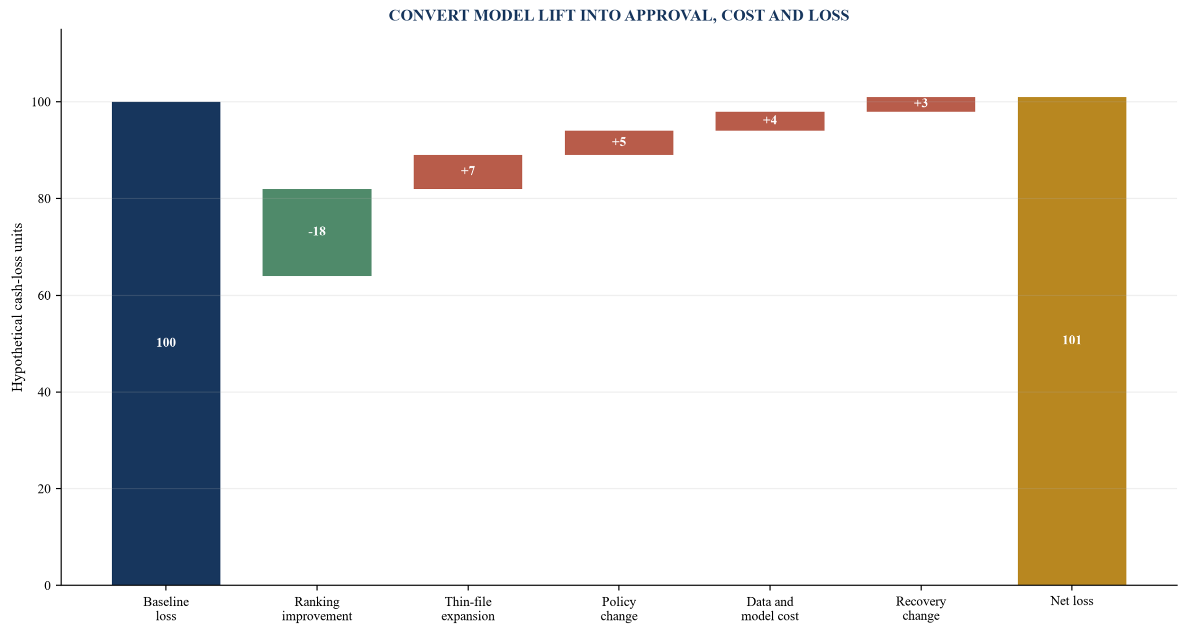

All values are hypothetical analytical assumptions and do not describe a lender, model or market forecast.

8. Test discrimination, calibration and realised loss together

Discrimination measures whether riskier borrowers receive worse scores. Calibration measures whether predicted probabilities correspond to observed outcomes. A model can rank correctly and systematically understate loss. Lenders need both measures because capital and pricing depend on the level of risk, not only order.

Validation should report area under the curve or Gini with confidence intervals, precision and recall where relevant, calibration by score band, Brier or log loss, and observed-to-expected outcomes. The analysis should use account and exposure weights and should show cash loss, not only default count. Recovery timing and cost matter for debt service.

Performance should be separated by development, out-of-time, production and stressed periods. The EBA's report on machine learning in internal-ratings models emphasises prudent use within the credit-risk framework and the interaction with data protection and AI rules.[20] CBUAE validation standards require quantitative review of stability, discriminatory power, sensitivity and calibration.[10]

Approval thresholds should be tested against capacity and liquidity. A model that admits more borrowers can change the pool's mix and overwhelm verification or collections. Financeability therefore links technical tests to operations and funding.

9. Evaluate thin-file inclusion without masking adverse selection

AI can identify borrowers whose cash activity is not represented in conventional files. BIS analysis notes that alternative data can improve default prediction for underserved groups and reveal borrowers whose conventional scores are weak signals.[3] IFC's work also describes models using mobile money, digital payments, platform records and business data.[16][17]

The lender should measure inclusion as a controlled cohort. New-to-credit or thin-file borrowers should have explicit definitions, separate approval, exposure, pricing, delinquency, loss and complaint reporting, and staged limits until performance seasons. Comparisons should adjust for amount, term, geography, channel and economic conditions.

Inclusion claims should distinguish access from outcomes. A larger number of approvals can coexist with unaffordable pricing, repeat refinancing or weak borrower understanding. Cash performance, total cost, repeat borrowing, complaints, hardship and restructuring should be monitored together.

The financing structure can apply lower advance rates or higher reserves to unseasoned segments and release them after verified cohort performance. This lets capital follow evidence without abandoning an inclusion objective.

10. Measure fairness and proxy risk in the actual decision process

Historical data can encode unequal access, treatment and outcomes. Alternative variables can act as proxies for protected or vulnerable characteristics. Fairness analysis should therefore cover data, labels, model output, policy thresholds, overrides, pricing, limits, servicing and collections.

The institution should identify legally permissible comparison groups with counsel. Useful tests can include approval, error rates, calibration, pricing and loss by group, with sample size and confidence. A statistical difference requires investigation of mechanism, business necessity, data quality and potential remediation. One parity metric cannot settle the assessment.

The FSB has warned that historical data and opaque alternative sources can perpetuate bias and make errors difficult for borrowers to correct.[2] The G20 MSME financing action plan notes that automated credit assessment can introduce distortions and that incomplete, unrepresentative data can affect underserved groups.[18]

The lender should require a feature register, proxy review, fairness testing, complaint route and change-control process. Where lawful group labels are unavailable, the institution should document the limitation and use other evidence such as feature sensitivity, geographic outcomes, manual review and borrower feedback.

11. Make explanations useful to borrowers, credit officers and validators

An explanation should match its audience. Borrowers need clear principal reasons and a route to correct data or seek review. Credit officers need decision drivers, sensitivity and policy context. Validators need reproducible methods, stability and limitations. Supervisors need traceability and evidence of accountability.

Global or local feature importance alone may be insufficient. Correlated variables can exchange importance, and post-hoc explanations may vary. The institution should test fidelity, stability and actionability. A reason should reflect the actual decision and avoid exposing sensitive security or fraud controls.

BIS Project Noor is developing explainable-AI methods that can help supervisors assess transparency, fairness and robustness while preserving privacy.[5] A 2026 BIS speech states that institutions remain accountable for explainability and data quality in high-impact decisions such as credit applications.[6] These initiatives support practical explanation while preserving institutional responsibility.

The lender should sample explanations against source data, score changes and policy decisions. A borrower-adverse decision that cannot be explained consistently should not qualify for automated expansion or capital benefit.

12. Govern human overrides as a measurable model component

Human judgement can add information that the model cannot observe. It can also weaken discipline, favour growth targets or hide a deteriorating segment. Each override should record direction, reason, evidence, authority, timestamp, model score, original decision, final terms and subsequent outcome.

Override analysis should distinguish model override from policy exception and data correction. Correcting a verified input error is different from approving a borrower outside risk appetite. Upward and downward overrides should be tracked separately by officer, branch, channel, product, geography and score band.

The hypothetical analysis should compare observed loss for non-overridden approvals, positive overrides and negative overrides at similar model risk. High override frequency can indicate a poorly fitted model or weak adoption. Concentrated override authority creates conduct and operational risk. A low override rate can also be concerning if staff follow an unsuitable model without challenge.

CBUAE validation guidance expressly requires review of rating overrides and the consistency of subjective inputs and potential bias.[10] The lender can set reporting triggers, independent sample review and limits, with temporary suspension where override performance materially exceeds approved loss.

Figure 3. Illustrative override frequency and realised loss by score band

TRACK JUDGEMENT, CONCENTRATION AND OUTCOME TOGETHER

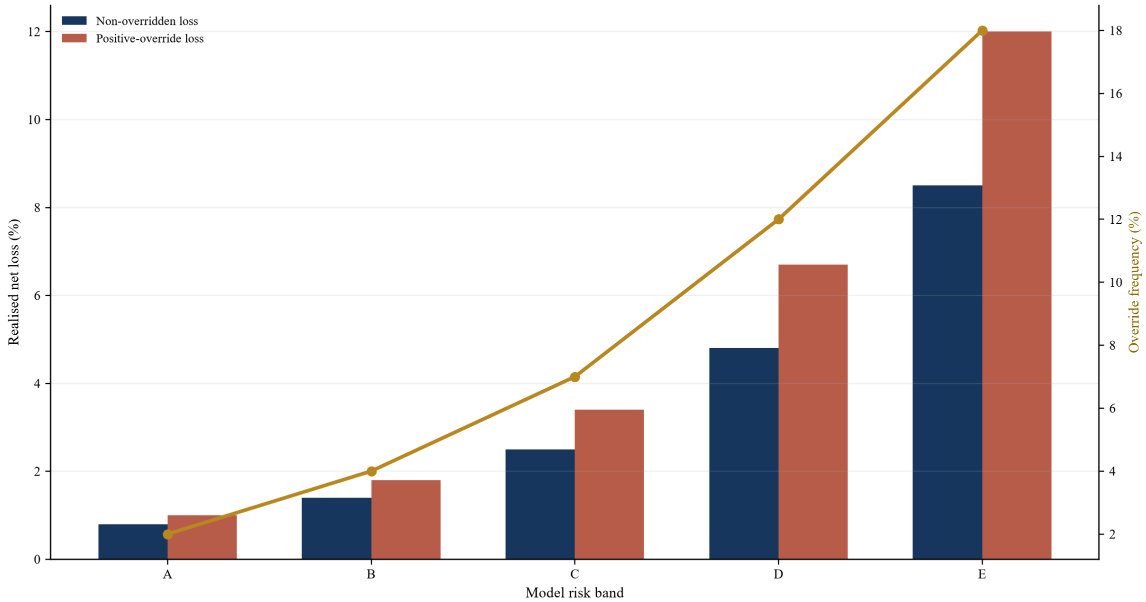

Frequencies and loss rates are hypothetical analytical assumptions and do not describe a lender or model.

13. Monitor data, score, decision and outcome drift

Drift occurs when production differs from development or prior approved periods. Data drift changes input distributions or coverage. Concept drift changes the relationship between inputs and default. Policy drift changes thresholds, rules or human behaviour. Outcome drift changes delinquency, recovery or loss. The lender should monitor all four because a stable score distribution can coexist with worsening cash.

Population stability index, characteristic stability, missingness, category movement and distribution distance can support data monitoring. Score-band approval, decline, pricing and limit changes support decision monitoring. Observed-to-expected loss, roll rates, vintage curves and recovery support outcome monitoring. Metrics should be shown by product, cohort, channel, geography and material data source.

Thresholds need an approved response. A warning can trigger investigation and enhanced reporting. A breach can restrict new originations, raise reserves, reduce advance rate, switch to fallback or require recalibration and validation. The response should consider sample size and seasonality rather than rely on one isolated statistic.

The dashboard below uses hypothetical values to demonstrate governance. It does not provide universal thresholds. Each institution should calibrate triggers to verified historical variability, portfolio risk and regulatory requirements.

Figure 4. Illustrative model-drift dashboard across the credit lifecycle

Scores are hypothetical analytical assumptions on a one-to-five monitoring scale and are not universal limits.

14. Test performance through macroeconomic and local shocks

Emerging-market credit portfolios can be exposed to inflation, currency depreciation, commodity prices, climate events, conflict, policy changes, platform interruptions and local employment shocks. Development data may contain few comparable periods. A model can remain statistically stable while borrowers' repayment capacity changes through variables it does not observe.

Stress testing should combine model and portfolio channels. Inputs can be shocked where a defensible relationship exists; policy thresholds and cash outcomes should also be tested directly. The lender can apply higher transition rates, lower cures, delayed recovery, lower income, higher refinance cost and data outages by cohort and geography. Scenarios should include correlated effects and operational constraints.

BIS research on AI and relationship lending found that AI investment supported credit supply in normal periods in the studied sample, while it did not provide additional credit or interest-rate protection during the Covid crisis.[4] This result is context-specific and illustrates why normal-period performance should not be assumed to persist through a new shock.

The transaction case should compare stressed loss and liquidity with capital, reserve and facility headroom. Model confidence should reduce, rather than increase, where the shock falls outside development experience.

15. Reconcile model monitoring to portfolio accounting and cash

Monitoring reports often use delinquency or default labels produced inside the lending system. The lender should reconcile those labels to contractual schedules, bank receipts, general ledger, provisions, write-offs and recoveries. A model can appear calibrated if arrears are restructured, rolled or written off outside the performance data.

The reconciliation should preserve original and revised schedules, gross and net write-offs, settlement, cure, recovery cost and timing. Account and exposure views should be shown together. Loss should be attributed to the model version and policy in force at origination, with later servicing actions recorded separately.

Capital and facility models may use expected loss, unexpected loss, concentration or stress measures that differ from the model's target. The bridge should state each definition and show how score bands inform probability of default, loss given default, exposure, reserve or advance rate. Unsupported transformations are a financeability gap.

Data adjustments should be controlled and reported. The Basel corporate-governance framework stresses the board's responsibility for data quality and the limitations of models and external inputs.[8] Reconciliation makes that accountability observable.

16. Assign board, management and control accountability

The governing body should approve material AI use, risk appetite and accountability. Senior management should own implementation, resources, competence, monitoring and remediation. Credit, risk, compliance, data, technology, security, legal, internal audit and business teams need defined roles with independent challenge.

The model owner should not approve independent validation. A model committee can review development, change, performance, overrides and retirement. A credit committee should retain authority over credit policy and portfolio risk. Material issues should escalate to board or board risk committee through defined thresholds.

The CBUAE enabling-technology guidance makes the governing body and senior management accountable for outcomes and requires appropriately skilled technical and credit specialists.[9] The South African Prudential Authority and Financial Sector Conduct Authority's 2025 report recommends comprehensive governance, data governance, model-risk management and board-level oversight.[21]

The lender should inspect minutes, challenge, decisions and issue closure rather than accept an organisation chart. Governance becomes financeable when authority changes behaviour, funding and deployment.

17. Require independent validation with reproducible evidence

Independent validation should review conceptual soundness, data, code, implementation, use, performance, limitations and outcomes. Independence can be organisational or achieved through external review, provided the validator has access, competence and freedom from development incentives. Materiality should determine depth and frequency.

Validation should reproduce samples, confirm target construction, test leakage, compare baselines, examine tuning and selection, run out-of-time and segment tests, assess calibration, fairness, explanations, overrides, stress, drift and fallback. Implementation validation should confirm that approved code, transformations, thresholds and data operate in production.

The validator should state findings, severity, compensating controls and conditions for use. A model with an open critical finding should not receive the same capital or facility treatment as a fully validated model. Remediation should be retested by an independent party.

CBUAE standards identify stability, discriminatory power, sensitivity, calibration, governance, overrides, documentation and data as validation areas.[10][11] The validation report should be delivered to the lender under confidentiality terms with sufficient detail for independent assessment.

Table 3. Independent-validation test programme

Test area	Minimum evidence	Failure signal	Financeability response
data and target	lineage, rights, completeness, leakage and default reconstruction	unsupported features or immature outcomes	exclude model or affected cohorts
performance	baseline, out-of-time, segment, calibration and cash loss	lift disappears or loss is understated	reserve, reprice or restrict originations
fairness and explanation	group outcomes, feature sensitivity and reason fidelity	unstable or misleading explanations	human review and remediation gate
implementation	code, transformation, version, threshold and access tests	production differs from approved model	suspend use and correct deployment
monitoring and fallback	drift, incidents, overrides, fallback and retirement	breach has no tested response	covenant, fallback activation and stop right

Scope and frequency should follow model materiality, jurisdiction and portfolio risk.

18. Diligence vendor models and concentrated technology dependencies

A vendor may supply scores, data, features, code, hosting or monitoring. The regulated institution and lender remain exposed to model and continuity risk even when intellectual property is external. The diligence team should identify which evidence, audit, validation and change rights exist under contract.

The agreement should cover data use, ownership, confidentiality, security, model changes, performance, service levels, incident notice, subcontracting, regulator access, lender access where required, termination, transition and record retention. A black-box service may limit the institution's ability to explain decisions or reproduce historical scores. Contractual access should be tested against actual vendor practice.

Concentration can arise when many lenders use the same data aggregator, cloud service or model. A common change or outage can affect several portfolios simultaneously. The FSB's proposed sound practices include third-party, cyber and technology risk.[1] NIST also treats third-party software, hardware and data across the lifecycle.[19]

The financing case should include replacement time, data portability, parallel operation and fallback. A vendor termination that disables underwriting should be a liquidity and business-continuity scenario.

19. Protect privacy, security and borrower recourse

Credit models can process sensitive financial, behavioural and device information. The institution should collect only data needed for an approved purpose, protect it, retain it for a justified period and give borrowers applicable access, correction and review rights. Consent should be meaningful where relied upon and should not conceal unnecessary collection.

Security diligence should map data at rest and in transit, identities, privileged access, encryption, keys, logs, development environments, model artefacts, backups and vendors. Model extraction, data poisoning, adversarial input and unauthorised feature changes should enter the threat model. Incidents should link to affected decisions and borrowers.

World Bank guidance on responsible financial access emphasises consumer protection for digital finance and risk-based supervision.[15] RBI's framework requires clear privacy policy, limited service-provider storage and regulated-entity oversight.[12][13] Kenya's regulations require confidentiality and complaint handling.[14]

The lender should require notification of material data, security or borrower-treatment incidents, with remediation and financial impact. Borrower recourse supports both legal compliance and credit data quality because corrected records improve future decisions.

20. Build a jurisdiction-specific regulatory matrix

Emerging markets differ in licensing, consumer protection, automated decisions, data localisation, bureau reporting, outsourcing, model standards and supervisory access. A regional platform should not assume that one model and consent flow can be deployed unchanged. The legal and operating perimeter should be mapped by country, product, borrower and entity.

The UAE has in-force model-management and enabling-technology standards covering governance, data, validation and lifecycle controls.[9][10][11] India places regulated-entity responsibility, consent, audit trail and digital-lending controls around service providers.[12][13] Kenya licenses and supervises digital credit providers and regulates credit-information use and customer protection.[14] Nigeria's open-banking guidelines include board-approved data governance and consideration of algorithmic systems.[22]

The matrix should identify approval, notification, documentation, explanation, data transfer, retention, audit and complaint requirements. Changes should be version-controlled and linked to system configuration. Local counsel and regulators may need to confirm interpretation.

A jurisdiction without prescriptive AI rules still requires credit, data, conduct, outsourcing and governance compliance. The lender should price uncertainty and require evidence rather than treat regulatory silence as permission.

21. Convert model evidence into capital-eligibility gates

Capital eligibility means that lender or investor exposure receives the agreed availability, advance rate, reserve, pricing or portfolio treatment because specified evidence is present. The gate should not rely on a vendor label or technical metric. It should combine legal, data, model, portfolio and operational controls.

The first gate confirms rights, inventory, governance and reproducibility. The second confirms independent validation, baseline lift and calibration. The third confirms production implementation, explanations, overrides and monitoring. The fourth confirms seasoned realised loss and stress performance. The fifth confirms continuing compliance, vendor access, fallback and covenant reporting.

A model can pass selected gates and receive limited treatment. Unseasoned cohorts can have lower advance rates or higher reserves. A breach can stop new eligibility while existing exposures amortise under controlled servicing. Restoration should require verified cure and independent sign-off.

The hypothetical gate diagram below shows progressive capital access. Percentages are analytical assumptions and do not represent market terms. Actual treatment depends on verified evidence, legal structure and risk appetite.

Figure 5. Illustrative evidence gates for model-supported capital eligibility

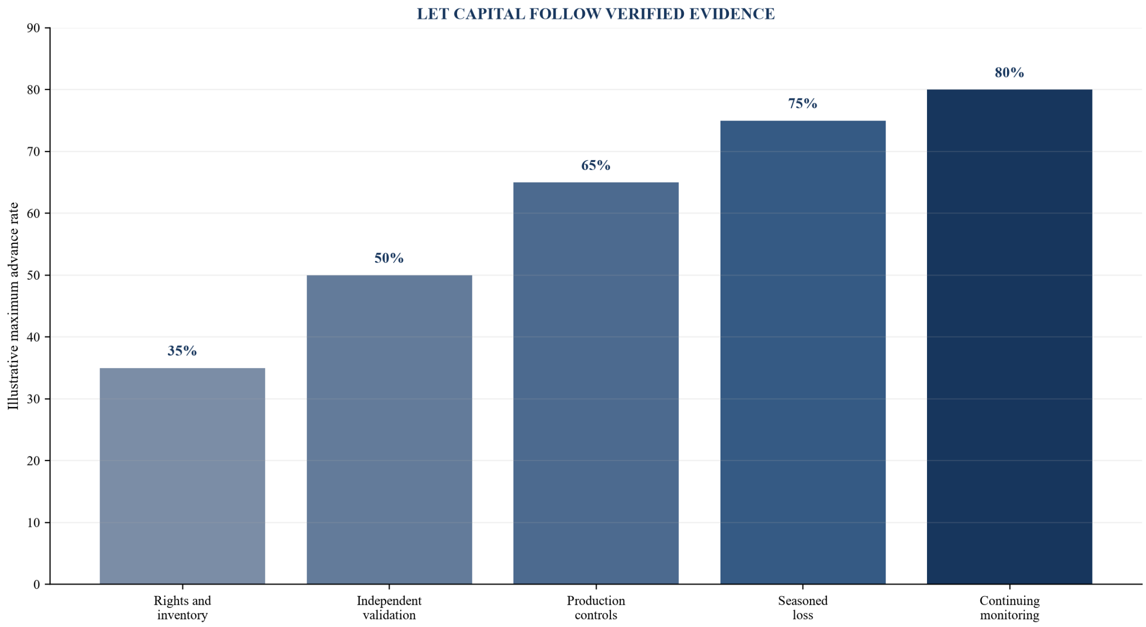

Advance rates are hypothetical analytical assumptions and do not represent lender terms or market quotations.

22. Embed model risk in facility structure and covenants

Facility documents should define eligible loans, model versions, approved products, data requirements, concentration, seasoning, arrears, restructuring, fraud and documentation. Model-supported eligibility can be conditional on validation, monitoring and access. The borrower should not change a material model, threshold or policy without agreed notice and approval where appropriate.

Reporting can include score distribution, approvals, overrides, data coverage, drift, observed-to-expected loss, vintage performance, fairness or conduct indicators, incidents and validation findings. Triggers should have defined actions: enhanced reporting, reserve, advance-rate reduction, pause, remediation or fallback. Materiality and cure periods require careful drafting.

Cash control remains essential. A model cannot substitute for receivable ownership, enforceability, controlled collections, borrowing-base reconciliation, liquidity and capital. The facility should connect model risk to the ordinary credit architecture.

The lender should receive audit and access rights sufficient to test the evidence while protecting borrower data and vendor intellectual property. Independent experts can support sensitive review under agreed confidentiality arrangements.

Table 4. Model-risk covenant and reporting framework

Control area	Reporting evidence	Illustrative trigger	Response
data supply	coverage, missingness, vendor status and consent exceptions	material source loss or unsupported feature	affected-loan exclusion and fallback
performance	calibration, vintages, roll rates and net loss	observed loss exceeds approved band	reserve increase and origination review
drift	input, score, decision and outcome dashboard	sustained warning or breach	investigation, recalibration and validation
overrides	rate, reason, authority and realised performance	concentration or excess override loss	enhanced review and authority restriction
governance	model changes, validation and issue log	unauthorised change or critical finding	eligibility stop and independent cure test

Terms are illustrative and require transaction-specific legal and regulatory advice.

23. Apply the framework in transaction and lender diligence

An investor or lender should request the model inventory, policies, data map, contracts, development files, validation, production logs, decision records, portfolio tape, cash and accounting reconciliation, complaints, incidents, regulator correspondence and committee minutes. Management presentations should be tested against source records and independently reproduced samples.

The diligence issue register should connect each finding to value, credit structure, closing condition, covenant, remediation owner and monitoring. A model that improves approval but lacks durable data rights may reduce value. A model with modest lift and strong governance may support controlled capital. The conclusion should state mechanism and evidence.

Transaction protections can include conditions for validation, data or vendor consent, remediation, escrow, indemnity and retention. Valuation should deduct required investment and use verified contribution and loss. Future model value should be credited through milestones rather than unsupported claims.

The review should include the institution's ability to operate if the model is unavailable. A controlled manual or simpler-scorecard fallback protects borrowers, collections and lender cash during outage or challenge.

24. Execute a 100-day financeability plan

The first 30 days should lock the inventory, decision lineage, data rights, model versions, validation findings, portfolio reconciliation and critical incidents. The institution should freeze unauthorised changes and establish governance, issue ownership and lender reporting. High-risk or unsupported cohorts can receive temporary limits.

Days 31 to 60 should reproduce performance, recalibrate where justified, test explanations and overrides, close data gaps, validate production, establish drift thresholds and rehearse fallback. Legal, privacy, security and vendor actions should proceed in parallel with controlled dependencies.

Days 61 to 100 should test remediated controls, season new cohorts, implement capital gates and covenants, confirm borrower recourse and obtain independent assurance. Expansion should follow verified results rather than the calendar.

Each initiative needs an owner, evidence, cost, decision gate and accepted residual risk. The plan should report cash and borrower outcomes alongside technical completion.

Table 5. Illustrative 100-day model-financeability plan

Period	Evidence priority	Credit action	Acceptance gate
days 1 to 15	inventory, data rights, versions, cash and decision samples	protect exposure and freeze unsupported change	reproducible score-to-cash lineage
days 16 to 30	validation issues, vendor access, incidents and overrides	segment limits and temporary reserves	approved remediation and fallback
days 31 to 60	out-of-time tests, calibration, fairness and production	reset thresholds and eligibility	independent validation of material fixes
days 61 to 100	drift controls, covenants, recourse and governance	staged advance rate and monitored expansion	committee and lender evidence pack
beyond day 100	seasoned loss, stress and continuing assurance	release or tighten capital treatment	verified performance against gate

Timing should follow risk, portfolio seasoning, regulatory requirements and operational readiness.

25. A decision framework for the investment committee

The investment committee should receive a connected answer to six questions. Does the institution have the right to use and continue receiving each material input? Can the model and decision be reproduced? Does performance exceed a transparent baseline out of time and by segment? Is probability calibrated to realised cash loss? Are explanations, overrides, fairness, drift and fallback governed? Does the facility convert those controls into enforceable capital gates?

The recommendation should distinguish verified evidence, open items and hypothetical analytical assumptions. Every material finding should affect value, advance rate, reserve, covenant, condition, monitoring or the decision to stop. A technical model score should never stand alone in an investment conclusion.

AI credit models can expand information and support inclusion where formal records are thin. Their value becomes durable when data rights, validation, policy, operations and cash outcomes remain connected through the lifecycle. Financeable governance makes that connection visible to borrowers, management, regulators and capital providers.

The strongest evidence pack is reproducible. A reviewer should be able to move from a facility or valuation output back to cohort performance, decision records, model version, inputs, consent, contracts and realised cash. That chain supports responsible scaling and protects capital when conditions change.

References

[1] Financial Stability Board, Sound Practices for Responsible Adoption of Artificial Intelligence: Consultation Report, 10 June 2026, https://www.fsb.org/2026/06/fsb-consults-on-sound-practices-for-the-responsible-adoption-of-artificial-intelligence-ai/

[2] Financial Stability Board, The Financial Stability Implications of Artificial Intelligence, 14 November 2024, https://www.fsb.org/2024/11/the-financial-stability-implications-of-artificial-intelligence/

[3] 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

[4] Bank for International Settlements, Artificial Intelligence and Relationship Lending, Working Paper No. 1244, 19 February 2025, https://www.bis.org/publications/working-paper-1244-artificial-intelligence-and-relationship-lending

[5] BIS Innovation Hub, Project Noor, https://www.bis.org/project/noor

[6] Bank for International Settlements, Regulation and Supervision of the Financial Sector in the Age of Artificial Intelligence, 20 May 2026, https://www.bis.org/speeches/20260520-regulation-and-supervision-financial-sector-age-artificial-intelligence

[7] Basel Committee on Banking Supervision, Principles for the Management of Credit Risk, 2025, https://www.bis.org/bcbs/publ/d595.pdf

[8] Basel Committee on Banking Supervision, Corporate Governance Principles for Banks, consolidated framework, https://www.bis.org/committees/bcbs/basel-consolidated-guidelines/module/cgo/10

[9] Central Bank of the UAE, Guidelines for Financial Institutions Adopting Enabling Technologies: Big Data Analytics and Artificial Intelligence, https://rulebook.centralbank.ae/sites/default/files/en_net_file_store/CBUAE_EN_2413_VER1.pdf

[10] Central Bank of the UAE, Model Management Standards: Monitoring and Validation, https://rulebook.centralbank.ae/en/rulebook/211-monitoring-and-validation

[11] Central Bank of the UAE, Model Management Standards, https://rulebook.centralbank.ae/en/rulebook/model-management-standards

[12] Reserve Bank of India, Guidelines on Digital Lending, 2 September 2022, https://systemhealth.rbi.org.in/Scripts/NotificationUser.aspx_Id%3D12382%26Mode%3D0.html

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

[14] Central Bank of Kenya, Digital Credit Providers Regulations 2022, https://www.centralbank.go.ke/wp-content/uploads/2022/03/L-.N.-No.-46-Central-Bank-of-Kenya-Digital-Credit-Providers-Regulations-2022.pdf

[15] World Bank, Responsible Financial Access and Financial Consumer Protection Resources, https://responsiblefinance.worldbank.org/en/responsible-finance

[16] International Finance Corporation, Cracking the Credit Code: Alternative Data and AI for Financial Inclusion, 2026, https://www.ifc.org/en/insights-reports/2026/cracking-the-credit-code-alternative-data-and-ai-for-financial-inclusion

[17] International Finance Corporation, MSME Banking in the Digital Era, 2025, https://www.ifc.org/content/dam/ifc/doc/2025/msme-banking-in-the-digital-era.pdf

[18] G20 Global Partnership for Financial Inclusion, Action Plan for MSME Financing, 2025, https://www.ifc.org/content/dam/ifc/doc/2025/gpfi-action-plan-for-msme-financing.pdf

[19] National Institute of Standards and Technology, Artificial Intelligence Risk Management Framework 1.0, https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf

[20] European Banking Authority, Follow-up Report on the Use of Machine Learning for Internal Ratings-Based Models, 4 August 2023, https://www.eba.europa.eu/publications-and-media/press-releases/eba-publishes-follow-report-use-machine-learning-internal

[21] South African Reserve Bank and Financial Sector Conduct Authority, Artificial Intelligence in the South African Financial Sector, 24 November 2025, https://www.resbank.co.za/content/dam/sarb/publications/media-releases/2025/sarbfsca-ai-report/Artificial%20Intelligence%20in%20the%20South%20African%20Financial%20Sector%20Final.pdf

[22] Central Bank of Nigeria, Operational Guidelines for Open Banking in Nigeria, https://www.cbn.gov.ng/Out/2023/CCD/Operational%20Guidelines%20for%20Open%20Banking%20in%20Nigeria.pdf

About the Author

Authored by Chennakeshav Adya, Independent Researcher

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 close.

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, leading 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.