THE CHARTERED INSTITUTE OF LOAN & RISK MANAGEMENT OF NIGERIA

PAPER PRESENTATION

BY

DR. CHINEDU OKEKE

Senior Manager, Zenith Bank PLC

A CRITICAL LOOK AT AI-ENABLED STRATEGIES FOR MANAGING THE LOAN LIFECYCLE

@

16TH ANNUAL RISK PROFESSIONAL INTERNATIONAL CONFERENCE

NICON LUXURY HOTEL,

AREA 11, GARKI

ABUJA.

NOVEMBER 27, 2026.

ABSTRACT

AI in lending has moved past the pilot stage, but the evidence of value is uneven. A 2026 survey by the Cambridge Centre for Alternative Finance of 628 organisations found that more than 80% of financial services firms use AI at some level, yet only 40% of respondents reported higher profitability and 76% of large institutions found AI value hard to measure. This paper examines AI-enabled strategies across four stages of the loan lifecycle: origination and underwriting, fraud detection, portfolio monitoring and early warning, and collections and recovery. It draws on McKinsey and IACPM research, the Reserve Bank of India's FREE-AI report, US and EU regulatory changes through November 2026, and institution-level case results, most of which are vendor-reported.

Reported gains are large: 20% to 60% analyst productivity in credit memo workflows, loan decisions shortened from days to minutes, and a 20-point rise in one bank's recovery rate. Evidence on credit losses, fraud-model durability, and collections outcomes is thinner than these headlines suggest. AI's most durable advantage is shortening the time between a signal and a decision, and that advantage is worth having only when decisions are explainable, auditable, and fair. Governance data point to a supervisory exposure: in the RBI survey, 18% of AI-using entities kept audit logs and 35% validated models for bias.

1. INTRODUCTION: THE LIFECYCLE AS THE UNIT OF ANALYSIS

Lenders have traditionally managed the loan lifecycle as a sequence of separate functions: origination, underwriting, disbursement, monitoring, and collections. Each has its own systems, KPIs, and risk appetite. AI changes this arrangement mainly by linking the stages, so that a signal captured at one point, such as a fraud flag at application or a cash-flow change during monitoring, can alter decisions at another. The operational case is substantial. Legacy origination workflows can take seven to fifteen business days from application to disbursement, while technology-first platforms complete the same process in hours (Jifiti, 2026). The same source reports that banks using AI-driven origination with transparent decision frameworks see approval rates about 20% higher than manual processes, with no corresponding rise in credit risk. That figure comes from a trade source and has not been checked against primary data. Fintech lenders' share of the US personal loan market rose from 5% in 2013 to 38%, while banks' share fell from 40% to 28%. A shift of this size is consistent with an access gap created by legacy processes, but it does not prove one. The threat has grown in parallel. Javelin estimates US losses from traditional identity fraud at $27.2 billion in 2024, up 19%, and $27.3 billion in 2025. TransUnion put US lender exposure to synthetic identities at $3.3 billion. Vendor data indicate that AI-generated document fraud rose about fivefold between April and December 2025. The technologies that widen credit access also widen the attack surface, so institutions that treat the two problems separately will manage neither well.

Sections 3 to 6 examine the four lifecycle stages. Sections 7 to 9 cover orchestration, governance, and implementation. Section 10 sets out what the evidence does not support, and Section 11 gives conclusions.

2. METHODOLOGY AND SOURCES

This paper synthesises evidence from four categories of source:

1. **Industry studies.** McKinsey and IACPM's 2025 survey of 44 financial institutions (including a multiagent credit memo case), McKinsey's 2026 analysis of AI-enabled risk operating models, and the Cambridge Centre for Alternative Finance's 2026 survey of 628 organisations.
2. **Regulatory and supervisory material.** The Reserve Bank of India's FREE-AI Committee Report (August 2025) and its survey of regulated entities, the CFPB's April 2026 Regulation B final rule, and the EU AI Act as amended by the Digital Omnibus (July 2026).
3. **Institution-level deployment data.** Published results for Blend, Harmoney, Garanti BBVA, Trius FCU, and Central Bank of India. These come from vendor case studies or vendor blogs.
4. **Research on fraud and credit risk.** Singh, Reddy, and Reddy (2025) in the Journal of Information Systems Engineering and Management, Javelin's identity fraud studies, and TransUnion data.

Each table labels its sources by type. Independent means a regulator, academic body, or peer-reviewed study. Industry means a survey or analysis from a consultancy, credit bureau, or industry body. Vendor means results published by the technology provider. Where a source reports a projection rather than an outcome, the text says so. Vendor results are subject to selection bias and are best read as upper bounds. Section 10 returns to these limits.

3. STAGE 1: ORIGINATION AND UNDERWRITING

3.1 The adoption curve

Reported adoption depends heavily on how it is defined. Table 1 shows a range from 12% to more than 80% across surveys that measure different things: any use, active development, or live production.

Table 1: AI adoption indicators in lending and financial services

Indicator	Reported figure	Source	Type
Financial services firms using AI at some level (global)	More than 80%	Cambridge CCAF, April 2026 (n=628)	Independent
Firms experimenting with agentic AI	52%	Cambridge CCAF, April 2026	Independent
Mortgage lenders with AI in live production	17% (over 80% evaluating)	Mortgage Collaborative, June 2026	Industry
North American credit respondents with any gen AI use case deployed	12%	McKinsey and IACPM, 2025	Industry

Regulated entities in India using or developing AI	20.8% (127 of 612)	RBI FREE-AI survey, Jan to May 2025	Independent

The gap between experimentation and scale is the central implementation problem. McKinsey's survey found that about half of senior leaders named gen AI a priority, but only a few, mainly larger institutions, were ahead, and most reported slower progress than expected. The constraint was missing skills, frameworks, and operating architecture rather than model capability. The Cambridge survey reports a deep execution gap between early experimentation and institution-wide integration.

3.2 Productivity and cycle-time effects

Table 2 summarises reported outcomes. The Blend results come from live production loans rather than simulations. Across 24 lender and loan-type cohorts covering more than 175,000 loans, Blend reports pull-through rates 10% to 15% higher, loan-cycle times shorter by 2 to 4 days, and about $600 saved in fulfilment costs per funded loan. Blend attributes this to an agent-first design that works a file in parallel, while most AI in lending speeds up individual steps in a process that still runs one step at a time. That is the vendor's own account of its product and has not been independently assessed.

Table 2: Reported productivity and cycle-time outcomes in origination and underwriting

Metric	Reported result	Context	Source	Type
Credit analyst productivity	20–60% gain	Credit memo automation at one bank	McKinsey and IACPM 2025	Industry
Decision speed	About 30% faster	Same case	McKinsey and IACPM 2025	Industry
Productivity across wholesale credit process	30–50%; output quality about 95%	Early implementations	McKinsey 2026, via ET BFSI	Industry
Loan cycle time	2–4 days shorter	Blend Autopilot, 175,000+ loans, 24 cohorts	Blend, Aug 2026	Vendor
Pull-through rate	10–15% higher	Same	Blend	Vendor
Fulfilment tasks automated	4.5 hours per loan	Same	Blend	Vendor
Fulfilment cost saving	About $600 per funded loan	Same	Blend	Vendor

Two points limit comparison across rows. The McKinsey credit memo figures describe one bank's workflow, and the 30% to 50% range comes from a later McKinsey analysis of wholesale lending. The vendor cohorts state no control group.

3.3 Approval rates and credit performance

Vendor deployments report that AI-enabled underwriting changes the distribution of decisions as well as their speed. Zest AI reports an average 25% increase in approvals with no added risk and a 20% reduction in defaults, and a 49% rise in approval rates for Latino borrowers. Trius FCU, a small credit union, automated 20% of consumer lending decisions and reports its delinquency rate falling from 1.72% to 1.00% over 14 months. These results are consistent with legacy underwriting having excluded some creditworthy borrowers, but they do not establish it. The cohorts lack a counterfactual for borrowers the old process declined, and delinquency can improve through portfolio mix, macroeconomic conditions, and servicing changes. The relevant test is the performance of approved borrowers who would previously have been declined, measured over a full credit cycle. Until that evidence is published, risk teams should treat the approval-rate effect as a hypothesis to test in their own portfolios.

3.4 Agentic orchestration and parallel decisioning

Harmoney, an Australia and New Zealand digital lender that processes about 12,000 applications a month, reports that average decision time fell from four days to 17 minutes and 40 seconds after it automated origination on Camunda's orchestration platform over roughly four years. A later phase added BankBot, an AI agent that handles complex income verification. Camunda reports that 60% to 80% of complex cases are now decided in under 15 minutes and that drop-off in that cohort fell by 20% to 30%. Two cautions apply. Camunda describes the headline result as roughly 90 times faster, but a 90-fold improvement implies a baseline of about 26 hours, not four days, and four days to 17 minutes would be more than 300 times faster. The two figures probably rest on different baselines or on the subset of applications that use the automated path, so the result should be cited with that qualification. Second, Camunda says its platform records every step, giving risk and compliance a complete record of each decision. That is a design property to test in any procurement, and it is not automatic.

4. STAGE 2: FRAUD DETECTION

4.1 The scale and composition of the threat

Fraud in lending has shifted from isolated document forgery to coordinated identity and document fraud. Fraudsters now assemble sets of documents, such as ID, proof of address, bank statement, and employer letter, designed to corroborate each other and pass reviews that check documents one at a time. The Cambridge survey adds a supervisory signal: 48% of respondents named adversarial AI a top concern. Javelin's 2026 study notes that flat identity fraud losses ($27.3 billion in 2025 against $27.2 billion in 2024) may mask shifts in fraud activity.

Table 3: Fraud loss and trend indicators in lending

Indicator	Reported metric	Source	Type
US lender exposure to synthetic identities	$3.3 billion (2024)	TransUnion	Industry
Traditional identity fraud losses, US (excluding scams)	$27.2 billion in 2024 (+19%); $27.3 billion in 2025	Javelin 2025 and 2026 studies	Industry
New-account fraud losses, US	$6.2 billion in 2024	Javelin 2025, via AARP	Industry
Forged or altered documents, share of fraud attempts	50% in 2024	Sphinx, citing Sumsub	Vendor
AI-generated document fraud	About 5× between April and December 2025	Inscribe, via Sphinx	Vendor
Synthetic identity document fraud, North America	+311%, Q1 2024 to Q1 2025	Sphinx	Vendor

4.2 Detection performance

Singh, Reddy, and Reddy (2025) describe a real-time contextual framework for consumer lending that combines graph neural networks and transformer-based anomaly detectors over 157 or more contextual signals. They report a 40% reduction in false positives, detection of 95% of sophisticated fraud within 300 milliseconds, and an AUC-PR of 0.92 on imbalanced data. These are results from a controlled evaluation, and AUC-PR on a curated dataset does not predict performance under adversarial drift. In lending, a false positive is a declined creditworthy borrower or a delayed disbursement, so reducing false positives improves cost and access together.

4.3 Embedding detection in the workflow

Successful deployments run fraud detection inside the decision workflow instead of as a separate review stage. When detection runs after underwriting, the institution has already paid to process a fraudulent application and has already delayed genuine applicants. Real-time architectures combine streaming data pipelines, low-latency feature engineering, and ensemble models in the origination flow, so fraud signals appear at the point of decision.

5. STAGE 3: PORTFOLIO MONITORING AND EARLY WARNING

5.1 The productivity evidence

McKinsey reports that banks are linking credit origination with ongoing borrower monitoring so that credit teams can spot deterioration earlier. Early implementations in wholesale lending show productivity improvements of 30% to 50% across the credit process, with output quality of about 95% (McKinsey, reported by ET BFSI on 11 November 2026). The same work describes event-triggered surveillance: tracking borrower filings, rating changes, market developments, and sector risks, and triggering reassessment of debt service capacity, covenant headroom, and facility limits. Automated scenario modelling lets institutions test macroeconomic and sector scenarios across borrowers and portfolios at a frequency manual review cannot match.

5.2 Calibrating automation to risk

Leading institutions are not automating uniformly. McKinsey describes a tiered model, shown in Table 4. Some European banks are extending automated credit assessment to higher exposures, but credit officers keep accountability for final approvals.

Table 4: Automation tiers in wholesale credit

Transaction type	Treatment
Standardised, lower-risk	Highly automated workflows
Moderately risky	Human-assisted processes
Complex, high exposure	Human-led assessment supported by AI-generated analysis

This calibration is a prudential design principle as well as a practical compromise. The 30% to 50% productivity gains come from concentrating automation where the marginal value of human judgement is lowest, and automating the complex tail would likely reduce risk quality.

5.3 The limits of early warning evidence

Published early warning results report model performance on historical datasets or vendor-curated cohorts. Evidence is missing on how early warnings arrive relative to covenant breaches, and on whether earlier intervention reduces loss given default in live portfolios. The claim that AI identifies deterioration before covenant breaches is plausible and mechanistically supported, but the causal chain from earlier signal to lower ultimate loss is not yet established at scale. This is a priority for empirical research.

6. STAGE 4: COLLECTIONS AND RECOVERY

6.1 From who to when

The dominant pattern in reported collections deployments is a shift from identifying who is at risk to choosing when and how to engage. A Nucleus Software blog describes South African university research in which forecasting a borrower's monthly cash flow let banks time outreach to income peaks, improving repayment and cutting collection costs by up to 25% in commercial portfolios. The blog does not name the research and it could not be located, so the figure is illustrative.

McKinsey reports that digitally active borrowers offered a practical digital route to tackle arrears are 20% to 30% more likely to pay than those reached by outbound calling alone.

6.2 Reported recovery outcomes

Garanti BBVA's collections unit was recovering less than half of overdue debt in the normal 8-to-91-day window. After deploying interaction analytics to identify customers who wanted to pay but had difficulty, and to make targeted restructuring or promise-to-pay offers, the recovery rate rose from 45% to 65% in the first year. The bank reports €800,000 of additional payments in the first full quarter and a projected €3 million annual benefit. The source is a vendor case study.

Table 5: Reported collections and recovery outcomes

Institution or context	Reported outcome	Source	Type
Garanti BBVA (Turkey)	Recovery rate 45% to 65% in first year; €800,000 extra in first quarter	NiCE case study, 2025	Vendor
Central Bank of India	SMA ratio 8% to 3%; recoveries up more than 60% in one year	Nucleus Software blog, Nov 2025	Vendor
South Africa, commercial portfolios	Collection costs down by up to 25% with income-timed outreach	Nucleus blog citing unnamed university research	Vendor
Digital arrears route	20–30% more likely to pay than outbound calling alone	McKinsey	Industry

6.3 The dignity dimension

Reported successes share a design feature that quantitative risk analysis often omits: respect for borrower circumstances. Garanti targeted customers who wanted to pay, the South African approach timed contact to income, and digital routes let borrowers act on their own schedule. These choices lower the risk of conduct failure and regulatory censure as well as raising recovery. A collections strategy that maximises recovery at the expense of borrower trust weakens future franchise value.

7. CROSS-CUTTING THEME: AGENTIC AI AND THE ORCHESTRATION LAYER

The main architectural development of 2025 to 2026 is the move from AI as a set of point solutions toward agentic AI as an orchestration layer. McKinsey describes banks testing agentic workflows across wholesale lending, covering retrieval of borrower data, financial analysis, credit policy checks, documentation, and exception management. Survey estimates of use vary with definition: the Cambridge survey found 52% of respondents experimenting with agentic AI, and an FT Longitude, SAS, and Microsoft survey of 300 credit, treasury, and risk leaders in 31 countries (June to July 2026) found 41% using agentic AI in decision-making.

Point-solution AI accelerates a step. Agentic orchestration can change the sequence. If each step in a sequential process takes 30% less time, cycle time falls 30%. If steps run in parallel, exceptions are routed adaptively, and the file is worked continuously, cycle time can fall further. Blend's 2 to 4 days of cycle-time reduction is an early, vendor-reported production measure of this effect. For risk teams, orchestration creates a design requirement: the audit trail must be built into the orchestration layer and not reconstructed afterward. Harmoney's case, as its vendor describes it, shows this is achievable. Institutions that attach agents to legacy workflows without orchestration-level logging will find that speed outpaces explainability.

8. GOVERNANCE, MODEL RISK, AND FAIR LENDING

8.1 The RBI FREE-AI framework as an operating model

The Reserve Bank of India's FREE-AI (Framework for Responsible and Ethical Enablement of AI) Committee Report, released on 13 August 2025, is a detailed supervisory blueprint for AI in lending outside the EU's horizontal AI Act. It sets out seven principles ("Sutras"), six pillars, and 26 recommendations covering innovation enablement and risk mitigation. These are committee recommendations and not yet binding regulation, but they signal supervisory direction. The RBI's surveys, fielded from January to May 2025, show a governance gap (Table 6). Of 612 regulated entities, 20.8% used or were developing AI, and 67% expressed interest in exploring use cases. Among the 127 AI users, 18% maintained audit logs and 35% validated models for bias and fairness, mostly before deployment. Credit underwriting accounted for 13.7% of 583 reported AI applications, and about one-third of respondents had a board-level AI oversight framework.

Table 6: RBI FREE-AI survey, deployment and governance indicators

Indicator	Reported figure	Basis
Entities using or developing AI	20.8%	127 of 612 entities
Entities interested in exploring AI use cases	67%	Surveyed entities
AI users maintaining audit logs	18%	127 AI users
AI users validating for bias and fairness	35%	127 AI users
Credit underwriting share of AI applications	13.7%	583 applications
Entities with board-level AI oversight	About one-third	Respondents

Source: RBI FREE-AI Committee Report (2025). All indicators are independent.

Most entities that use AI lack the audit and fairness controls supervisors expect. The RBI warns of the cost of falling behind, and the survey data point to the opposite risk as well: deploying AI faster than governance can absorb it.

8.2 Fair lending and regulatory divergence

Fair lending exposure has changed shape. In the United States, the CFPB withdrew its algorithmic adverse-action circulars on 12 May 2025. On 22 April 2026 it finalised amendments to Regulation B, effective 21 July 2026, that remove the effects test and state that the Equal Credit Opportunity Act does not authorise disparate-impact liability. Disparate treatment remains prohibited, including intentional use of neutral criteria as proxies for protected characteristics. The rule does not end disparate-impact exposure. It remains available under the Fair Housing Act for mortgage lending. State attorneys general continue to act: in July 2025 the Massachusetts Attorney General announced a $2.5 million settlement with Earnest Operations LLC over AI-driven student loan discrimination. Law firms advising lenders also expect litigation over whether the rule is consistent with the statute, and some state requirements conflict with the federal position.

In the European Union, credit scoring is a high-risk use under Annex III of the AI Act. The Digital Omnibus on AI, published in the Official Journal on 24 July 2026, moved the application date for standalone Annex III high-risk obligations from 2 August 2026 to 2 December 2027. Obligations on general-purpose AI model providers and most transparency obligations were not deferred.

Fair lending risk has therefore not receded. The enforcement surface has moved from federal guidance to state attorneys general, private litigation, the Fair Housing Act, and EU conformity requirements. Institutions that treat bias testing as a compliance checkbox, or that drop disparate-impact testing because of the federal change, are exposed in every jurisdiction where it still applies.

8.3 Explainability as a precondition of supervision

A model that cannot give a reason for an adverse decision cannot be supervised, defended in a dispute, or improved through feedback. FREE-AI's "understandable by design" principle reflects this. Explanation techniques such as SHAP, LIME, and counterfactual explanations are necessary but insufficient: the explanation must be meaningful to the borrower and stable across the model's lifecycle. Drift detection, automated bias monitoring, and continuous model inventory are core lending infrastructure and belong outside the category of model risk add-ons.

9. IMPLEMENTATION ROADMAP AND RETURN ON INVESTMENT

9.1 The Productionisation Gap

Across the evidence, the gap between pilot and production is the binding constraint on value. McKinsey finds that most institutions focus narrowly on simple use cases instead of end-to-end workflows, for lack of skills, frameworks, and operating architecture. The Cambridge survey reaches a similar conclusion. The five-stage sequence below is this paper's synthesis of the McKinsey and deployment evidence:

1. **Align stakeholders and build capability.** Internal investment in AI literacy and early external engagement reduce inertia and build the trust that higher-stakes use cases need.
2. **Standardise data.** Unify structured and unstructured data in an operational layer that supports end-to-end experimentation. Cambridge respondents name data quality as the leading obstacle to adoption.

3. **Adopt a modular architecture.** Interchangeable components allow parallel use-case development and avoid monolithic systems that cannot adapt to model or regulatory change.
4. **Start with low-risk use cases.** Internal bots, document processing, and credit memo drafting deliver measurable early results before borrower-facing decisions are automated.
5. **Roll out agentic orchestration with the audit trail built in.** Move from static applications to orchestrating agents across the journey, using the logging and governance described in Section 7.

9.2 ROI evidence and its limits

ROI evidence is limited and mostly vendor-reported (Table 7). Three caveats apply. Vendor figures typically measure gross efficiency gains and not returns net of implementation, integration, and governance costs. The independent Cambridge survey is more cautious: 40% of respondents reported higher profitability, 43% reported no change, and 76% of large institutions found AI value hard to measure. Finally, institutions that treat AI as capability building, and not as a series of pilots justified quarter by quarter, are better placed to capture durable value. That last point is this paper's judgement and not a finding from the sources.

Table 7: Reported ROI and cost-saving metrics

Metric	Reported value	Source	Type
Fulfilment saving per funded loan	About $600	Blend	Vendor
Respondents reporting higher profitability from AI	40% (43% no change)	Cambridge CCAF 2026	Independent

10. CRITICAL REFLECTIONS: WHAT THE EVIDENCE DOES NOT YET SUPPORT

A paper for a risk institute should be as clear about the limits of the evidence as about its strengths. Five limits matter most.

First, the early warning causal chain is incomplete. Productivity and data-coverage claims for AI-enabled monitoring are well supported. The claim that earlier warning lowers ultimate credit losses is plausible but has not been shown at portfolio scale with out-of-sample validation.

Second, fraud benchmarks come from controlled evaluations. AUC-PR figures and false-positive reductions are measured on curated data. Fraud is a dynamic adversarial problem, and a model that performs well on 2024 data may degrade as generative tools evolve. Institutions should invest in drift detection and adversarial robustness testing at least as heavily as in accuracy.

Third, collections evidence is dominated by vendor-published successes. Garanti BBVA and Central Bank of India are reported successes, and the base rate of collections projects that fail to improve recovery is unknown because failures are rarely published. Published recovery improvements are best treated as upper bounds.

Fourth, the governance data are weak. In the RBI survey, 18% of AI-using entities maintained audit logs and 35% validated for bias. The survey covers India in early 2025 and includes many small cooperative banks, so it should not be read as a global estimate. Even so, it suggests that most surveyed AI deployments could not be audited to the standard prudential supervision requires, which makes the implementation gap a supervisory exposure as well as a value gap.

Fifth, source quality and definitions limit comparison. Adoption figures range from 12% to more than 80% depending on the survey and the definition used. Most lifecycle outcome figures are vendor-reported, and several figures in circulation could not be traced to primary sources while this paper was prepared, so they were excluded. For any AI lending result, readers should ask who measured it, against what baseline, over what period, and whether failure rates were reported.

11. CONCLUSION: FIVE TAKEAWAYS FOR LOAN AND RISK PROFESSIONALS

1. **AI works best as an orchestration layer.** The largest reported gains, such as Harmoney's drop in decision time from days to minutes and 30% to 50% productivity gains across wholesale credit workflows, come from redesigning the sequence of decisions and not from speeding up isolated steps. Most of these figures are vendor- or consultancy-reported.
2. **Credit risk and fraud risk are converging.** Synthetic identities, AI-generated documents, and coordinated fraud packages mean fraud detection belongs inside the underwriting workflow. The signals that detect fraud can also improve credit decisions.
3. **Collections is an emerging area of credit strategy.** Garanti BBVA reports a 20-point rise in recovery rate. Timing and segmentation matter more than aggressive outreach, and reported results are best treated as upper bounds.
4. **Governance speeds deployment when it is designed in.** FREE-AI offers a workable operating model for explainable, accountable, fair AI. Institutions that build audit trails, bias validation, and model risk management into the orchestration layer from the start should scale with less supervisory friction. Regulatory divergence between US federal rules, US states, and the EU makes that design more important.
5. **The implementation gap is the main risk.** More than 80% of firms use AI at some level, yet 55% of industry respondents find its value hard to measure and the governance indicators lag deployment. Over the next 24 months the central challenge is productionisation with accountability: AI governed like credit risk, measured like operational performance, and anchored in borrower fairness.

REFERENCES

1. McKinsey & Company and IACPM, Banking on gen AI in the credit business: The route to value creation, July 2025. Survey of 44 financial institutions.

2. McKinsey & Company, AI-enabled risk operating model analysis (30–50% productivity in wholesale credit, about 95% output quality), reported in ET BFSI, 11 November 2026.

3. McKinsey & Company, Banking on AI: How a new operating model can transform distressed credit management (digital arrears route, 20–30% more likely to pay).

4. Cambridge Centre for Alternative Finance, Cambridge Judge Business School, 2026 Global AI in Financial Services Report: Adoption, Impact and Risks, April 2026. 628 respondent organisations.

5. FT Longitude, SAS and Microsoft, survey of 300 credit, treasury and risk leaders at banks in 31 countries, June to July 2026, as reported by Fintech News Singapore.

6. The Mortgage Collabourative, Pulse of the Network survey, June 2026, as reported by National Mortgage News.

7. Jifiti, The Loan Origination Technology Gap Is Now a Revenue Problem for Banks, 7 May 2026.

8. Blend Labs, Early Production Results for Blend's Autopilot Show What Agentic AI Means For Lending, 20 August 2026.

9. Camunda, 90x Faster Lending Decisions: How Harmoney Scaled Loan Approvals With Agentic Orchestration, 15 May 2026.

10. Zest AI, Trius FCU Customer Success Story, 2025.

11. Nucleus Software, Reducing NPLs in Emerging Markets with AI-Powered Collections, 18 November 2025.

12. NiCE, Case Study: Garanti BBVA, 2025.

13. Singh, J., Reddy, V., and Reddy, D., "Real-time Contextual AI for Proactive Fraud Detection in Consumer Lending: Architectures, Algorithms, and Operational Challenges," Journal of Information Systems Engineering and Management, 2025.

14. Sphinx HQ, Document Fraud in Banking: How to Detect Fake Documents, 30 July 2026. Cites Sumsub and Inscribe data.

15. Javelin Strategy & Research, 2025 Identity Fraud Study: Breaking Barriers to Innovation, 25 March 2025, and 2026 Identity Fraud Study, April 2026.

16. TransUnion, synthetic identity exposure for US lenders ($3.3 billion), as cited in industry compilations.

17. Reserve Bank of India, Framework for Responsible and Ethical Enablement of Artificial Intelligence (FREE-AI) Committee Report, 13 August 2025.

18. Consumer Financial Protection Bureau, final rule amending Regulation B (Equal Credit Opportunity Act), issued 22 April 2026, effective 21 July 2026.

19. Massachusetts Attorney General, settlement with Earnest Operations LLC over AI-driven student loan discrimination, July 2025, $2.5 million.

20. European Union, Digital Omnibus on AI, published in the Official Journal of the European Union, 24 July 2026, amending Regulation (EU) 2024/1689.

This paper is prepared for presentation at the Annual Risk Conference of the Chartered Institute of Loan and Risk Management of Nigeria. The analysis rests on publicly available data and published case evidence as of 18 November 2026. Institution-level results are vendor-reported unless otherwise indicated, and readers should interpret them with caution.

Reviewer Notes (remove before submission)

This page records what was changed in the paper and why. It is not part of the paper. Items marked "Action for you" need a decision or a check before submission.

Corrections and removals

#	Where	Finding	Action taken
1	Abstract, Table 1	The headline 79% (up from 34%) came from an unnamed "market study" and traces only to a trade article. No source in the reference list.	Replaced with the Cambridge CCAF 2026 survey (named, n=628). Added definition spread (12% to 80%+).
2	Table 5, Section 6, Conclusion	Punjab National Bank "farm-loan NPAs reduced to 0.4%, RBI-reported" traces to a Nucleus Software marketing blog. PNB's reported net NPA ratio was 0.40% at March 2025, which looks like the real origin. Agricultural NPA ratios across Indian banks are far higher. The crop-harvest repayment claim has the same source.	Removed from text, table, and conclusion. Action for you: do not reinstate without a primary RBI or PNB source.
3	Section 3.4	Harmoney "90× faster" conflicts with its own endpoints. Four days to 17 min 40 s is over 300×; 90× implies a baseline near 26 hours. Camunda also says the improvement came over about four years.	Reworded with the inconsistency stated. Action for you: ask Camunda or Harmoney for the baseline definition if you want to keep 90×.
4	Section 9.2	Text said Freddie Mac saves "up to 600" per loan, mixing a Freddie Mac $1,500 figure with Blend's $600.	Freddie Mac figure removed (see item 12); Blend $600 retained and labelled vendor-reported.
5	Tables 3 and 7, Section 7	Unsourced or untraceable figures: 72% rule-based failure rate; 67% fewer false positives (also conflicts with the 40% figure); synthetic identity 11% of global fraud; deepfake +3,900%; 7.2-month "calculated" payback; >$3 million credit memo savings; 70% of banks using agentic AI; "contact bounce" down 25–30%; >60% QuantumBlack estimate.	All removed. The 70% agentic claim was replaced with CCAF (52%) and FT Longitude/SAS/Microsoft (41%).

6	Table 3, Ref 15	$27.2 billion is Javelin's figure for traditional identity fraud excluding scams (total with scams was $47 billion in 2024). The paper called it "overall". Ref 15 also credited the $3.3 billion synthetic exposure to Javelin; it is TransUnion's.	Relabelled; added Javelin's 2025 figure ($27.3 billion). Attribution fixed.
7	Sections 3.2, 5.1	The 30–50% productivity and ~95% quality figures come from a later McKinsey risk-operating-model analysis, not the July 2025 44-institution survey. The 20–60% and ~30% figures describe one bank's credit memo case.	Sources separated in text and Table 2.
8	Section 8.1, Table 6	20.8% means "using or developing" AI (127 of 612). The 18% and 35% are shares of the 127 AI users. Credit underwriting 13.7% is of 583 applications. One summary gives 15% for audit logs; that is interpretability tools. FREE-AI is recommendations, not binding rules.	Denominators and status stated. Action for you: re-check against the RBI PDF; I verified through KPMG, Khaitan & Co and press summaries.
9	Section 8.2	Out of date. The CFPB finalised the Regulation B rule on 22 April 2026 (effective 21 July 2026), going beyond the May 2025 guidance withdrawal. The EU AI Act credit scoring obligations moved to 2 December 2027.	Section rewritten with both updates.
10	Sections 3.3, 6.3, 9.2	Causal and sweeping claims not supported by the sources: manual underwriting was "systematically conservative"; collections AI cuts conduct risk; ROI abandonment claim; "most enthusiastic institutions least prepared".	Reworded as hypotheses or the paper's judgement, or removed.
11	Section 6.1	South Africa 25% cost reduction cites unnamed "university research" via a vendor blog.	Retained as illustrative and labelled Vendor. Action for you: locate the study or drop it.
12	Tables 1, 2, 3, 7	Figures that could not be traced to a primary source: IDC 2.3× ROI, McKinsey 2.8× ROI, 45–60% labour cost reduction, Freddie Mac $1,500 saving per loan, Fannie Mae 55%	All removed at your request, along with the Stealth Agents reference.

		projection with the 2023 and 2024 mortgage adoption figures, FDIC $6.8 billion synthetic identity losses.	
13	Throughout	Mixed UK and US spelling. Em dashes removed.	Standardised to UK spelling at your request.
14	Title block	Institute name and author needed confirming.	Set to Chartered Institute of Loan and Risk Management of Nigeria; author Dr. Chinedu Okeke.
15	Tables	Source quality was invisible to the reader.	Added a Type column (Independent, Industry, Vendor).

Verification status

Checked against sources: RBI survey figures (through secondary summaries), McKinsey 20–60% and ~30% credit memo case, McKinsey 30–50% and ~95%, Harmoney figures, Javelin 2024 and 2025 figures, CFPB Regulation B rule and dates, EU Digital Omnibus dates, Cambridge CCAF figures, PNB NPA figures.

Not independently checked: Blend production figures, Garanti BBVA and NiCE figures, Trius FCU and Zest AI figures, Singh et al. results, Sphinx and Inscribe fraud figures, TransUnion $3.3 billion (confirmed only in secondary compilations), CFPB 12 May 2025 withdrawal date, Massachusetts Earnest settlement details, FT Longitude survey details.