

AI-Enabled Agricultural Finance: Reimagining Asset Management and Sustainable Wealth Creation in Rural India

Akshathraj Jain

Asst. Professor

Department of Management Studies

Mangalore Institute of Technology & Engineering Mijar Moodbidri

email: akshathraj@mite.ac.in

Dr. Anjali Ganesh

Professor

Department of Business Administration

St Joseph Engineering College

Mangalore, Karnataka, India

anjali@sjec.ac.in

Abstract

Smallholder farmers in the Global South face persistent barriers in accessing formal credit, limiting their ability to adopt climate-resilient and environmentally sustainable agricultural practices. Conventional agricultural credit systems often struggle to assess the creditworthiness of smallholders due to the lack of formal credit histories and high transaction costs, resulting in widespread financial exclusion.

This paper synthesises empirical evidence on the developmental impacts of Financial Technology (FinTech) adoption, drawing on multi-stage and stratified survey studies. Using insights from Structural Equation Modelling (SEM) and related econometric approaches, the paper examines how digital financial services—including mobile payments, microloans, and digital crop insurance—contribute to rural economic resilience and influence farm-level environmental decision-making.

The analysis indicates a significant pathway through which FinTech adoption supports formal financial inclusion, which in turn acts as a mediator for transitions toward more sustainable farming practices. The paper offers policy-relevant insights for designing inclusive, context-sensitive digital financial ecosystems aligned with the United Nations Sustainable Development Goals (SDGs).

FinTech Adoption; Financial Inclusion; Sustainable Agriculture; Technology Acceptance Model (TAM); Digital Literacy; SDGs

Introduction

Indian agriculture is characterised by large numbers of farmers operating on very small plots. The Agriculture Census 2015–16 recorded approximately 146 million operational holdings, of which around 86 per cent were small or marginal, with an average holding size of 1.08 hectares (Ministry of Agriculture and Farmers Welfare, 2015–16). Incomes on such holdings are seasonal, exposed to weather and price shocks, and typically insufficient to build savings without external support.

Institutional credit has expanded substantially in recent years. By December 2024, there were 7.72 crore operative Kisan Credit Card (KCC) accounts with an outstanding amount of Rs 10.05 lakh crore (Ministry of Finance, 2025). However, a loan finances a crop cycle; it does not by itself create a balance sheet. A farm household may be “included” in the credit system yet hold few productive assets, limited insurance cover, and no liquid buffer, thereby cycling through debt from one season to the next.

This paper originates from a conference theme on technology-enabled rural development and sustainable finance. The core idea; shifting the focus from loan disbursement to household asset formation—emerged from observing that agricultural finance discussions often emphasise credit volume and repayment, while giving less attention to whether farm households accumulate productive assets and build resilience over time.

Three developments make it timely to revisit the design of agricultural finance. First, India is building digital public infrastructure for agriculture. Under AgriStack, more than 10.31 crore Farmer IDs had been created by 3 August 2026, and the Digital Crop Survey covered 648 districts and over 31.3 crore plots in the 2025–26 Rabi season (Ministry of Agriculture and Farmers Welfare, 2026). Second, the regulator has revised the rules for the main farm-credit product: the Reserve Bank of India's KCC Directions, 2026, effective 1 January 2027, introduce a six-year composite facility, a Flexi KCC of Rs 10,000–50,000 for marginal farmers, a collateral waiver up to Rs 2 lakh, and standardised Scale of Finance norms (Reserve Bank of India, 2026a). Third, the regulator has articulated principles for AI in finance through the FREE-AI framework, with seven guiding principles and 26 recommendations (Reserve Bank of India, 2025).

Data, rules, and analytical capability are therefore converging. What remains underdeveloped is a coherent design that turns these elements into wealth creation rather than merely faster lending.

This paper proposes three shifts in how AI-enabled agricultural finance is conceived. The first is from a loan-centred to a balance-sheet-centred view, in which financial assets, productive assets, natural capital, and liabilities of the farm household are considered together. The second is from one-time underwriting to continuous, consent-based monitoring, in which data flow from registries, satellites, and payment trails remains under the farmer's control. The third is from gatekeeping to empowerment, in which a risk flag triggers advice, insurance, or staged disbursement rather than a bare refusal.

The paper is guided by three research questions:

- **RQ1.** What does the existing evidence show about the effect of FinTech and AI on access, sustainable practices, and lending decisions for smallholders, and how reliable is that evidence?
- **RQ2.** Which technical and governance pitfalls determine whether AI-based credit decisions help or harm smallholder farmers?
- **RQ3.** What architecture and phased roadmap can convert India's digital rails into asset management and durable wealth creation at the household, collective, and institutional levels?

The paper makes four contributions:

1. It appraises a small but diverse evidence base with attention to internal consistency.
2. It re-derives the reported performance of a widely accessible loan-acceptance classifier to show how imbalance can mislead.
3. It proposes a layered framework with seven testable propositions.
4. It offers a roadmap and scorecard that lenders and policymakers can adapt.

Section 2 reviews the literature, Section 3 describes the approach, Section 4 assesses India's enabling infrastructure, Section 5 presents the framework, Section 6 reports the analysis, Section 7 sets out the roadmap and implications, Section 8 discusses risks and limitations, Section 9 outlines future research, and Section 10 concludes.

Background and Literature Review

Theoretical anchors

Two strands of theory frame the argument. Asset-based approaches to welfare hold that durable asset holding, not income alone, shapes household security and long-term outlook (Sherraden, 1991), while the livelihoods

literature treats natural, physical, financial, human, and social capital as interacting stocks that households build and draw down (Chambers & Conway, 1992). Both imply that the test of agricultural finance is whether it raises the household's net asset position over several seasons.

On the technology side, the Technology Acceptance Model attributes adoption chiefly to perceived usefulness and ease of use (Davis, 1989), and diffusion theory adds relative advantage, compatibility, and social networks (Rogers, 2003). Both are used in the farm-level studies reviewed below.

FinTech, financial inclusion and sustainable practices: evidence from southern India

Two recent studies examine FinTech among smallholders in southern India. Rathi and Sahoo (2026) surveyed 395 smallholders across eight districts of South Tamil Nadu and estimated a PLS-SEM that extends the Technology Acceptance Model with financial inclusion and sustainable agricultural practices as outcomes. They report that FinTech adoption is positively associated with financial inclusion ($\beta = 0.612$), that financial inclusion is associated with sustainable practices ($\beta = 0.548$), and that the indirect path through financial inclusion accounts for 55.6 per cent of the total effect (partial mediation). Perceived usefulness was the strongest predictor of behavioural intention ($\beta = 0.295$), ahead of ease of use ($\beta = 0.146$) and trust ($\beta = 0.099$). The practical reading is that farmers adopt digital finance when it demonstrably helps them, and tolerate moderate complexity if it does.

Vasudevan et al. (2025) surveyed 670 farmers using stratified random sampling and drew on diffusion-of-innovation theory, with digital literacy as a moderator. They conclude that FinTech supports access to credit, market information, and resilience to climate and market shocks; that effectiveness depends on infrastructure, legislative support, and digital literacy; and that the hypothesis that FinTech reduces the carbon footprint of farming was not supported. This last result matters for policy: digital finance does not automatically make farming greener.

Both studies are cross-sectional and self-reported, and neither observes assets, wealth, or repayment over time. Section 6.1 shows that some of the reported statistics are internally inconsistent, which argues for reading them as indicative of direction rather than magnitude.

Machine learning for agricultural lending

Goswami et al. (2026) provide the most comprehensive synthesis: a PRISMA-based systematic review of 214 studies on machine-learning and AI approaches to agricultural loan recommendation. They organise the field along a pipeline of feature specification, field-data acquisition, and model design. Determinants of loan decisions span farmer-centred factors (gender, household, education, experience), financial factors, geographic factors (location, terrain, soil, yield), and environmental factors (natural disasters, climate change). Data acquisition options range from wireless sensor networks and IoT devices to UAVs and satellite imagery.

On modelling, the review observes that loan datasets are usually tabular and lack sequential structure, so simpler statistical and machine-learning models may be more suitable than deep networks, whereas deep learning excels on imagery from remote sensing. It cites examples relevant to India, including the Kisan Rin Portal, the YES-TECH and CROPIC initiatives for yield estimation, a satellite-based underwriting product that cut assessment time from weeks to minutes, and multilingual assistants that guide farmers through schemes and documentation.

The review is equally candid about limitations. Much of the evidence is proof-of-concept; large-scale deployment studies in real smallholder settings are scarce; ground-truth data from marginalised and climate-vulnerable communities are limited; fairness auditing is underused; and models trained on data from well-served regions can create feedback loops in which underserved areas become algorithmic blind spots. It recommends treating loan recommenders as empowerment mechanisms, for example by pairing a high-risk flag with agronomic advice or subsidised inputs rather than denying credit outright, and by using cooperatives as ground-truth validators and credit aggregators.

Akça and Sevlı (2022) compare four support vector machine kernels for predicting whether bank customers will accept a personal-loan offer, using a public dataset of 5,000 customers in which 9.6 per cent accepted. They report the best accuracy for the polynomial kernel (97.2 per cent) and the lowest for the sigmoid kernel (83.3 per cent), and acknowledge that precision and recall are depressed by class imbalance. The study is not agricultural, and its postal-code field follows a five-digit format that differs from India's six-digit PIN codes, so transferability to rural India is unproven. Its relevance here is methodological: product take-up prediction (which farmers will respond to a savings, insurance, or investment offer) is a natural use of AI in an asset-management setting, and the evaluation issues it raises are general.

Fraud, identity and collateral integrity

Ajah and Inyama (2011) analyse loan fraud and non-performing loans in Nigeria and identify structural causes: weak identification, no shared central record of borrowers' indebtedness, inability to detect fake collateral or the same collateral pledged to several lenders, unreported bad loans that encourage multiple lending, and subjective loan-officer discretion. They propose a comprehensive IT system combining biometric identity, a national identity register, GPS/GIS capture of collateral coordinates, neural-network scoring, intelligent agents, and data mining.

The paper is conceptual and predates present-day digital public infrastructure, and its proposal that any borrower lacking a national identity number be refused a loan illustrates a risk that this paper takes seriously: a hard identity gate can exclude the very farmers, such as tenants and sharecroppers, that inclusion policy seeks to reach. Its typology of integrity risks nonetheless remains a useful checklist, and we revisit it in Section 6.3.

Governance, fairness and explainability

Goswami et al. (2026) note that credit-scoring systems are treated as high-risk under the EU AI Act and describe three fairness criteria—independence, separation, and sufficiency—alongside interpretability tools such as SHAP. In India, the FREE-AI report sets seven “sutras” (trust, people first, innovation over restraint, fairness and equity, accountability, understandability by design, and safety, resilience and sustainability) and six pillars (infrastructure, policy, capacity, governance, protection, and assurance). Commentary on the report notes that fewer than a quarter of surveyed regulated entities were deploying AI, concentrated in larger banks and non-banking finance companies, with little adoption among co-operative banks (Mondaq, 2025).

Because co-operative banks handle a large share of farm credit—just over half of the roughly 1,999 lakh KCC applications processed on the national platform came through them (United News of India, n.d.)—the readiness gap sits where the volume is. Equality-of-opportunity criteria (Hardt et al., 2016) and model-agnostic explanations (Lundberg & Lee, 2017) provide the technical vocabulary for the audits the framework requires.

Research gaps

Six gaps follow:

- None of the reviewed studies measures household wealth or asset formation over time.
- The Indian farm-level evidence is cross-sectional and self-reported.
- Lender-side evidence, including cooperative-bank readiness, is thin.
- Evaluation practice for imbalanced credit data is weak.
- No integrated architecture links registry data, scoring, insurance, and collective asset management.
- Tenants and women are under-addressed in model design and evaluation.
- The framework in Section 5 is designed to speak to each of these gaps.

Research Design and Approach

The study is an integrative, theory-building review with two analytical components. The evidence base consists of five studies provided for the review (Table 1) and a set of official Indian sources accessed in September 2026, namely parliamentary replies, ministry statements, and RBI directions and reports. Each study was appraised on design, sample and sampling, measurement, internal consistency of reported statistics, and transferability to smallholder finance in India.

One further article on satellite data and AI for agricultural disaster financing (Chen et al., 2022) was supplied but carries a retraction notice on every page. It was excluded from the evidence base, and none of its claims is used. It is listed in the references only for transparency.

The first analytical component is a consistency check of the reported path estimates in Rathi and Sahoo (2026): for each path we recomputed the t-statistic as β/SE , the implied two-tailed p-value, and the 95 per cent confidence interval ($\beta \pm 1.96 SE$) and compared them with the reported values. The second is a metric re-derivation for Akça and Sevli (2022): the reported confusion matrices are means over five cross-validation folds, each with 1,000 test instances, so accuracy, precision, recall, F1, specificity, balanced accuracy, and the Matthews correlation coefficient (Chicco & Jurman, 2020) can be recomputed and compared with the majority-class baseline. The framework and propositions are then developed by mapping the evidence and the enabling infrastructure onto the three research questions.

India's Enabling Infrastructure for AI-Enabled Agricultural Finance

A framework is only as credible as the infrastructure it can stand on. Table 1 summarises the components that matter most, their status as reported in official sources, and the residual gap that AI-enabled design must not paper over. Figures are as reported at the dates shown and are likely to have moved.

Table 1. Infrastructure readiness for AI-enabled agricultural finance

Component	Status (as reported)	Residual gap relevant to design
Farmer and land identity	AgriStack Farmer Registry: over 10.31 crore Farmer IDs at 3 Aug 2026; target of 11 crore by 2026–27; federated, state-owned data; consent manager and unified service interface planned as part of the design	Coverage of tenants and sharecroppers; quality of underlying land records; duplicate or stale entries
Crop-sown data	Digital Crop Survey: 648 districts and 31.3 crore plots in Rabi 2025–26	Accuracy of survey entries; need for satellite validation before use in scoring
Core credit product	KCC: 7.72 crore operative accounts, Rs 10.05 lakh crore (Dec 2024); RBI KCC Directions 2026 effective 1 Jan 2027 (composite six-year facility, Flexi KCC, collateral waiver up to Rs 2 lakh, Scale of Finance)	Small average ticket sizes; conversion of working capital into asset-building investment credit; last-mile digitisation of co-operative banks
ligibility and inclusion rules	RBI KCC Directions for regional rural banks list owner cultivators, tenant farmers, oral	Documentation of tenancy; women's land rights; group-lending data


Component	Status (as reported)	Residual gap relevant to design
	lessees, sharecroppers, and self-help or joint liability groups as eligible	
AI governance	RBI FREE-AI report (13 Aug 2025): seven sutras, six pillars, 26 recommendations	Uneven adoption capacity across lenders; audit skills in smaller institutions
Advisory layer	Multilingual AI advisory announced as part of AgriStack; assistants such as Jugalbandi and Farmer.CHAT described in the literature	Accuracy, safety, and accountability of generated advice; local-language coverage

A Framework for AI-Enabled Asset Management and Wealth Creation

Architecture

Figure 1 presents the framework as four layers under a cross-cutting layer of governance and trust, with a feedback loop from farmers and cooperatives back to the data foundation.

Figure 1. Conceptual framework for AI-enabled asset management and wealth creation in rural India



Layer 1 – Data foundation. The base layer assembles consented data from farmer and land registries, crop-sown records, satellite, weather and soil signals, payment trails, and credit-bureau information. Goswami et al. (2026) note that in sensor-poor areas proxy features such as vegetation indices from open satellite data must substitute for field sensors, which introduces structural variation in data quality that models must respect. Data minimisation and provenance tracking are design requirements, not afterthoughts.

Layer 2 – Intelligence. Four families of models sit here: credit scoring and loan recommendation; yield and climate-risk estimation; fraud, duplicate-pledge, and identity checks; and multilingual advisory assistants. For tabular credit data, interpretable models with calibrated probabilities should be the default, with deep learning reserved for imagery (Goswami et al., 2026). Every score should be accompanied by an explanation in plain language and a measure of uncertainty that routes ambiguous cases to human review.

Layer 3 – Financial and asset-management services. Products include working-capital and investment credit, index-based crop insurance, savings and warehouse-receipt finance, and tools for cooperatives and self-help

groups to track shared assets and portfolios. The layer's distinctive feature is sequencing: the same intelligence that prices credit also proposes what to protect (insurance), what to store (savings), and what to build (investment) so that finance strengthens the balance sheet.

Layer 4 – Outcomes.

Outcomes are defined at multiple levels: household income stability and resilience, productive-asset formation, adoption of sustainable practices, and institutional portfolio quality.

Governance and trust. This layer applies to all others and comprises consent and data rights, fairness audits and explainability, human-in-the-loop review, grievance redress, and cyber and model-risk controls, consistent with the FREE-AI sutras of trust, people first, fairness, accountability, and understandability.

Ground-truth feedback. Farmers and cooperatives are treated not merely as data subjects but as validators who confirm or contest yield estimates and crop-failure flags, improving retraining and localisation in low-data regions (Goswami et al., 2026).

Reimagining asset management at three levels

Household. The unit of analysis becomes the farm-household balance sheet: financial assets (deposits, insurance cover), productive assets (irrigation, equipment, livestock, storage), natural capital (soil health, water access), and liabilities. With the farmer's consent, AI can assemble a view of this balance sheet from registry, payment, and bureau data and support a sequence of decisions—liquidity buffer, risk cover, productivity-raising investment, savings—under human advisory oversight.

We define sustainable wealth creation as a sustained increase in the net worth of the farm household, including productive and natural capital, over multiple seasons, achieved without rising over-indebtedness. This is distinct from growth in loan volume or in a single season's income.

Collective

Farmer producer organisations (FPOs), self-help groups, and cooperatives can act as data intermediaries and portfolio aggregators. Pooling validated yield and practice records allows collectives to negotiate portfolio-level credit, finance shared assets such as storage or machinery, and lower the per-borrower cost of monitoring (Goswami et al., 2026). Collectives also offer a route to include tenants and women whose individual documentation is thin.

Institutional

For lenders, asset management means portfolio surveillance: satellite- and weather-based early warning, geographic concentration analysis, climate stress-testing, and continuous fairness monitoring.

Institutions that do this can lend more to small farmers without accumulating hidden risk, because risk is observed during the season rather than after default.

From gatekeeping to empowerment: a three-route decision policy

We propose that AI-supported credit decisions route each application to one of three outcomes: approve; approve with support, meaning staged disbursement, an insurance bundle, agronomic advice, or input subsidy where the model flags elevated risk; or refer for human review, used when uncertainty is high, when the applicant is a first-time borrower, or when an adverse decision is being considered. Automated denial without human review is avoided, in line with the "people first" principle. This design follows the recommendation in Goswami et al. (2026) that recommenders should work as empowerment tools rather than gatekeepers.

Propositions

The framework yields seven propositions (Table 2). They are theory-building statements to be tested, not findings of this paper.

Table 2. Propositions, evidential basis and testable indicators

No.	Proposition	Evidential basis	Testable indicator
P1	Consent-based integration of registry, crop-sown, and payment data raises the coverage of thin-file farmers more than added model complexity does.	Goswami et al. (2026); infrastructure in Section 4	Share of registered farmers with a scored or pre-approved limit, by model version and data source
P2	For tabular credit data, interpretable models with calibrated probabilities come close enough to complex models that data quality and governance dominate performance differences.	Goswami et al. (2026), Section 5.7	Head-to-head calibration and cost-weighted loss on the same lender data
P3	Evaluations that report only accuracy on imbalanced data overstate decision value; balanced, cost-weighted, and group-wise metrics are needed.	Akça & Sevlı (2022), re-derived in Section 6.2	Gap between accuracy and balanced accuracy or MCC; group-wise error rates
P4	Routing risk flags to advice, insurance, or staged disbursement raises net inclusion and repayment relative to binary denial.	Goswami et al. (2026), Section 6.3.1	Approval, repayment, and default in matched pilots (cluster-randomised where feasible)
No.	Proposition	Evidential basis	Testable indicator
P5	Digital access to finance raises sustainable-practice adoption through investment capacity, conditional on digital literacy, connectivity, and extension support, but does not by itself reduce emissions.	Rathi & Sahoo (2026); Vasudevan et al. (2025)	Change in verified practices (e.g., water-efficient irrigation) by literacy and connectivity strata
P6	FPO- and SHG-mediated data pooling lowers monitoring cost and increases the effective inclusion of tenants and women relative to individual digital onboarding alone.	Goswami et al. (2026), Section 6.3.2	Inclusion rates and per-borrower monitoring cost, collective vs. individual channel
P7	Cross-source verification of identity, parcel, and crop declarations reduces duplicate-pledge and misdeclaration losses, whereas single-identifier gating increases exclusion.	Ajah & Inyama (2011); RBI eligibility rules	Fraud and misdeclaration loss rates; rejection rate for applicants lacking a single identifier

Analysis

What the farm-level evidence supports

Taken together, the two southern-India surveys converge on four points. First, digital finance is associated with wider access to credit, insurance, and market information. Second, the benefit is conditional on digital literacy, infrastructure, and institutional support. Third, in the Rathi and Sahoo (2026) data, perceived usefulness outweighs ease of use as a driver of adoption, so products should be designed around demonstrable economic value. Fourth, environmental outcomes are not automatic: the carbon-footprint hypothesis was not supported in Vasudevan et al. (2025), and the sustainable-practice measure in Rathi and Sahoo (2026) captures adoption of practices rather than measured environmental effect.

Reliability requires qualification. Table 3 recomputes the test statistics and confidence intervals for the seven direct paths reported by Rathi and Sahoo (2026). Six paths are internally consistent. The trust path (H3) is not: the reported t-value of 1.783 implies a two-tailed p of about 0.075, not the reported p

< 0.001, and the interval implied by $\beta = 0.099$ and $SE = 0.054$ is $[-0.007, 0.205]$, which includes zero, whereas the reported interval is $[0.017, 0.181]$. The evidence for a trust effect is therefore weaker than the paper's text suggests.

Other concerns compound this. The target population is defined as farmers familiar with mobile phones, although 7.34 per cent of respondents report owning none, which points to selection on digital awareness. The financial-inclusion construct includes digital-payment usage, which overlaps conceptually with FinTech adoption and may inflate the adoption-to-inclusion path. Some figures in the paper's structural-model diagram also differ from the tables (for example, the reported R^2 for behavioural intention is 0.688 in the table but 0.648 in the diagram).

In Vasudevan et al. (2025), the reporting of unit-root, cointegration, and panel-stationarity tests and macro-level controls such as GDP per capita and energy consumption is difficult to reconcile with a one-time farm-level survey, and one hypothesis (carbon footprint, $\beta = 0.009$, $p < 0.05$) is listed as rejected while the text reports no observable effect.

Table 3. Consistency check of direct-path estimates reported by Rathi and Sahoo (2026)

Path	β	SE	t (reported)	Implied two-tailed p	95% CI (reported)	95% CI ($\beta \pm 1.96$ SE)
H1 PU $\rightarrow$ BI	0.295	0.059	5.039	< 0.001	[0.179, 0.411]	[0.179, 0.411]
H2 PEOU $\rightarrow$ BI	0.146	0.042	3.450	0.001	[0.064, 0.228]	[0.064, 0.228]
H3 PT $\rightarrow$ BI	0.099	0.054	1.783	0.075	[0.017, 0.181]	$[-0.007, 0.205]$
H4 PGS $\rightarrow$ ATU	0.195	0.059	3.308	0.001	[0.080, 0.310]	[0.079, 0.311]
H5 Adoption $\rightarrow$ FI	0.612	0.048	12.750	< 0.001	[0.518, 0.706]	[0.518, 0.706]
H6 FI $\rightarrow$ SAP	0.548	0.062	8.839	< 0.001	[0.426, 0.670]	[0.426, 0.670]
H7 Adoption $\rightarrow$ SAP	0.267	0.055	4.855	< 0.001	[0.159, 0.375]	[0.159, 0.375]

Classifier evaluation under class imbalance

Akça and Sevli (2022) report per-kernel confusion matrices averaged over five folds. Because each fold has 1,000 test instances and 96 positives on average (9.6 per cent), a rule that predicts “reject offer” for every customer would score 90.4 per cent accuracy. Table 4 recomputes the metrics from the reported matrices and adds specificity, balanced accuracy, and MCC.

Table 4. Metrics recomputed from the reported confusion matrices (mean of five folds)

Kernel	Accuracy	Precision	Recall	F1	Specificity	Balanced accuracy	MCC
Linear	0.951	0.864	0.581	0.695	0.990	0.786	0.685
Polynomial	0.972	0.873	0.829	0.850	0.987	0.908	0.835
RBF	0.968	0.892	0.756	0.818	0.990	0.873	0.804
Sigmoid	0.833	0.011	0.008	0.010	0.921	0.465	−0.081
Majority-class baseline	0.904		0.000		1.000	0.500	0.000




Figure 2. Accuracy, balanced accuracy and recall by kernel, recomputed from the reported confusion matrices. The dashed line is the accuracy of a rule that rejects every offer; the balanced-accuracy value for such a rule is 0.5.

Three observations follow. First, the polynomial, RBF, and linear kernels all report accuracy above the 90.4 per cent baseline, but their advantage is smaller than the headline suggests, and the linear kernel misses 42 per cent of acceptors (recall 0.581). Second, the sigmoid kernel’s accuracy of 83.3 per cent is below the do-nothing baseline; its balanced accuracy of 0.465 is below chance and its MCC is negative, meaning it is systematically worse than random assignment on the class that matters. Third, the sigmoid precision recomputed from the reported matrix is $0.8 / (0.8 + 71.6) = 0.011$, whereas the paper’s table and abstract give 0.108; the reported F1 of 0.009 is consistent with 0.011, not 0.108, so the latter appears to be a decimal slip. The ranking of kernels is unchanged, with the polynomial kernel best on every balanced measure, but the exercise illustrates the risk of relying on accuracy.

For agricultural finance the lesson is direct. Default, fraud, and take-up are all rare-event problems, and the costs of errors are asymmetric: wrongly denying a creditworthy smallholder forecloses a livelihood investment, while wrongly approving a defaulter costs the lender. Decisions should therefore be evaluated with balanced accuracy or MCC, precision–recall curves (Saito & Rehmsmeier, 2015), calibration, cost-weighted loss, and group-wise error rates, consistent with Proposition P3.

Mapping fraud and integrity risks to contemporary controls

Table 5 maps the integrity risks identified by Ajah and Inyama (2011) to controls now technically feasible within India’s digital public infrastructure, and notes the exclusion or error risk each control carries. The mapping is analytical: it identifies risk types to be tested rather than asserting documented incidence.

Table 5. Integrity risks, 2011 proposals and contemporary analogues

Risk (Ajah & Inyama, 2011)	2011 proposal	Contemporary analogue in India	Residual or exclusion risk
Multiple or false identities	Biometrics and a national identity number as a mandatory gate for any loan	Farmer ID linked to land records, verified through consented identity checks	Duplicate or stale registry entries; hard single-identifier gates exclude tenants and sharecroppers
No shared record of indebtedness; multiple lending	Central database and credit-bureau checks	Bureau data plus consented account and benefit data; registry-linked loan records	Thin or absent files for first-time borrowers; incomplete co-operative bank reporting
Fake or repeatedly pledged collateral	GPS/GIS capture of property coordinates	Geo-referenced village maps and plot registry; satellite cross-check of declared crops	Boundary disputes; land-record errors; cloud cover and misclassification in imagery
Subjective officer discretion	Neural-network scoring	Explainable scoring with uncertainty-based referral and audit trail	Automation bias; unfairness embedded in historical data
Hidden fraud patterns in bank-wide data	Data mining and monitoring agents	Anomaly detection over loan and transaction graphs	False positives against honest borrowers; need for
Risk (Ajah & Inyama, 2011)	2011 proposal	Contemporary analogue in India	Residual or exclusion risk
			appeal and grievance redress

Synthesis

Five design lessons emerge:

1. Data infrastructure is now more binding than algorithm choice, and the value of AI depends on registry quality, consent, and interoperability.
2. Interpretable models and good evaluation are worth more than marginal accuracy gains.
3. Credit alone is not wealth creation; the design must sequence liquidity, protection, and investment, and must measure assets and net worth.

4. Enabling conditions—digital literacy, connectivity, and institutional support—determine whether benefits materialise, and the co-operative banks that carry much of farm credit need targeted capability building.
5. Integrity controls must be paired with inclusion safeguards, because control mechanisms that are too rigid can exclude the intended beneficiaries.

Implementation Roadmap and Implications

Phased roadmap

Table 6 proposes a three-phase roadmap that emphasises learning before scale, given that much of the evidence is proof-of-concept.

Table 6. Phased implementation roadmap

Phase	Horizon	Key actions	Exit criteria
1. Foundation and pilots	0–12 months	Consent-based data integration in a small set of districts; baseline of household balance sheets; model-risk policy aligned with FREE-AI; onboarding of co-operative banks and FPOs; ground-truth validator networks	Data-sharing agreements and consent flows working; baseline complete; first fairness audit published
2. Scale-out and product integration	12–36 months	Three-route decision policy in production; index-based insurance and savings linked to credit; FPO portfolio tools; multilingual advisory with human escalation; routine group-wise error monitoring	Improved coverage of thin-file, tenant, and women farmers without deterioration in portfolio quality
3. Asset-management maturity	36–60 months	Household balance-sheet views with farmer consent; longitudinal evaluation of asset formation; shared-asset financing through collectives; incentive schemes linked to verified sustainable practices	Demonstrated multi-season rise in productive-asset index and insurance cover; stable repayment

Implications for stakeholders

- **Lenders and co-operative banks.** Adopt interpretable, calibrated models; report balanced and group-wise metrics; use satellite and weather data for portfolio surveillance; pair risk flags with advisory and insurance; invest in staff capability, particularly in co-operative banks.
- **Regulators and policymakers.** Extend FREE-AI implementation guidance to smaller lenders; require fairness audits that include tenancy, gender, and region; keep tenants and sharecroppers eligible and visible in registries; support shared data and audit tooling so that small institutions are not left behind.
- **FPOs, SHGs, and cooperatives.** Develop capacity as data intermediaries and validators; negotiate portfolio-level credit and shared-asset finance; help members understand and contest their scores.
- **Technology providers.** Design for low literacy and intermittent connectivity, in local languages and with voice interfaces; expose explanations and confidence; support consent management and audit logs.
- **Researchers.** Move from cross-sectional attitude surveys to longitudinal, asset-based, and lender-side evaluations with pre-specified metrics.

Monitoring scorecard

Table 7. Scorecard for AI-enabled agricultural finance

Dimension	Indicators	Why it matters
Coverage	Share of registered farmers with a scored or pre-approved limit; share of tenants, women, and marginal farmers among new borrowers	Tests whether inclusion is real, not only aggregate volume
Decision quality	Balanced accuracy or MCC; calibration error; precision–recall; error rates by gender, tenancy, and region	Prevents accuracy illusions and hidden bias
Financial outcomes	Repayment and default rates; ticket-size trends; share of investment credit	Guards against over-indebtedness and short-termism
Wealth and resilience	Change in a productive -asset index; insurance coverage; liquid savings buffer over several seasons	Measures asset formation rather than credit use
Sustainability	Verified adoption of water-efficient irrigation, soil conservation, and diversification	Recognises that digital finance is not automatically green
Trust	Consent revocations; explanation requests; grievance resolution time; human-override rate	Signals legitimacy and safe use

Risks, Ethical Considerations and Limitations

Risks

Several risks warrant attention. Over-indebtedness is possible if a system tuned for approval volume pushes credit onto households that cannot service it; affordability checks, staged disbursement, and asset-based outcome metrics are needed. Feedback loops may arise if models trained on well-served regions neglect underserved ones, creating algorithmic blind spots (Goswami et al., 2026).

The digital divide remains significant. In the Rathi and Sahoo (2026) sample, 37.2 per cent of respondents used a basic phone or had no phone, and 63.8 per cent lived more than five kilometres from a bank branch; digital-only channels would leave such households behind, which is why assisted and collective channels are part of the framework. Privacy and consent are critical: continuous data flows raise the stakes for consent management, purpose limitation, and data-protection compliance.

Generated advice also poses risks. Multilingual assistants can err; advice that affects money or crop decisions needs guardrails and escalation to humans. Finally, concentration and vendor risk matter: dependence on a few data or model providers can create systemic fragility.

Limitations of this paper

This is a secondary, integrative study with no primary data. The evidence base is small, and two of its five studies are cross-sectional surveys from a single southern state, so generalisation to other agro-ecological and institutional settings is uncertain. Official statistics are moving quickly and were taken from government statements as reported in September 2026; the 2015–16 census shares are the latest all-India figures used here and should be updated when newer census results are published.

The framework and propositions are conceptual and untested. The consistency checks establish internal inconsistencies but cannot determine their cause without access to the original data.

From the author's perspective, a key limitation is that the proposed framework has not yet been tested in collaboration with any lending institution or farmer producer organisation. Future work will aim to address this through pilot studies.

Directions for Future Research

- Cluster-randomised or matched evaluations of the three-route decision policy against binary denial, with repayment, inclusion, and asset outcomes.
- Longitudinal panels that measure productive-asset indices, insurance and savings buffers, and net worth of farm households over several seasons.
- Fairness audits of agricultural credit models by gender, tenancy, farm size, and region, using independence, separation, and sufficiency criteria.
- Comparative cost-effectiveness of satellite-based, sensor-based, and field-based verification, including performance under cloud cover and small plot sizes.
- Studies of AI readiness in co-operative banks and regional rural banks, including data quality, staff skills, and governance.
- Evaluation of accuracy, safety, and accountability of multilingual generative-AI advisory in credit and insurance decisions.
- Analysis of whether digital finance can be linked to verified environmental outcomes, given the null carbon-footprint result reported in the literature.

Conclusion

India has assembled much of what AI-enabled agricultural finance needs: farmer and crop registries at scale, a renewed farm-credit framework, and a regulator's blueprint for responsible AI. The existing evidence, however, is thin on the outcome that matters most, namely sustained wealth creation among smallholders, and shows that digital finance helps mainly where literacy, connectivity, and institutional support are in place.

Reappraisal of the available studies also shows that headline statistics can mislead: an accuracy figure that falls below a do-nothing baseline, and a trust effect whose statistics do not support its reported significance, are cautionary examples. The framework proposed here treats AI as an instrument for managing the household, collective, and institutional balance sheet, rather than as a faster loan-approval engine.

Its distinguishing features are consent-based data flows, interpretable models evaluated with appropriate metrics, a decision policy that offers support instead of bare refusal, and collectives that act as validators and aggregators. If lenders, regulators, and farmer organisations adopt these principles and test the seven propositions through carefully designed pilots, AI-enabled finance can help shift rural households from cycles of seasonal borrowing towards durable asset ownership and sustainable prosperity.

References

- Ajah, I., & Inyama, C. (2011). Loan fraud detection and IT-based combat strategies. *Journal of Internet Banking and Commerce*, 16(2).
- Akça, M. F., & Sevlı, O. (2022). Predicting acceptance of the bank loan offers by using support vector machines. *International Advanced Researches and Engineering Journal*, 6(2), 142–147. <https://doi.org/10.35860/iarej.1058724>

Chambers, R., & Conway, G. R. (1992). Sustainable rural livelihoods: Practical concepts for the 21st century

(IDS Discussion Paper 296). Institute of Development Studies.

Chen, S., Lin, H., & Yang, G. (2022). Efficient agricultural disaster financing using satellite data and artificial intelligence. *Computers and Electrical Engineering*, 103, 108394. <https://doi.org/10.1016/j.compeleceng.2022.108394> [Retracted; excluded from the evidence synthesis.]

Chicco, D., & Jurman, G. (2020). The advantages of the Matthews correlation coefficient (MCC) over F1 score and accuracy in binary classification evaluation. *BMC Genomics*, 21, 6. <https://doi.org/10.1186/s12864-019-6413-7>

Davis, F. D. (1989). Perceived usefulness, perceived ease of use, and user acceptance of information technology. *MIS Quarterly*, 13(3), 319–340.

Goswami, Y., Bangde, S., Singh, S., & Khosla, A. (2026). Revolutionizing agricultural loan recommendation systems via machine learning and artificial intelligence: A systematic literature review. *Computers and Electronics in Agriculture*, 241, 111231. <https://doi.org/10.1016/j.compag.2025.111231>

Hardt, M., Price, E., & Srebro, N. (2016). Equality of opportunity in supervised learning. *Advances in Neural Information Processing Systems*, 29.

Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. *Advances in Neural Information Processing Systems*, 30.

Ministry of Agriculture and Farmers Welfare. (2015–16). Agriculture Census 2015–16 [All-India results; average holding size of 1.08 hectares also stated in replies to Parliament]. Government of India.

Ministry of Agriculture and Farmers Welfare. (2026, August 7). Written reply in the Rajya Sabha on AgriStack Farmer IDs and Digital Crop Survey [as reported by BioVoice News, “India creates over 10.31 crore Farmer IDs under AgriStack”]. Government of India.

Ministry of Finance. (2025, February 26). Kisan Credit Card amount crosses Rs 10 lakh crore, benefits 7.72 crore farmers [Press statement as reported by IANS]. Government of India.

Mondaq. (2025). RBI releases report on responsible and ethical enablement of AI in the financial sector [Legal commentary on the FREE-AI report].

Rathi, R., & Sahoo, A. K. (2026). Impact of FinTech adoption on financial inclusion and sustainable agricultural practices: Extending the Technology Acceptance Model among smallholder farmers in South Tamil Nadu, India. *International Journal of Economic Practices and Theories*, 2026(1), 1444–1471.

Reserve Bank of India. (2025, August 13). Report of the Committee to develop a Framework for Responsible and Ethical Enablement of Artificial Intelligence (FREE-AI) in the Financial Sector.

Reserve Bank of India. (2026a, June 19). Reserve Bank of India (Commercial Banks – Kisan Credit Card (KCC) Scheme) Directions, 2026 (RBI/FIDD/2026-27/402).

Reserve Bank of India. (2026b, June 19). Reserve Bank of India (Regional Rural Banks – Kisan Credit Card (KCC) Scheme) Directions, 2026 (RBI/FIDD/2026-27/404).

Rogers, E. M. (2003). *Diffusion of innovations* (5th ed.). Free Press.

Saito, T., & Rehmsmeier, M. (2015). The precision-recall plot is more informative than the ROC plot when evaluating binary classifiers on imbalanced datasets. PLoS ONE, 10(3), e0118432. <https://doi.org/10.1371/journal.pone.0118432>

Sherraden, M. (1991). Assets and the poor: A new American welfare policy. M. E. Sharpe.

United News of India. (n.d.). KCC outstanding credit crosses Rs 10.2 lakh crore nationwide. UNI.

Vasudevan, A., Rani, P. J., Raja, N., Nedumaran, G., Arokiasamy, A. R. A., & Qian, C. (2025). Fintech for sustainable agriculture: Insights from Tamil Nadu, India. Frontiers in Sustainable Food Systems, 9, 1614553. <https://doi.org/10.3389/fsufs.2025.1614553>