International Journal of Multidisciplinary Research and Growth Evaluation.

Beyond More Data: Do Agricultural Signals Improve Financial Inclusion Prediction?

Arthur Brick Kamya [1*], Dr. Nasake Julius Alochere [2]
[1] PhD Candidate in Law, Student and Researcher, Kampala International University (KIU), Kampala, Uganda
[2] College of Education, Open and Distance Learning, Kampala International University, Kampala, Uganda

* Corresponding Author: **Arthur Brick Kamya**

Article Info

ISSN (Online): 2582-7138
Volume: 07
Issue: 06
Received: 24-08-2026
Accepted: 26-09-2026
Published: 28-10-2026
Page No: 144-153

Abstract
Agricultural livelihoods are often treated as a promising source of information for predicting financial inclusion. Yet their incremental predictive value is rarely separated from broader development signals or tested across entirely unseen countries. This study links Global Findex account-ownership estimates with rural water access, urbanization, population, agricultural trade indicators, and agricultural production values. The resulting unbalanced panel contains 294 country-year observations from 70 low- and lower-middle-income economies over 2011–2024. Ridge regression, random forests, and XGBoost are compared using five-fold country-grouped validation, an agricultural feature-block ablation, and a separate 2024 temporal holdout. A Uganda National Panel Survey household analysis provides a distinct descriptive perspective on saving channels. Core random forests and XGBoost achieve country-held-out mean absolute errors of approximately 14.69 percentage points, compared with 17.55 for a training-mean benchmark. Adding six agricultural variables changes random-forest error by +0.05 points, with a country-cluster bootstrap interval of −0.44 to +0.56; none of the three models shows a clear pooled improvement. In the 2024 holdout, a last-observed account-rate benchmark achieves 7.81 points of error, substantially below the best contemporaneous structural model at 13.36. Among 2,987 Ugandan households with complete saving responses, institution-or-mobile saving is less prevalent among farming-main-income households than other households, while savings-group use is more prevalent. These results support a cautious conclusion: agricultural context remains relevant to service design, but more agricultural data do not automatically improve transferable national predictions.

Keywords: Financial inclusion, agricultural livelihoods, Global Findex, country-grouped validation, random forest, XGBoost, savings groups, Uganda

1. Introduction

Financial inclusion is easiest to count when a person owns an account, but harder to understand when that account must support an irregular livelihood. A farmer may receive cash only after harvest, save through a village association, and use a mobile-money agent without maintaining a bank balance. An account-ownership statistic captures part of this experience, not the whole of it. The Global Findex Database 2025 brings financial access and digital connectivity into a common empirical frame, while the digital-inclusion literature emphasizes that access, use, and enabling infrastructure are related but distinct concepts (Klapper et al., 2025; Khera et al., 2022) [6, 7]. The present study begins with that distinction rather than assuming a single indicator represents every dimension of inclusion.

Agriculture is an appealing source of additional predictive information because it links income generation, market participation, and exposure to seasonal risk. Household evidence from Zimbabwe identifies education, off-farm income, financial literacy, and distance among the correlates of smallholder financial inclusion (Mhlanga & Dunga, 2020). Evidence from Pakistan also connects agricultural

participation with financial inclusion and digital finance, although its outcome is agricultural participation rather than account ownership (Mumtaz, 2024). These findings motivate an agricultural feature block, but they do not establish that national trade or production aggregates will improve predictions of national account ownership. That missing step is an empirical question, not a conclusion that can be imported from household studies.

The methodological concern is equally important. Repeated observations from the same country can appear on both sides of an ordinary random split, allowing a model to benefit from familiar national structure. Such performance may be useful for a narrow interpolation task, but it is not evidence of successful transfer to unseen economies. Leakage and mismatched evaluation settings can inflate the apparent value of machine learning in scientific applications (Kapoor & Narayanan, 2023) [5]. Here, all observations from each country remain together during the principal validation exercise. A separate temporal evaluation asks whether relationships estimated before 2024 transport to a later wave, and a persistence benchmark tests whether complex structural models improve on a simple previous observation.

Three questions guide the article. First, how accurately can a parsimonious set of development signals predict national account ownership in held-out low- and lower-middle-income economies? Second, do agricultural indicators add measurable predictive value on the same observations and evaluation folds? Third, what does a separate household view reveal about saving arrangements that national account ownership cannot describe? The contribution is the combination of these questions in a transparent workflow, not an unverified claim of being the first study of its kind. The analysis deliberately broadens the benchmark topic to include lower-middle-income economies and narrows the predictive target to account ownership, making both the geographic scope and the outcome explicit.

2. Literature Review

2.1. Financial Inclusion Is Not a Single Financial Behavior

A central lesson of the literature is that financial inclusion cannot be reduced to possession of one product. Khera et al. (2022) [6] construct a digital financial inclusion index that recognizes multiple aspects of digital financial development across emerging and developing economies. This multidimensional perspective matters because a person may own an account without using it regularly, or may save through a group without owning a conventional bank account. The present article therefore treats national account ownership as a defined, limited target. Its Uganda component addresses saving channels separately rather than relabeling every non-bank arrangement as exclusion. This conceptual separation also prevents a household saving indicator from being presented as an independent replication of a national account-ownership model.

2.2. Socioeconomic and Rural Determinants

Studies of sub-Saharan Africa establish a useful development context but differ in their level of analysis. Bashiru et al. (2023) [1], using a dynamic panel approach, report positive relationships for literacy and financial globalization and a negative relationship for rural population growth. Jama et al. (2025) [4] investigate drivers of financial inclusion across sub-Saharan African countries, contributing to the regional determinants literature. Neither type of study removes the need to assess out-of-sample predictive performance for a newly assembled dataset. Explanatory coefficients, associations, and transferable predictions answer different questions. This article consequently uses regional and income-group context as baseline information while refusing to treat their predictive importance as a demonstrated mechanism for expanding financial access.

2.3. Agricultural Livelihoods and Financial Participation

Farmer-focused studies provide a more direct livelihood connection. In Zimbabwe, Mhlanga and Dunga (2020) show that financial inclusion among smallholders is associated with characteristics extending beyond crop production itself, including off-farm income, education, and proximity to financial institutions. Mumtaz (2024) examines financial inclusion, digital finance, and agricultural participation in Pakistan using household data and a nonlinear decomposition. Together, these studies suggest that agricultural participation is embedded in household resources, services, and technology access. They also caution against equating a country's agricultural export share with the banking position of its farmers. Aggregate agricultural indicators may describe the economic setting without carrying enough additional information to predict financial access once broad development context is already observed.

Research on resilience adds another reason to take rural finance seriously without confusing outcomes. Negera et al. (2025) analyze Ethiopian rural households using a multidimensional financial inclusion index and an instrumental-variable strategy, reporting a positive relationship with climate resilience. Their study addresses a substantive development outcome rather than the predictive usefulness of agricultural covariates. The implication for the present article is therefore conceptual: financial inclusion can matter to livelihoods even if a national agriculture block does not improve prediction. An instrumental-variable design depends on assumptions about its instruments, whereas the present exercise makes no identification claim. The two forms of evidence answer different questions: one investigates a proposed relationship and the other evaluates whether covariates carry additional information outside the fitted sample.

2. Literature Review and Analytical Framework

2.4. Digital Services, Prediction, and Appropriate Benchmarks

Experimental evidence is particularly helpful for separating plausible policy mechanisms from observational prediction. Batista and Vicente (2020) [2] study improvements in access to savings through mobile money among African smallholder farmers. Their experimental design concerns a specific intervention and population, not a universal effect of national agricultural trade on account ownership. The study illustrates why service relevance and predictive relevance need not coincide. A saving product may benefit farmers because it fits their circumstances, even when a country-level agriculture variable adds little to an algorithm's accuracy. The present article accordingly uses household saving profiles to identify questions about service design, while reserving causal language for studies whose designs support it rather than for its own descriptive comparisons.

Machine learning provides flexible tools for examining financial inclusion, but flexibility is not evidence of

superiority in every setting. Maehara et al. (2024) [8] compare several algorithms for financial inclusion in Peru, including random forests, XGBoost, and generalized linear approaches. Saidy and Hassan (2025) apply machine learning to fintech adoption in sub-Saharan Africa. These studies support the relevance of predictive methods to financial-service research, but their populations, targets, and available features differ from those used here. Direct comparison of accuracy scores would therefore be misleading. This article compares models only within a shared dataset and validation design, reporting errors in percentage points and retaining a transparent mean benchmark alongside the more flexible estimators.

2.5. A Testable Distinction: Context Versus Incremental Information

The analytical framework separates three layers. The first is national development context: survey year, income-group label, region, population, urbanization, and rural water services. The second is agricultural structure: export intensity, trade openness, commodity-level trade summaries, and agricultural production value relative to the adult population. The third is household financial practice, examined separately in Uganda through institution-or-mobile saving and group saving. No causal arrow between these layers is estimated. Instead, the national exercise asks whether the second layer improves prediction beyond the first, while the household exercise asks whether saving channels differ across livelihood and residence groups. This structure permits an informative negative predictive result without erasing the practical relevance of rural livelihoods.

The core proposition is intentionally modest: agricultural signals are useful to a transferable national model only if they lower held-out error on comparable observations. A second proposition is that a model fitted to an earlier wave should be judged against information already available for previously observed countries, including the last observed account rate. A third is that policy interpretation must respect the measurement unit. A national model cannot identify an individual farmer who lacks an account, and a household saving profile cannot validate national account predictions. Together, these propositions turn the original broad topic into an auditable study of incremental value, temporal transportability, and measurement boundaries rather than a search for the most impressive algorithmic score.

3. Methodology

3.1. Data Provenance, Scope, and Integration

The archive contains 383 files, including data tables, metadata, software assets, and multiple snapshots. It is not one rectangular dataset. National account-ownership estimates come from the Global Findex file; only rows labeled “all” in both subgroup fields are retained. ISO-recognized countries with low-income or lower-middle-income labels in the incomegroupwb24 field define the sample. Those labels are held fixed across waves and should not be interpreted as historical income classifications. After requiring an observed target, the analysis includes 25 low-income and 45 lower-middle-income economies. The country-year is the modeling unit, and every national join is checked to be one-to-one, preventing subgroup or commodity duplication from multiplying outcome rows.

Table 1: Dataset-family audit and analytical decisions

Family	Role or decision
Global Findex 2025	National account ownership and adult population; all/all rows only (Klapper et al., 2025) [7].
FAOSTAT trade indicators	Retained: two aggregate trade measures and three medians over available commodity indicators.
FAOSTAT production values	Retained: Agriculture aggregate, constant 2014–2016 US$, element code 58. Streamed rather than fully extracted.
JMP processed repository	Retained: rural basic water and urban population share. Sanitation excluded: 96.3% missing in the selected snapshot.
World Development Indicators	Retained: rural safely managed drinking water; uploaded header updated July 13, 2026 (World Bank, 2026).
FAOSTAT value shares / agriculture census	Excluded from models: selected food-at-home agriculture-share series has no sample overlap; census observations are not an annual panel.
FAOSTAT raw trade / detailed matrix / trade indices	Inventoried, not fully modeled: alternate commodity/bilateral representations require separate harmonization and risk redundant or double-counted aggregates.
Uganda household survey modules	Separate weighted saving-channel analysis; household records are not appended to country-year outcomes (Uganda Bureau of Statistics, 2021).

FAOSTAT supplies the agricultural sources; the FAO yearbook provides database context (Food and Agriculture Organization of the United Nations [FAO], 2025). JMP release context and the processing package are documented separately (World Health Organization [WHO] & United Nations Children’s Fund [UNICEF], 2025; Schöbitz, 2026).

“Combine all datasets” is implemented as an audit of all substantive families, not an indiscriminate merge. The national tables are matched by country code and observation year without forward filling. UN M49 codes in the agricultural tables are mapped to ISO3 codes; non-country regional aggregates are excluded. The Uganda files are linked by household identifier and remain a separate descriptive sample. The processed JMP snapshot is dated August 20, 2026 and represents the 2025 release, but its sanitation fields are largely missing. That pattern is treated as a limitation of the processed file, not as evidence that the countries lack sanitation data. No replacement national data are downloaded to manufacture coverage.

3. Methodology: Outcome and Predictors

3.2. Target and Feature Blocks

The dependent variable is the account estimate multiplied by 100, yielding national adult account ownership on a percentage scale. It includes the account concept used in the harmonized Findex file and is not a measure of credit access, financial literacy, or active saving. Because the target is continuous, random forests and XGBoost are fitted as regressors rather than classifiers. The core block contains seven raw predictors and the agricultural extension contains six. Findex financial-use variables are not used as predictors, avoiding features that are closely entangled with the target. Mobile and internet measures that are not consistently

observed across the waves are also omitted; the resulting baseline is deliberately parsimonious rather than a complete model of every inclusion determinant.

Table 2: Predictor definitions and observed coverage (N = 294)

Block	Predictor	Observed n	Missing %
Core	Survey year	294	0.0
Core	Fixed income-group label	294	0.0
Core	Region label	294	0.0
Core	log(1 + adult population)	294	0.0
Core	Urban population share (%)	293	0.3
Core	Rural at-least-basic water (%)	283	3.7
Core	Rural safely managed water (%)	196	33.3
Agriculture	Agricultural export share of GDP	214	27.2
Agriculture	Agricultural trade openness index	214	27.2
Agriculture	Median available commodity terms of trade	188	36.1
Agriculture	Median available commodity import dependency	209	28.9
Agriculture	Median available commodity revealed advantage	212	27.9
Agriculture	log(1 + constant-US$ agriculture value/adult)	242	17.7

Raw database units are retained for export share, openness, and commodity ratios. The production numerator is thousand constant 2014–2016 US$, multiplied by 1,000 before division by adult population. This denominator is a scale adjustment, not output per agricultural worker or farm productivity.

Aggregate trade measures are taken from the "Crops and livestock products" item. Terms of trade, import dependency, and revealed comparative advantage are commodity-specific; their country-year medians summarize only the available commodity basket and are not aggregate national equivalents. The medians are unweighted, and changing basket coverage can alter their interpretation. Self-sufficiency is excluded because it is mechanically related to import dependency. The import and export concentration indicators refer to the World aggregate, not individual countries, and are therefore excluded. These decisions are made before model comparison and documented in the reproducible code so that broad database labels cannot quietly become misleading country-level features.

3. Methodology: Validation and Estimation

3.3. Models, Preprocessing, and Evaluation

Four estimators are compared: a training-fold mean, ridge regression, a random forest, and XGBoost. Settings are fixed in advance rather than optimized against test results. Numeric missing values are replaced by training-fold medians, with missingness indicators added where applicable; categorical predictors are imputed using the training-fold mode and one-hot encoded. Numeric scaling is fitted inside the same pipeline. Each preprocessing step is re-estimated using only training observations. Predictions are clipped to the admissible 0–100 range before evaluation. Appendix A records the principal settings. This design controls evaluation leakage, but it is not an exhaustive hyperparameter search, so the findings apply to these transparent model specifications rather than all possible implementations of the algorithms.

The primary test is five-fold GroupKFold validation with country as the grouping variable. Every observation for a country is assigned to the same held-out fold, and country identifiers are not predictors. Core and expanded models use exactly the same observations and folds. Errors are pooled across the 294 out-of-fold predictions, giving country-years equal weight rather than weighting countries by adult population. Mean absolute error (MAE) and root mean squared error (RMSE) are expressed in percentage points; out-of-sample R^2 is reported as a complementary fit measure. This task estimates transfer to countries absent from model training. It differs from an ordinary random row split and should not be interpreted as an evaluation of individual-level financial-exclusion detection.

Incremental agricultural value is measured as expanded-model MAE minus core-model MAE. A negative difference indicates improvement. A paired bootstrap resamples countries with replacement 2,000 times and carries all their observed years together; percentile intervals describe uncertainty in the error difference, conditional on the fitted out-of-fold predictions. They do not include retraining or model-selection uncertainty. Joint permutation of predefined predictor blocks is performed on held-out folds of the full random forest, with 20 repeats per block and fold. The average increase in MAE assesses predictive reliance, not direction or causality. The fold range is displayed as a stability diagnostic, not a confidence interval, and correlated blocks can limit the interpretability of permutation scores.

3.4. Temporal and Household Checks

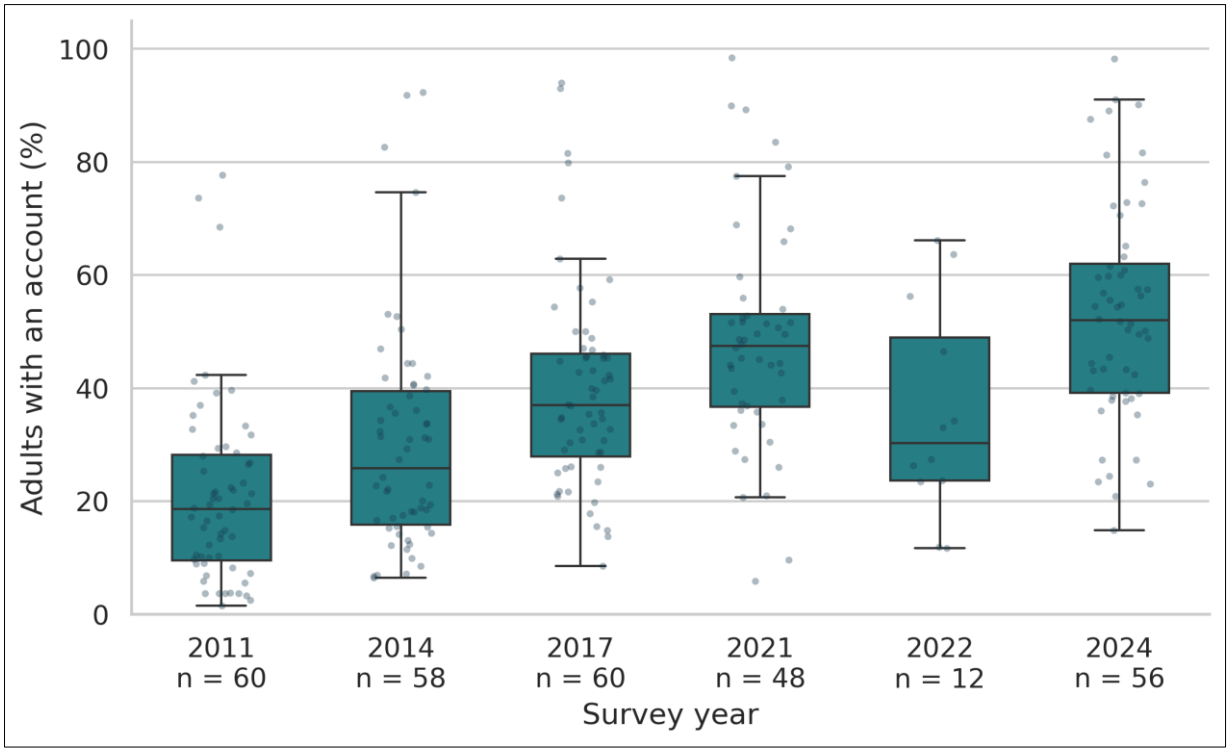
The temporal exercise trains on all observations before 2024 and tests on the 56 observations from 2024. It is an out-of-period prediction using contemporaneous covariates, not a real-time forecast: the datasets contain later revisions and their historical publication lags are not reconstructed. All 56 test countries also have a previous Findex observation, permitting a last-observed-rate benchmark. A separate sensitivity analysis repeats country-grouped validation on 172 observations from 45 countries with all six agricultural values observed. Finally, Uganda household saving profiles use the household weights and complete coded responses; they are descriptive, permit overlapping saving channels, and do not attach significance tests or causal interpretations to group differences.

4. Results and Discussion

4.1. Coverage and the Distribution of Account Ownership

The sample is both geographically broad and uneven over time. Sub-Saharan Africa contributes 152 observations from 37 countries, while the remaining 142 observations come from five other regions. There are 94 low-income and 200 lower-middle-income country-years. National account ownership ranges from 1.52% to 98.46%, with a pooled median of 35.78%. The spread is large enough that a single

average is a weak prediction for many economies. Basic rural water and urbanization measures have relatively strong coverage, but safely managed rural water is missing for one-third of observations. Agricultural variables have missingness ranging from 17.7% to 36.1%, making both imputation and a same-sample comparison important to the analysis.



Note. Each point is a country-year estimate. Boxes show the interquartile range and median; whiskers follow the standard 1.5-interquartile-range convention. Sample sizes differ across years. Source: authors' calculations from the Global Findex file.

Fig 1: Account-ownership distributions in the linked country-year sample.

Account ownership is generally higher in later waves, but the figure is not a balanced-panel estimate of change. The equal-country-year mean is 20.99% in 2011 and 53.23% in 2024, while the 2022 mean is lower than the 2021 mean. That 2022 pattern should not be described as a global reversal: only 12 selected economies appear in the 2022 rows, compared with 48 in 2021, and the composition differs. Survey year consequently has room to act as a strong predictive signal. At the same time, variation within each wave remains substantial, implying that a calendar trend alone cannot describe the diversity of financial access represented in the linked sample.

4. Results: Country-Held-Out Predictions

4.2. Flexible Models Improve on the Mean, but Remain Imprecise

Core random forests and XGBoost both achieve an MAE of approximately 14.69 points. Their unrounded difference is about one-thousandth of a point and should not be presented as a meaningful algorithmic victory. The random forest reduces MAE by 16.3% relative to the training-fold mean benchmark at 17.55, while core ridge regression achieves 14.97. These gains demonstrate that the observed development context contains some transferable information. However, pooled predictive R^2 remains around 0.18 for the two core tree models, and RMSE is close to 20 points. The models therefore leave substantial uncertainty in national account ownership, limiting their suitability for fine-grained ranking or confident policy allocation without additional evidence.

Table 3: Pooled country-grouped out-of-fold performance

Model	Features	MAE	RMSE	R^2
Mean	Core	17.55	22.08	-0.026
Ridge	Core	14.97	20.12	0.148
Random forest	Core	14.69	19.68	0.184
XGBoost	Core	14.69	19.70	0.183
Ridge	Core + agriculture	15.23	20.41	0.123
Random forest	Core + agriculture	14.74	19.68	0.184
XGBoost	Core + agriculture	14.90	19.87	0.168

Note. N = 294 country-years across 70 countries in five held-out country folds. MAE and RMSE are percentage points. The mean benchmark is identical under both feature sets, so it is listed once. Lower errors and higher R^2 indicate better performance. Settings are fixed; no test-fold tuning is conducted.

The expanded random forest achieves an MAE of 14.74 and the expanded XGBoost model 14.90. Expanded ridge regression reaches 15.23, also slightly worse than its core specification. A model with more variables can still look plausible in an importance plot or fit the observed training sample, yet provide no improvement for unseen countries. The matched-fold ablation exposes that distinction. It also cautions against comparing the present scores with household-classification studies: the target, unit, sample, and validation task are different. These results measure errors in national percentages, not the probability of correctly classifying a household as financially included. Accuracy percentages, F1 scores, and household-level claims would be inappropriate summaries of this exercise.

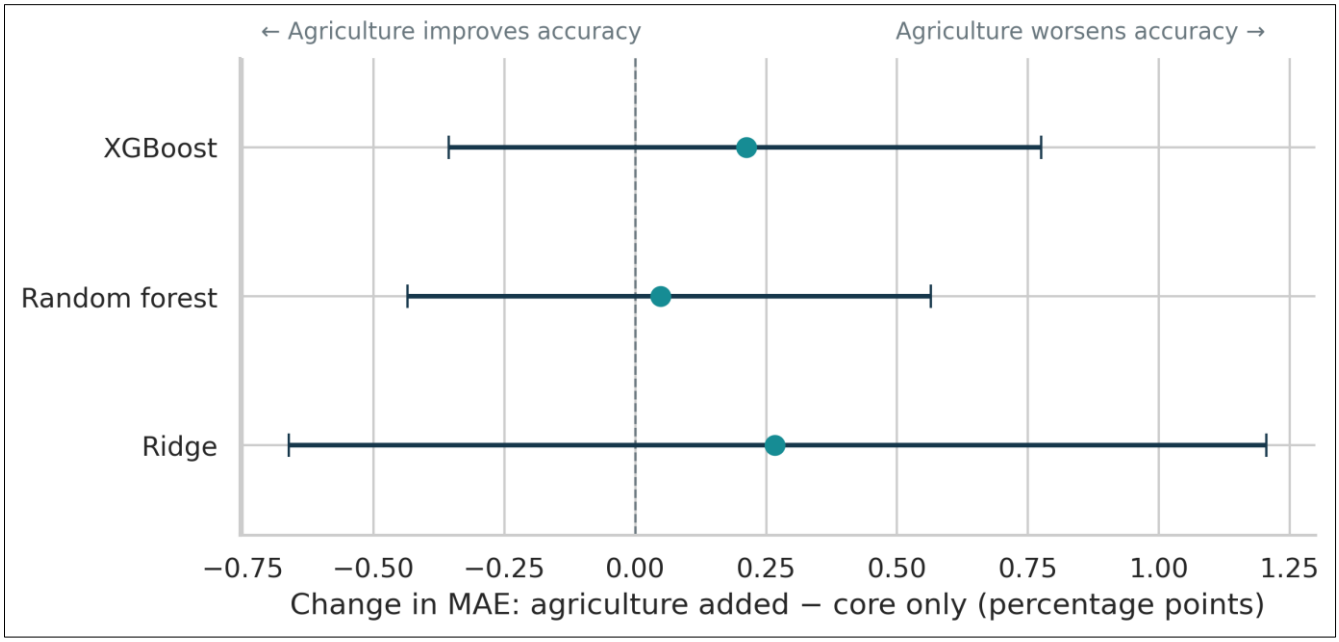
4.3. Performance Differs Across Income-Label Groups

Subgroup error checks reinforce the need for cautious interpretation. For the core random forest, MAE is 12.53 points among low-income observations and 15.71 among lower-middle-income observations. Core ridge regression is lower-error within the low-income subset at 11.23 points, showing that the pooled ranking does not hold uniformly across groups. These comparisons are descriptive evaluations of the same out-of-fold predictions, not separate tuned models or evidence that one population is intrinsically easier to predict. Target variability, geography, sample composition, and the distribution of missing covariates all differ. The broader lower-middle-income scope improves the reach of the study, but it also makes a single pooled performance number an incomplete description of model reliability.

4. Results: Incremental Agricultural Value

4.4. Agricultural Variables Do Not Show a Clear Pooled Gain

The random-forest error difference is +0.048 points, with a bootstrap interval from −0.435 to +0.565. For ridge regression, the difference is +0.266, with an interval from −0.663 to +1.205; for XGBoost it is +0.213, with an interval from −0.356 to +0.775. All three intervals cross zero. These findings neither demonstrate a useful pooled improvement nor prove exact equivalence between feature sets. Instead, they show that the observed gain is too uncertain and too small to justify a general claim that the agricultural block improves cross-country prediction. This inference is narrower than saying agriculture has no relationship with financial inclusion: the test concerns additional predictive information in six particular aggregate indicators.



Note. Points show expanded-minus-core MAE. Bars are 95% percentile intervals from 2,000 country-cluster resamples of paired out-of-fold errors. Positive values indicate worse predictions after adding agriculture. The intervals are conditional on the fitted predictions; they do not include model retraining.

Fig 2: Paired agricultural-block ablation with country-cluster bootstrap intervals.

The complete-agriculture sensitivity analysis helps assess whether missingness alone drives this result. On 172 observations from 45 countries, expanded random-forest MAE is 15.18 compared with 15.18 after rounding for the core model; the unrounded difference is +0.008. Ridge again worsens, from 14.47 to 14.75. XGBoost improves from 15.83 to 15.44, a gain of 0.39 points. This exception matters: agricultural value is not uniformly absent across algorithms and samples. Nevertheless, there is no consistent gain across the tested specifications. The complete-case subset has a different country composition and is not a definitive missing-data correction. It serves as a diagnostic against attributing the entire pooled finding to imputation.

4. Results: Temporal Transportability

4.5. A Simple Previous Observation Beats the Structural Models

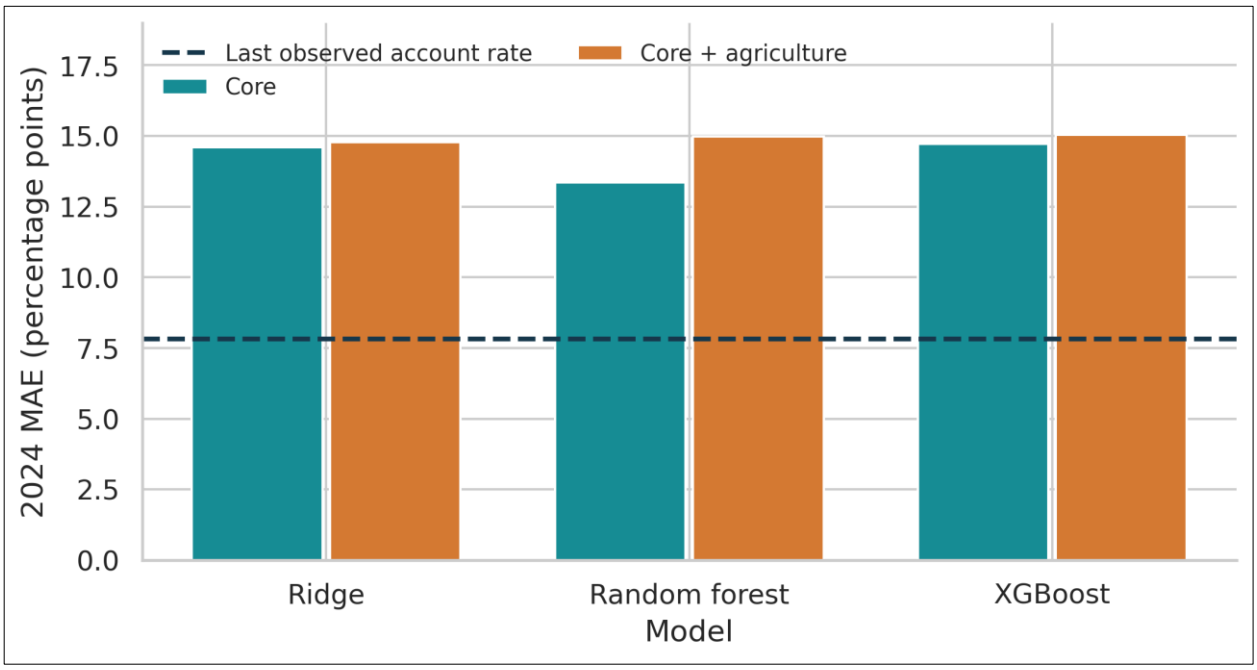
In the later-wave holdout, the core random forest performs best among the structural models, with MAE of 13.36 points and RMSE of 17.26. Adding agricultural variables increases its MAE to 14.97; the expanded XGBoost model reaches 15.03 and expanded ridge 14.78. The full tree models have slightly negative holdout R^2, meaning that their squared-error performance is worse than predicting the test sample's own mean, although that test mean is not an available operational predictor. Because the agricultural extension fails to improve this additional evaluation, the results do not support its routine inclusion without a clearer use case, stronger data coverage, or a revised representation of rural livelihoods.

The last-observed-rate benchmark achieves MAE of 7.81 points, RMSE of 9.85, and R^2 of 0.736. Its advantage is not evidence that simple methods always outperform machine learning. Rather, a previous measurement of the same outcome carries country-specific information that the structural feature sets intentionally lack. The comparison makes the operational distinction visible: cold-start prediction for an unseen country and updating estimates for a previously surveyed country require different baselines. A practical monitoring system should not discard a recent survey measurement merely to use a more complex model. Conversely, persistence cannot solve the cold-start problem. These findings favor task-specific benchmarks and suggest that future updating models should test incremental value beyond the most recent observed account rate.

4. Results: Predictive Reliance and Its Limits

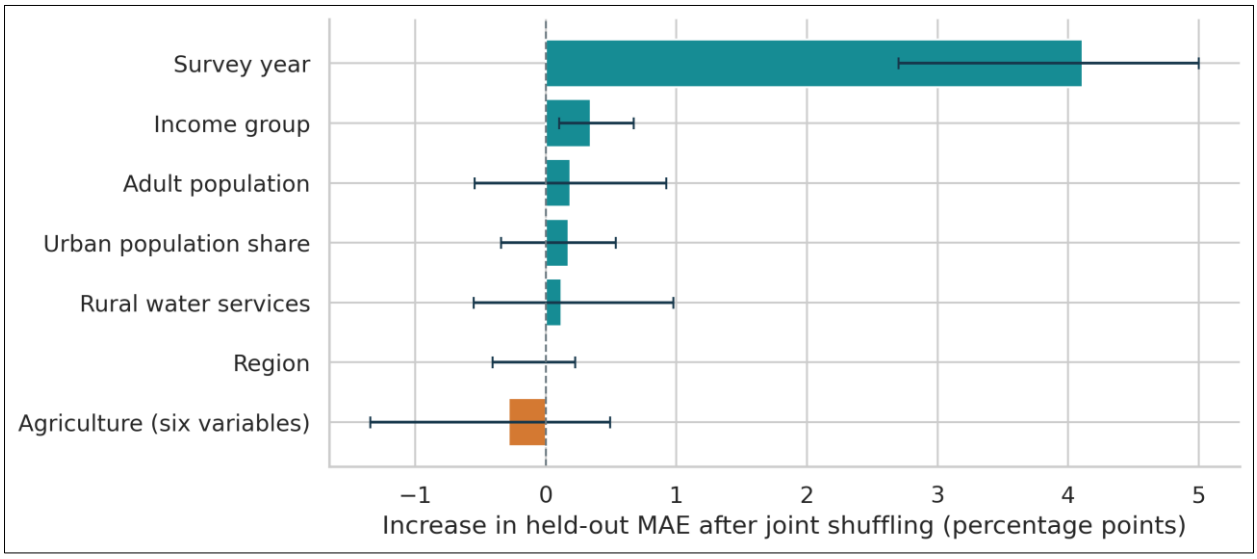
4.6. Calendar Period Dominates the Full Random Forest

Jointly shuffling survey year increases held-out MAE by 4.11 points on average, substantially more than shuffling any other block. The mean increase is 0.35 for income group, 0.19 for adult population, 0.18 for urban population share, and 0.12 for rural water services. The regional block is approximately zero, while the agricultural block has a negative mean importance of −0.28 points. Fold ranges reveal instability in several of the smaller scores. The most defensible interpretation is that the full random forest relies primarily on the broad period pattern, with limited stable additional reliance on the agriculture block in this evaluation. The plot does not identify which predictor raises account ownership or quantify a policy effect.



Note. Models train on the 238 pre-2024 observations and predict the 56 observations from 2024. The dashed benchmark uses each country's latest earlier Findex account rate: 2021 for 45 countries and 2022 for 11. This benchmark addresses known countries and is not available in the unseen-country task.

Fig 3: Errors for the 2024 holdout and the last-observed-rate benchmark.



Note. Bars show mean MAE increases after joint permutation within held-out folds. Whiskers span the minimum and maximum fold-mean importance across five folds, not confidence intervals. Each fold uses 20 repeats. Negative values mean the perturbation reduced held-out error.

Fig 4: Held-out block permutation importance in the expanded random forest.

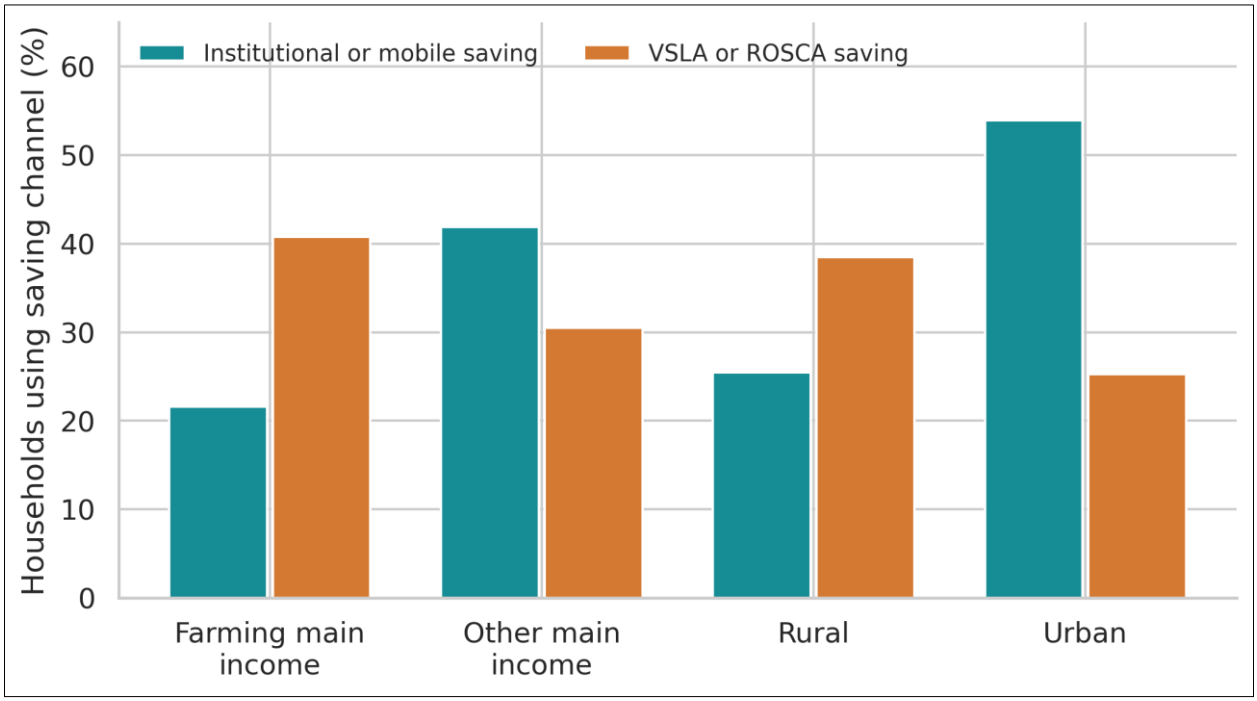
A negative agricultural permutation score can occur when fitted relationships are weak, redundant, or harmful for held-out observations. It is not evidence that agriculture reduces inclusion. Water measures likewise represent contextual service conditions rather than direct financial-service interventions; their modest scores do not justify treating water provision as an estimated inclusion policy. Correlation between development indicators can redistribute apparent importance, and a joint block shuffle creates combinations that need not resemble real countries. Consequently, the ablation and temporal tests carry more weight than a feature-ranking story alone. The combined evidence supports

revisiting the representation of livelihoods, perhaps through aligned household data and service-distance measures, rather than interpreting aggregate feature rankings as instructions for targeted credit decisions.

4. Results: The Uganda Household Lens

4.7. Saving Channels Reveal Livelihood Differences Hidden by One Indicator

The household exercise starts with 3,078 saving-module records and retains 2,987 households with positive household weights and complete 1/2-coded responses to the selected saving items. Institution-or-mobile saving is reported by a weighted 21.61% of farming-main-income households, compared with 41.91% of households with other main income sources. Group saving moves in the opposite direction: 40.78% versus 30.51%. The corresponding sample sizes are 1,344 and 1,643 households. Rural households also report less institution-or-mobile saving than urban households, 25.43% versus 53.95%, while group saving is more prevalent in rural areas. These descriptive contrasts do not identify the effect of farming or residence; they show how the observed saving arrangements differ across household groups.



Note. Institution-or-mobile saving includes commercial banks, deposit-taking microfinance institutions, SACCOs, other microfinance institutions, or mobile money (CB06B/C/D/E/H). Group saving includes VSLAs or ROSCAs (CB06F/G). Categories overlap. Farming main income comprises subsistence or commercial farming (s11q01 = 1 or 2). N = 2,987 complete-response households; comparisons are descriptive.

Fig 5: Survey-weighted saving channels in Uganda, 2019/20.

The contrast is informative precisely because the saving channels are not interchangeable. A farmer using a VSLA is not necessarily excluded from financial arrangements, while institutional or mobile saving is not the same as owning a formal account. SACCOs and microfinance institutions are included in a broad institutional category without assuming they all have identical regulatory status. National Uganda account ownership in the 2024 Findex row is 72.83%, but it should not be compared directly with the household saving percentages: dates, populations, units, and financial behaviors differ. The Uganda lens therefore complements the national prediction exercise rather than externally validating it. It points toward measuring suitable use and service design alongside ownership, not replacing one concept with another.

4. Discussion and 5. Conclusions

4.8. Policy Relevance, Limitations, and a Research Agenda

The strongest practical lesson is to match the analytical tool to the decision. For monitoring a previously surveyed country, the last observed account rate is a stronger starting point than the tested contemporaneous structural models. For countries absent from training, the remaining national error is large enough to discourage fine-grained allocation based on model scores alone. For livelihood-oriented service design, the Uganda saving-channel differences justify closer examination of group-based and institution-linked arrangements, not the automatic classification of farmers as financially excluded. Such questions are consistent with the importance of suitable savings access in farmer-focused research (Batista & Vicente, 2020) [2], but they require field evaluation before any particular product or outreach intervention can be recommended as effective.

Several limitations bound the findings. The panel is small, unbalanced, and defined by fixed income labels. National covariates cannot establish household mechanisms, and processed-file missingness may encode both statistical coverage and pipeline quality. Commodity medians are selective baskets; production value per adult is not farm productivity. The predictors omit many individual capabilities and local service conditions. Later-revised covariates make the temporal exercise unsuitable as a real-time forecasting demonstration. Fixed algorithm settings and one grouped partition limit conclusions about model optimality, while the bootstrap intervals condition on already fitted predictions. Uganda estimates use household weights

but no full survey-design variance calculation, so their contrasts are not accompanied by inferential claims. These limitations are part of the interpretation, not merely technical footnotes.

5. Conclusions

This study tests whether more agricultural information improves national financial inclusion prediction rather than assuming that it does. Across 294 country-years and 70 low- and lower-middle-income economies, core random forests and XGBoost improve on a mean benchmark but retain substantial held-out error. Six agricultural variables provide no consistent pooled improvement, and a later-wave holdout reinforces that caution. A complete-agriculture subset produces a modest XGBoost gain but not a consistent cross-model pattern. For the 2024 updating task, the last observed account rate clearly outperforms the structural models. The appropriate conclusion is therefore task- and representation-specific: these agricultural aggregates have limited demonstrated incremental value, not that agricultural livelihoods are irrelevant to financial inclusion.

The separate Uganda evidence shows why that distinction matters. Farming households rely more on savings groups and less on institution-or-mobile saving than other households in the observed weighted profiles. National ownership alone cannot reveal those arrangements or determine whether available services fit seasonal livelihoods. Future research should link household agricultural activities to verified financial-use measures within the same survey, evaluate genuinely forward-looking data vintages, and compare models with strong prior-outcome baselines. It should also assess calibration, stability across geographic groups, and the consequences of any targeting rule before deployment. The article's contribution is a reproducible boundary on predictive claims and a more careful bridge between national monitoring and livelihood-sensitive financial-service research.

References

1. Bashiru S, Bunyaminu A, Yakubu IN, Al-Faryan MAS. Drivers of financial inclusion: Insights from sub-Saharan Africa. Economies. 2023;11(5):146. doi:10.3390/economies11050146
2. Batista C, Vicente PC. Improving access to savings through mobile money: Experimental evidence from African smallholder farmers. World Development. 2020;129:104905. doi:10.1016/j.worlddev.2020.104905
3. Food and Agriculture Organization of the United Nations. (2025). *World food and agriculture—Statistical yearbook 2025.* https://doi.org/10.4060/cd4313en
4. Jama, A., Eshetu, F., & Ali, H. (2025). Drivers of financial inclusion in sub-Saharan Africa countries. *Cogent Social Sciences, 11*(1), Article 2590235. https://doi.org/10.1080/23311886.2025.2590235
5. Kapoor, S., & Narayanan, A. (2023). Leakage and the reproducibility crisis in machine-learning-based science. *Patterns, 4*(9), Article 100804. https://doi.org/10.1016/j.patter.2023.100804
6. Khera, P., Ng, S., Ogawa, S., & Sahay, R. (2022). Measuring digital financial inclusion in emerging market and developing economies: A new index. *Asian Economic Policy Review, 17*(2), 213–230. https://doi.org/10.1111/aepr.12377
7. Klapper L, Singer D, Starita L, Norris A. The Global Findex Database 2025: Connectivity and financial inclusion in the digital economy. World Bank; 2025. doi:10.1596/978-1-4648-2204-9
8. Maehara, R., Benites, L., Talavera, A., Aybar-Flores, A., & Muñoz, M. (2024). Predicting financial inclusion in Peru: Application of machine learning algorithms. *Journal of Risk and Financial Management, 17*(1), Article 34. https://doi.org/10.3390/jrfm17010034
9. Mhlanga, D., & Dunga, S. H. (2020). Measuring financial inclusion and its determinants among the smallholder farmers in Zimbabwe: An empirical study. *Eurasian Journal of Business and Management, 8*(3), 266–281. https://doi.org/10.15604/ejbm.2020.08.03.008
10. Mumtaz, M. Z. (2024). Financial inclusion, digital finance and agricultural participation. *Agricultural Finance Review, 84*(2/3), 93–113. https://doi.org/10.1108/AFR-10-2023-0132
11. Negera, M., Alemu, T., Hagos, F., & Haileslassie, A. (2025). Does financial inclusion enhance farmers' resilience to climate change? Evidence from rural Ethiopia. *Sustainable Development, 33*(2), 3008–3022. https://doi.org/10.1002/sd.3277
12. Saidy, S., & Hassan, A. D. (2025). Modeling drivers of fintech adoption in sub-Saharan Africa using machine learning. *Digital Finance, 7*(4), 1065–1091. https://doi.org/10.1007/s42521-025-00161-4
13. Schöbitz, L. (2026). *jmpdata: Input data and indicator estimates from the WHO/UNICEF Joint Monitoring Programme* (Version 0.1.0) [Computer software; snapshot released August 20, 2026]. https://github.com/openwashdata/jmpdata
14. Uganda Bureau of Statistics. (2021). *National Panel Survey 2019–2020* (UGA_2019_UNPS_v03_M) [Data set]. World Bank, Development Data Group. https://doi.org/10.48529/nqzx-f196
15. World Health Organization, & United Nations Children's Fund. (2025). *Progress on household drinking-water, sanitation and hygiene 2000–2024: Special focus on inequalities.* World Health Organization. https://www.who.int/publications/i/item/9789240115149
16. World Bank. (2026). *People using safely managed drinking water services, rural (% of rural population)* [Data set; World Development Indicators, extract updated July 13, 2026]. https://data.worldbank.org/indicator/SH.H2O.SMDW.RU.ZS

Appendices

Appendix A. Fixed Model Settings and Reproducibility

Table A1: Principal Settings

Model	Settings
Mean / ridge	Training-fold mean; ridge alpha = 10. Numeric values standardized inside each training pipeline.
Random forest	400 trees; minimum leaf size = 4; max_features = 0.8; random_state = 2026.
XGBoost	250 trees; depth = 2; learning_rate = 0.03; subsample = 0.8; colsample_bytree = 0.8; reg_lambda = 10; squared-error objective; seed = 2026.
Validation	Five country-grouped folds; 2,000 country bootstrap resamples; 20 permutation repeats/block/fold; predictions clipped to 0–100.
Software	scikit-learn 1.6.1; XGBoost 3.4.1. Analysis code reads the original uploaded ZIP and writes all derived results.

Reproduction order: install the listed Python dependencies; place the analysis script beside the original uploaded archive; run the script; compare its CSV outputs with the analysis bundle. No downloaded financial dataset is substituted for an uploaded file. Source hashes and an archive inventory are included in the bundle. The full code, not the abbreviated setting table, defines all defaults.

Appendix B. Weighted Saving Profiles and Measurement Boundaries

Table B1: Uganda Household Saving Profiles (%)

Group	n	Inst./mobile	VSLA/ROSCA	Bank	Mobile	No saving
All households	2987	34.46	34.28	16.84	14.24	13.14
Farming main income	1344	21.61	40.78	5.92	9.60	12.06
Other main income	1643	41.91	30.51	23.18	16.93	13.77
Rural	2268	25.43	38.47	9.57	11.64	14.14
Urban	719	53.95	25.23	32.56	19.84	10.99
Poor	417	4.22	33.09	1.89	1.02	25.41
Non-poor	2567	38.11	34.41	18.69	15.79	11.66

Inst./mobile = any CB06B/C/D/E/H = 1; VSLA/ROSCA = any CB06F/G = 1; bank = CB06B = 1; mobile = CB06H = 1; no saving = CB06Z = 1. Farming main income = s11q01 in {1,2}. Urban/rural uses gsec1 urban in {1,0}; household weights use gsec1 wgt; poverty uses pov2019_20 poor_2020. Three complete-saving households have no poverty label, so poor/non-poor counts do not sum to 2,987. Channel columns are not mutually exclusive. Household interviews occurred in 2019 and 2020.

Appendix C. Audit Trail and Scope of Claims

The bundle contains predictor coverage, pooled and fold-level performance, paired bootstrap intervals, temporal predictions, persistence predictions, complete-agriculture sensitivity results, income-label errors, and held-out permutation repeats. The country-year modeling table is included; raw household microdata are not redistributed. Retained bibliographic details were checked against DOI metadata, official source pages, or the package citation. Negera et al. is cited as the final 2025 issue, not the 2024 online-first version. The Uganda dataset DOI records publication year 2021, although version 03 was posted later. Unused suggested references are not presented as verified evidence. FAOSTAT snapshot publication dates are not inferred from observation years. All substantive results in this article are computed, descriptive, or predictive—not causal.

How to Cite This Article

Kamya AB, Alochere NJ. Beyond more data: do agricultural signals improve financial inclusion prediction? International Journal of Multidisciplinary Research and Growth Evaluation. 2026;7(6):144–153.

Creative Commons (CC) License

This is an open access journal, and articles are distributed under the terms of the Creative Commons Attribution NonCommercial-ShareAlike 4.0 International (CC BY-NC-SA 4.0) License, which allows others to remix, tweak, and build upon the work non-commercially, as long as appropriate credit is given and the new creations are licensed under the identical terms