INTERNATIONAL JOURNAL OF NOVEL RESEARCH AND DEVELOPMENT (IJNRD)
© 2026 IJNRD | Volume 11, Issue 10, October 2026 | ISSN: 2456-4184 | IJNRD.ORG

MERCHANT ACCEPTANCE TARGETING UNDER LIMITED SURVEY INFORMATION

Abdul Muqsit Khan, Arivuchelvan Gnanasekaran
Independent researchers, Richmond, Virginia, United States
Corresponding author: Abdul Muqsit Khan; abdulmuqsit.khan@icloud.com
Arivuchelvan Gnanasekaran: arivuganbu@gmail.com

Abstract: Payment providers need to decide where limited diagnostic outreach is most likely to encounter merchant acceptance problems. A controlled computational experiment tests whether the denominator used in a survey ranking, and smoothing of sparse estimates, change the accuracy of that decision. The experiment uses 57,408 account owners in 87 low- and middle-income economies from the 2024 survey wave of Global Findex 2025. Within 174 economy-by-residence cells, 500 randomized repetitions separate training and evaluation respondents and reveal four nested amounts of training information. Seven allocation rules are compared under identical samples and contact budgets. At the smallest information budget, 2,792 training responses, and a contact capacity of 20% of the evaluation population, ranking by the joint incidence of definite digital nonuse and a reported acceptance constraint yields 19.99 confirmed reports per 100 planned contacts. Ranking by the acceptance share among cash nonusers yields 18.25. The paired mean advantage is 1.74 reports, with a Monte Carlo standard error of 0.048; the central 95% of split-specific differences ranges from −0.16 to 4.00. Prespecified moderate smoothing improves on the conditional rule but underperforms the unsmoothed joint rule. Strong smoothing increases the rural contact share from 62.07% to 91.43% while lowering yield. The contribution is a reproducible decision benchmark quantifying denominator choice, information requirements and allocation shifts. These are experimental results about decision methods on observed survey data, not treatment effects on merchants or payment adoption.

IndexTerms: digital merchant payments; computational experiment; survey targeting; financial inclusion; information budget; merchant acceptance.

I INTRODUCTION

A payment provider considering expansion needs information about where customers encounter merchants that accept only cash. It may have enough resources for diagnostic contact with only a fraction of prospective customers. A survey can guide that effort, but different summaries of the same responses need not identify the same markets. The business decision is therefore not simply whether merchant acceptance matters. It is how to allocate a fixed contact budget when the underlying information is incomplete and geographically uneven.

The Global Findex 2025 report identifies gaps between financial access and everyday digital use and provides consumer explanations for cash use [1]. Those explanations offer an observable starting point for diagnostic targeting. However, the share of cash nonusers citing acceptance is a composition statistic. A provider contacting account owners more broadly needs the incidence of that explanation among the people it can contact. A place can have a high conditional share while containing relatively few cash nonusers. Smoothing sparse estimates creates an additional choice: it can stabilize scores but also move priorities toward the characteristics of the pooling group.

This study asks how these choices perform when they must make decisions about responses withheld from estimation. We conduct a controlled computational experiment using the authorized public-use microdata [2]. The manipulated factors are the information available to a decision rule, the rule itself and the permitted contact share. Actual survey responses supply the outcomes; new customer responses are not generated. All methods receive the same revealed training records and face the same evaluation records in each randomized repetition.

The contribution has three parts. First, we quantify the held-out contact yield associated with conditional and joint acceptance measures. The underlying probability identity is familiar; its decision consequences under sparse survey information require empirical evaluation. Second, we estimate an information learning curve across four survey budgets. Third, we assess whether a transparent smoothing rule improves targeting and document its effect on the rural share of contacts. The study contributes an application benchmark, rather than a new estimator or an assertion that a particular intervention works.

The findings favor a simple joint-incidence measure in the principal setting. With 2,792 revealed training responses, it identifies 19.99 confirmed reports per 100 planned contacts, compared with 18.25 using a conditional measure and 8.97 under uniform allocation. Moderate smoothing performs better than the conditional measure but worse than the unsmoothed joint measure. This negative result matters because smoother estimates may look more dependable while changing the contact population in ways that do not improve the chosen outcome.

II RELATED RESEARCH AND THE DECISION GAP

2.1 Acceptance evidence and intervention evidence

Global Findex measures whether adults use digital merchant payments and, among respondents routed to the cash-reason question, the main explanation for paying cash [1,3,4]. This supports a consumer-reported diagnostic endpoint. It does not directly audit merchant equipment, acceptance policies or transaction success. The distinction matters because selecting customers who report an obstacle is a different task from selecting merchants that will benefit from an intervention.

Allen et al. use country and merchant information with machine learning to examine incentives for digital payment acceptance and usage [5]. Their analysis connects infrastructure, transactions and incentives across the payment system. Kankanhalli et al. conduct a randomized field experiment with cash-only retailers in Mexico, testing customer success management during digital payment adoption [6]. That field experiment concerns merchant responses to an intervention. The present study concerns the performance of survey-based allocation rules before such an intervention is chosen.

CGAP's merchant-payments research explains why cash can remain attractive when digital acceptance is technically possible, including the practical value of the payment service to a business [7]. Consequently, high acceptance-report yield should lead to further diagnosis rather than an automatic recommendation to install terminals. A cash-only response could coexist with price, reliability, liquidity or habit barriers that a single-main-reason question cannot fully represent.

2.2 What this experiment adds

A targeted literature search on 13 September 2026 combined merchant acceptance, Global Findex, cash-only, targeting, ranking, holdout and cross-validation, and checked the primary studies above. The reviewed work does not establish how the present conditional and joint consumer-reason scores perform under randomized survey-information budgets. This is a bounded application gap. The search does not justify a claim that no related unpublished or unindexed research exists.

The experiment follows the principle of explicitly specifying aims, mechanisms, targets, methods and performance measures for computer experiments [8]. Its mechanism is randomized information masking and sample splitting on existing records. It evaluates decision rules against responses excluded from their fitting. Unlike a model that assumes a merchant intervention effect and then simulates the resulting gains, this benchmark does not need to invent an effect on adoption. Its practical output is an estimate of diagnostic concentration under specified decision constraints.

III DATA AND MEASUREMENT

3.1 Source and eligibility

The source is the World Bank's Global Findex 2025 public-use microdata, reference WLD_2024_FINDEX_v02_M, covering the 2024 survey wave [2]. The authorized Stata release contains 144,090 records from 140 economies. Those file counts differ from the headline report coverage and define the actual data available for this study. Income classification comes from the archived official country-level database, using its 2024 World Bank income-group field [9]. We include low- and middle-income economies and do not infer their classification from observed payment outcomes.

Eligible respondents have constructed account ownership equal to one, age at least 15, a finite positive survey weight, valid rural or urban classification, and a recorded response to the in-store digital payment question, fin25e2, in categories 1–4. Requiring a module response avoids treating a structural skip as a negative answer. This initial frame contains 59,133 respondents. We retain economies with at least 60 eligible respondents in each residence group, giving a final frame of 57,408 respondents in 87 economies and 174 cells. Five economies fail the cell-size rule: Lao PDR, Lebanon, Madagascar, Niger and Eswatini. No inclusion rule requires acceptance-positive cases.

Table 1 summarizes the frame and the randomized information budgets. The minimum cell size leaves at least 30 records in each training pool and evaluation group. The full frame includes 31,815 rural respondents and 25,593 urban respondents. These are unweighted counts; they do not determine the evaluation distribution. The underlying survey uses country-specific sampling and weighting procedures, and the public-use file does not supply the original design identifiers needed to reproduce every national variance calculation [10].

Table 1. analysis frame and randomized information budgets

Item	Count	Definition
Source records	144,090	Actual authorized Stata release
Eligible before cell threshold	59,133	Account owners with valid module, age, residence and weight
Final analysis frame	57,408	87 economies; 174 economy-by-residence cells
Confirmed acceptance reports	4,161	Definite nonuse and cash-only merchants as main reason
Uncertain responses	2,280	Unknown use status or nonuse without a substantive reason
Evaluation records per repetition	28,742	Never used to score geographic cells
Revealed training records	2,792 / 7,100 / 14,287 / 28,666	10% / 25% / 50% / 100% of the training pool

Source: original eligibility and split calculations from [2], with income classification from [9]. Counts are unweighted. Every repetition uses the same counts but randomly changes which records are revealed and evaluated.

3.2 Confirmed reports and uncertain responses

The primary endpoint Y equals one if a respondent reports definite digital in-store nonuse, fin25e2=2, and selects merchants accepting only cash as the main reason, fin25e4=1. It equals zero otherwise. The nonuse indicator C equals one for fin25e2=2. There are 4,161 confirmed acceptance reports in the analysis frame. Because Y identifies a confirmed report, a zero is not a verified statement that every encountered merchant accepts digital payment.

The analysis retains 2,280 respondents with uncertain use status or a missing, uncertain or refused substantive cash-use reason. Keeping them in the denominator prevents silent complete-case selection. It also limits interpretation: the primary outcome is the concentration of confirmed reports, not the prevalence of every latent acceptance constraint. A prespecified sensitivity assigns all such uncertain responses to acceptance-positive status. For uncertain use status, it also assigns presumed nonuse so that an acceptance case remains within the presumed nonuser group. This deliberately extreme scenario is not an imputation model or a confidence bound for the primary paired effect.

The questionnaire asks for one main reason [3]. The endpoint cannot measure the full set of obstacles faced by a person, the number of cash-only merchants, or the proportion of transactions affected. These limits apply equally to all compared methods and should accompany any business use of the ranking.

IV EXPERIMENTAL DESIGN

4.1 Randomized information availability

The internal experiment protocol was recorded on 13 September 2026 after checking eligibility counts and before running the method comparisons. It was not registered in an external registry. Earlier work had already examined the same microdata's rural associations and reason distributions. The present exercise therefore supplies internal held-out evaluation within an already studied dataset, rather than a previously untouched external validation sample.

In each of 500 repetitions, records are randomly permuted separately within each economy-by-residence cell, without using their outcomes. The first floor(n/2) records form a candidate training pool; the rest form the evaluation group. Training pools total 28,666 records and evaluation groups total 28,742 records in every repetition. The generator uses PCG64 with a deterministic seed derived from master seed 20260913 and the repetition number. The full 500-repetition count and seed were fixed before comparison results were observed.

Four nested information conditions reveal 10%, 25%, 50% or 100% of each training pool, rounded down with a minimum of two records. These produce 2,792, 7,100, 14,287 and 28,666 revealed records, respectively. The percentages refer to the training half, not the entire analysis frame. Within a repetition, every rule sees identical revealed records at a given information budget. Evaluation labels remain unavailable to the scoring rules and to the selection of smoothing strength.

4.2 Decision rules and the denominator

Table 2 defines the competing rules. The conditional acceptance score estimates the weighted share of confirmed acceptance reports among definite cash nonusers. The joint score estimates the weighted share of confirmed reports among all observed module respondents. For each cell, the two are linked by Equation 1 when the conditional denominator is nonzero. Thus different nonuse rates can reverse their rankings.

$$\Pr(Y = 1 \mid g) = \Pr(C = 1 \mid g) \times \Pr(Y = 1 \mid C = 1, g) \quad (1)$$

Here g denotes an economy-by-residence cell, Y the confirmed report and C definite nonuse. Since Y=1 implies C=1 in the primary coding, the identity is exact. Its novelty is not claimed. The experiment asks how much choosing a denominator aligned with the contact population changes out-of-sample allocation under finite information. Conditional ranking remains a reasonable target for a different operation restricted to known cash nonusers; this experiment evaluates a broader account-owner contact frame.

Uniform allocation and nonuse-only ranking provide additional reference rules. If a revealed cell contains no definite nonusers, the conditional score uses a pooled same-residence conditional rate computed from revealed data across economies. This fallback prevents undefined rankings in sparse samples. The pooled ratio uses the means of the cell joint and nonuse rates; it does not access evaluation responses.

Table 2. prespecified decision rules

Rule	Score or allocation	Role
Uniform	Same contact fraction in every geographic cell	Reference
Nonuse	Weighted definite nonuse share	Broad diagnostic baseline
Conditional acceptance	Weighted acceptance-positive total divided by definite nonuse total	Composition measure
Joint acceptance	Weighted acceptance-positive share of all module respondents	Incidence measure
Smoothed joint	Joint share pooled toward same-residence mean using Equation 2	k = 20 primary; 5 and 100 sensitivity

Source: experiment protocol. All scores use revealed training responses only. Conditional scores with no revealed nonusers use the same-residence pooled conditional rate. The three smoothing strengths are separate experimental arms, giving seven rules in total.

4.3 Smoothing and geographic allocation

The smoothed joint score moves a cell estimate toward the mean joint rate across revealed cells with the same rural or urban status. Equation 2 gives the rule. The effective sample size uses the Kish weight concentration expression. The fixed strength k=20 is the primary smoothing specification; k=5 and k=100 are prespecified weaker and stronger alternatives. These are heuristic regularized scores, not exact posterior probabilities from a fitted hierarchical model.

$$\tilde{p}_g = \frac{n_{eff,g}\hat{p}_g + k\,\hat{\mu_r}}{n_{eff,g} + k}, \quad n_{eff,g} = \frac{(\Sigma w_i)^2}{\Sigma w_i^2} \qquad (2)$$

In Equation 2, p-hat is the revealed cell's weighted joint share, mu-hat is the simple mean of revealed joint shares across economies with residence status r, and n-eff is the cell's effective sample size. Pooling means are recalculated from training records in every repetition. The experiment does not tune k using evaluation outcomes. Because rural and urban pooling means can differ, stronger smoothing can change the residence composition of selected contacts as well as their accuracy.

All methods rank the 174 geographic cells and assign a constant contact probability within each cell. The target distribution gives each economy equal total evaluation weight while retaining the source-weighted residence mix within that economy. Contact capacity is 20% of this evaluation distribution in the primary setting; 10% and 40% are secondary capacities. Higher-score cells are filled first. At the boundary, all cells tied at the same score share the remaining contact fraction proportionately. This avoids arbitrary selection by country name and permits exact budget matching.

4.4 Outcomes and uncertainty

The principal performance measure is confirmed acceptance reports per 100 planned contacts, calculated in Equation 3. Evaluation weights v sum to one, and a is the method's contact probability for a held-out respondent. The contact constraint makes the denominator equal to the assigned capacity q. Additional measures are the fraction of all positive reports covered and the rural share of contacts. The algorithm uses evaluation-frame membership and weights to set capacity, but never uses evaluation outcomes to rank cells.

$$\text{Yield} = 100 \times \frac{\Sigma(v_i a_i Y_i)}{\Sigma(v_i a_i)}, \quad \Sigma(v_i a_i) = q \qquad (3)$$

The primary prespecified contrast is the moderately smoothed joint rule minus conditional ranking, at 2,792 revealed records and 20% contact capacity. Joint minus conditional and smoothed minus unsmoothed comparisons are also reported, regardless of sign. All seven rules and all four information budgets are retained. There is no significance-based stopping or removal of settings that favor a comparator.

For each method, we report the mean across repetitions and the empirical 2.5th and 97.5th percentiles of split-specific performance. Paired method differences use the same split in both arms. Monte Carlo standard errors equal the standard deviation across repetitions divided by the square root of 500 [8]. They measure computational precision conditional on this empirical file. The split percentiles describe allocation instability. Neither measure is a population confidence interval, and 500 repetitions do not constitute 500 independent field studies.

V RESULTS

5.1 Accuracy under scarce survey information

Table 3 and Figure 1 report the primary setting. Uniform allocation yields 8.97 confirmed reports per 100 planned contacts. Nonuse ranking raises yield to 16.87, and conditional acceptance ranking to 18.25. The unsmoothed joint rule reaches 19.99. Its mean paired advantage over conditional ranking is 1.74 reports per 100 contacts, or approximately 9.5% relative to the conditional rule's mean yield. This relative change concerns a reported diagnostic endpoint, not increased digital usage or revenue.

Table 3. held-out performance with 2,792 revealed records and 20% contact capacity

Rule	Reports per 100 contacts	Central 95% of splits	MCSE	Positive reports covered %
Uniform	8.97	8.63–9.30	0.008	20.00
Nonuse	16.87	14.12–19.38	0.061	37.64
Conditional acceptance	18.25	15.30–20.94	0.065	40.71
Joint acceptance	19.99	17.46–22.06	0.054	44.58
Smoothed joint k = 5	19.75	17.26–21.91	0.055	44.06
Smoothed joint k = 20	19.43	16.77–21.68	0.057	43.35
Smoothed joint k = 100	18.14	15.91–20.69	0.056	40.46

Source: original experiment on [2]. Means across 500 paired randomized splits. Central ranges are empirical split percentiles; MCSE is the Monte Carlo standard error. Neither is a population confidence interval. Each economy receives equal total evaluation weight.

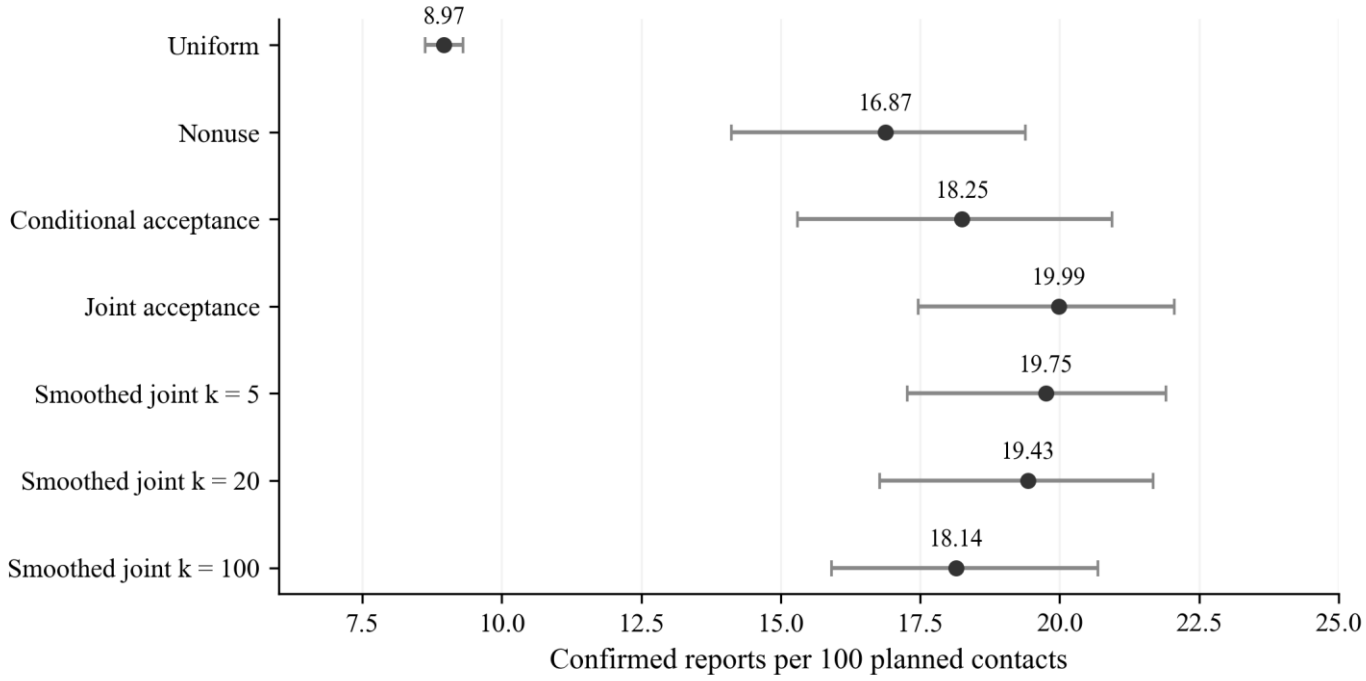

Figure 1. mean diagnostic yield and variation across randomized splits

Source: original experiment on [2]. Points are mean yields; horizontal bars span the central 95% of randomized split outcomes, not population confidence intervals. All rules use 2,792 revealed records and 20% contact capacity.

The primary prespecified moderately smoothed rule yields 19.43 reports per 100 contacts, 1.18 more than conditional ranking. Its paired Monte Carlo standard error is 0.050, and the central 95% of split differences runs from −0.94 to 3.52. The same smoothed rule yields 0.55 fewer reports than unsmoothed joint ranking. We retain this unfavorable comparison: the experiment does not support recommending moderate smoothing solely because its estimates appear less variable.

The unsmoothed joint advantage over conditional ranking is positive in 96.6% of randomized repetitions. The central 95% of its paired differences is −0.16 to 4.00, with Monte Carlo standard error 0.048. Thus the average advantage is precisely computed, while an individual limited survey can still produce a reversal. Reporting only the much smaller Monte Carlo error would obscure this operational uncertainty.

At a 20% contact budget, joint ranking covers an average 44.58% of confirmed positive reports, compared with 40.71% for conditional ranking and exactly 20% for uniform allocation. These figures quantify concentration within the chosen evaluation distribution. They do not establish that a provider can physically contact this cross-country population at uniform cost.

5.2 Information learning curves

Table 4 and Figure 2 show how performance changes as more training information is revealed. The joint rule increases from 19.99 to 22.24, 23.33 and 23.96 reports per 100 contacts as revealed records increase from 2,792 to 7,100, 14,287 and 28,666. The conditional rule follows the same broad pattern, reaching 22.85 at the largest budget. The joint advantage is positive on average at every prespecified information level, although it narrows to 1.11 at the largest level.

Table 4. reports per 100 planned contacts as survey information increases

Rule	2,792 records	7,100 records	14,287 records	28,666 records
Uniform	8.97	8.97	8.97	8.97
Nonuse	16.87	17.94	18.45	18.80
Conditional acceptance	18.25	20.65	21.99	22.85
Joint acceptance	19.99	22.24	23.33	23.96
Smoothed joint k = 5	19.75	22.14	23.29	23.96
Smoothed joint k = 20	19.43	21.97	23.21	23.94
Smoothed joint k = 100	18.14	20.80	22.67	23.79

Source: original experiment on [2]. Contact capacity is 20%; entries are mean held-out yields across 500 repetitions. Training records are nested within each repetition. The complete replicate file contains split variability for every entry.

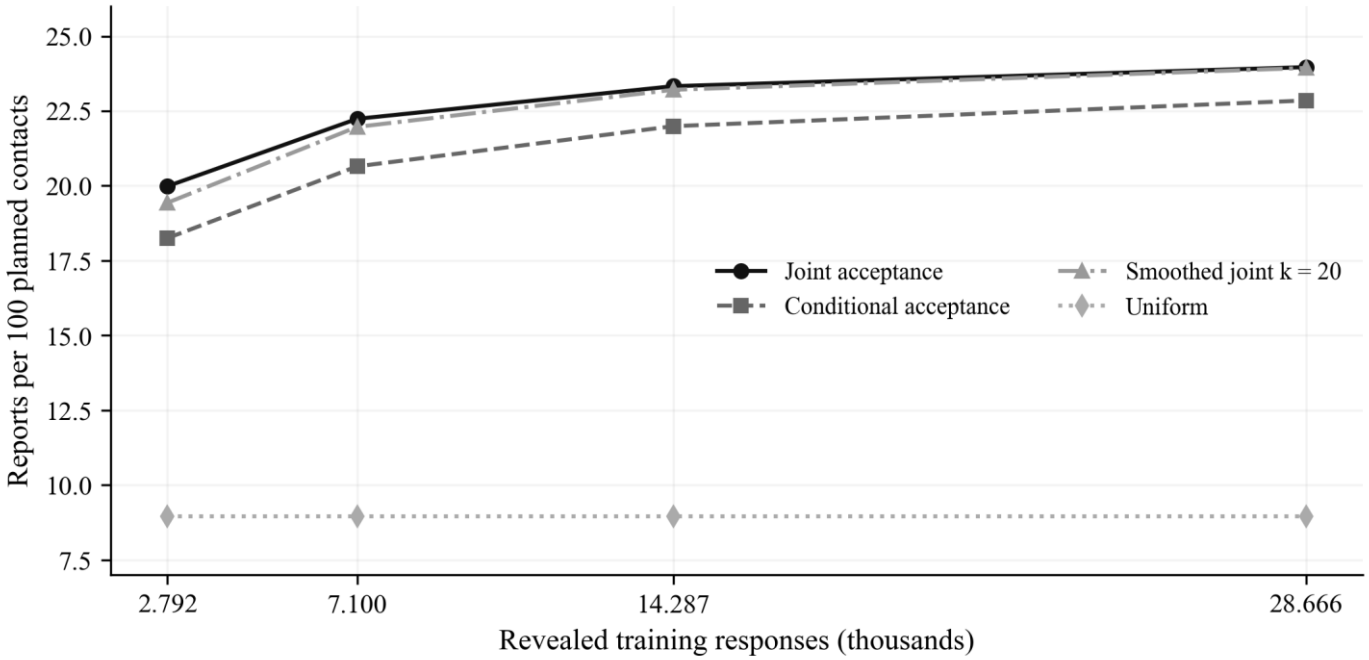

Figure 2. diagnostic yield as training information increases

Source: original experiment on [2]. Means over 500 paired repetitions at 20% contact capacity. All seven rules are reported in Table 4; the plot shows four reference curves. Information levels are the actual numbers of revealed training responses.

The marginal improvement becomes smaller as information increases. Moving from 2,792 to 7,100 revealed records raises mean joint yield by 2.25 reports per 100 contacts. Moving from 14,287 to 28,666 raises it by 0.64. These changes show diminishing gains within this design. They cannot determine an economically optimal survey size without actual interview costs, outreach costs and the value of a successful diagnostic contact. The counts are a reproducible information benchmark, not a recommendation to survey exactly these numbers in a new country.

Weak smoothing approaches the unsmoothed rule as training information increases. Moderate and strong smoothing also close much of their initial gap at the largest budget. This pattern is consistent with their formula, which gives the cell's own estimate more weight when its effective sample size is larger. The empirical contribution is the size and geographic consequence of the tradeoff in this dataset, rather than the general mathematical fact that a fixed penalty matters less with more information.

5.3 Rural allocation and robustness

Figure 3 shows a substantial residence shift under stronger smoothing. At the smallest information budget, the joint rule assigns 62.07% of contacts to rural respondents. The corresponding shares are 65.43%, 72.88% and 91.43% under k=5, 20 and 100. Mean yield falls to 19.75, 19.43 and 18.14, respectively. Because the pooling means are residence-specific, strong smoothing makes residence-level differences more influential and suppresses differences among economies within a residence group. This is a property of the tested rule, not evidence that rural targeting is undesirable.

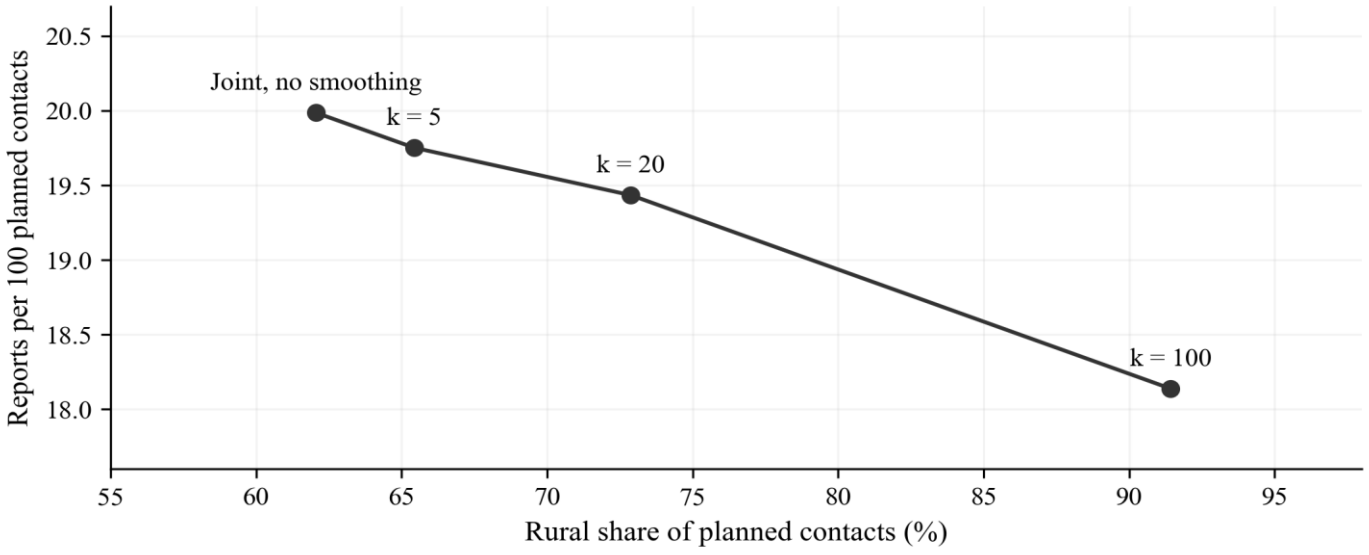

Figure 3. rural allocation and yield under different smoothing strengths

Source: original experiment on [2]. Means over 500 repetitions with 2,792 revealed records and 20% contact capacity. Greater same-residence smoothing changes both the rural contact share and confirmed-report yield. The connecting line orders smoothing strengths and does not trace a fitted response curve.

Rural inclusion can be a separate policy objective. If a provider values it independently of confirmed-report yield, the allocation should state that objective explicitly. A smoothing parameter otherwise risks functioning as an implicit geographic preference. The experiment measures this change; it does not attach a social welfare value to the resulting contact mix.

Table 5 reports the prespecified contact-capacity and uncertain-response sensitivities. With 10% and 40% contact capacity, the joint-minus-conditional mean advantages are 3.02 and 0.76 reports per 100 contacts, respectively. The larger advantage at the tighter budget shows that denominator choice matters most when the rule must concentrate strongly on high-score cells. Under the extreme uncertain-response coding, the advantage remains positive on average at 2.22, while yield levels rise mechanically because more responses are defined as positive. The scenario does not identify the true status of uncertain respondents.

Table 5. paired contrasts and prespecified sensitivity conditions

Setting	Paired contrast	Mean difference	Central 95% of splits	MCSE
Primary contrast	Smoothed k = 20 minus conditional	1.18	-0.94–3.52	0.050
Joint comparison	joint minus conditional	1.74	-0.16–4.00	0.048
Smoothing comparison	Smoothed k = 20 minus joint	-0.55	-2.26–0.84	0.034
10% contacts	joint minus conditional	3.02	-0.45–6.52	0.083
40% contacts	joint minus conditional	0.76	0.03–1.58	0.019
Uncertainty upper scenario	joint minus conditional	2.22	-0.16–4.94	0.057

Source: original experiment on [2]. All settings reveal 2,792 training records. Differences and MCSEs are reports per 100 planned contacts. Contact capacity is 20% unless specified. The upper scenario codes every uncertain response as an acceptance-positive presumed nonuser; it is not an estimated missing-data correction.

5.4 Verification

An independent respondent-level calculation reproduced all seven rules across all prespecified budgets and coding scenarios for repetitions 0, 17 and 499. The 273 matched comparisons agreed to less than 0.0000000001 percentage points. Separate checks confirmed training/evaluation disjointness, proportional treatment of tied scores, exact contact capacity and the equivalence of uniform scoring to uniform allocation. The complete output contains 45,500 method-setting-repetition records, including every specified sensitivity. Aggregate inputs, programs, seeds, protocol and results accompany the manuscript; respondent records are excluded from the distributable supplement.

VI BUSINESS INTERPRETATION AND LIMITATIONS

The actionable lesson concerns how to prepare diagnostic outreach. A dashboard intended to guide contact with account owners should show acceptance reports per account owner alongside the share among cash nonusers. The latter answers a valid descriptive question but can misalign the ranking with the contact population. In this experiment, the joint measure improves average held-out yield without requiring more training responses or a more complex model. A provider can test that comparison on its own current contact frame before adopting it.

The results also support reporting allocation consequences alongside accuracy. A method can produce plausible, smooth scores while changing which populations receive attention. Here, stronger residence-specific smoothing sharply increases rural contact allocation and reduces average confirmed-report yield. A transparent implementation would show both statistics and allow the business to choose an explicit rural-inclusion constraint if one is desired. The experiment has not established an optimal fairness constraint, and its unconstrained ranking should not be interpreted as one.

Several restrictions limit transfer. First, the intervention is on the information supplied to a decision algorithm. No customers or merchants were newly assigned to treatment, contacted, or followed for payment outcomes. These results are computational experimental evidence, not a randomized estimate of the effect of onboarding, subsidies, education or merchant support. Second, the endpoint is a consumer's main explanation, not audited acceptance. Concentrating reports can improve a diagnostic starting point while failing to identify customers whose behavior would change after an intervention.

Third, the experiment uses random internal splits from the same 2024 survey wave and the same economies. It does not test a later period, a previously unseen economy or a local commercial customer base. Earlier examination of the full file also means that a formally untouched external test set was not available. Randomized splitting protects each rule from fitting its own evaluation responses, but cannot undo prior knowledge used to choose the research question.

Fourth, all countries receive equal total evaluation weight. The results represent an equal-economy benchmark across the included countries, not a worldwide adult-population estimate. Economy-by-residence cells are broad, contact costs are assumed equal, the contact frame is treated as known, and capacity can be divided fractionally within a cell. Actual operations face consent, reachability, location, acquisition-cost and service-coverage constraints absent from the data. Incorporating those constraints may change the preferred ranking.

Fifth, Monte Carlo precision does not solve survey-design or external-validity uncertainty. Original primary-sampling-unit and stratum identifiers are unavailable in the analysis file, and the included countries are not a random sample of all possible markets. The empirical split ranges quantify one controlled source of instability. They should not be used as population intervals or converted into strong claims about business returns.

A subsequent field study could first compare the diagnostic contact yield of joint and conditional rankings in a current, consented contact frame under equal budgets. After validating the reported obstacle against merchant conditions, a separate randomized intervention could test onboarding assistance or another specified remedy. That sequence would connect the present methodological

evidence to actual business outcomes while keeping the two questions distinct. The current study supplies the completed computational benchmark; it does not report that future field study as conducted.

VII CONCLUSION

This controlled computational experiment shows that denominator choice has measurable consequences for merchant-acceptance diagnostic targeting under limited survey information. Across 500 randomized repetitions on 57,408 existing survey records, the joint acceptance measure improves mean held-out contact yield over the conditional measure at every prespecified information budget. Moderate smoothing improves on conditional ranking but underperforms the unsmoothed joint rule in the primary setting, and strong smoothing substantially shifts contacts toward rural respondents. The research contribution is the quantified decision comparison, learning curve and geographic tradeoff. Its findings support clearer measurement and local validation before a provider commits resources to an intervention.

VIII DECLARATIONS

Data and code availability: The official microdata are available from the World Bank catalog under its access terms [2]. The supplement includes the experiment protocol, reproducible programs, aggregate results, chart data and verification records. Raw respondent records are not redistributed. Reproduction requires obtaining the specified authorized Stata release and the official aggregate classification file [9].

Research ethics: The study uses anonymized public-use secondary survey records and controlled computational allocation. No new participants were recruited, no respondent was contacted, and no real-world treatment was administered. No new ethics approval or participant consent is represented as having been obtained.

Funding: None.

Competing interests: The authors declare no competing interests.

REFERENCES

[1] Klapper, L., Singer, D., Starita, L., & Norris, A. (2025). The Global Findex Database 2025: Connectivity and Financial Inclusion in the Digital Economy. Washington, DC: World Bank. doi: 10.1596/978-1-4648-2204-9

[2] World Bank, Development Research Group. (2025). World: The Global Findex Database 2025, 2024 survey wave [Anonymized public-use microdata]. Reference WLD_2024_FINDEX_v02_M. Stata release findex_microdata_2025_labelled_update112425.dta. Retrieved 9 September 2026. doi: 10.48529/bk9n-8r43

[3] World Bank. (2025). Global Findex 2025 questionnaire: English. https://thedocs.worldbank.org/en/doc/be6615202d1f08a25855c8ac2d615122-0050012025/related/Findex-English.pdf

[4] World Bank. (2025). Global Findex 2025 microdata codebook. Documentation for WLD_2024_FINDEX_v02_M. https://microdata.worldbank.org/catalog/7860/download/351895

[5] Allen, J., Carbo-Valverde, S., Chakravorti, S., Rodriguez-Fernandez, F., & Pinar Ardic, O. (2022). Assessing incentives to increase digital payment acceptance and usage: A machine learning approach. PLOS ONE, 17(11), e0276203. doi: 10.1371/journal.pone.0276203

[6] Kankanhalli, S., Anderson, S. J., Iacovone, L., & Narayanan, S. (2026). Driving merchant adoption of digital payment solutions through customer success management. Management Science. Advance online publication, 12 June. doi: 10.1287/mnsc.2024.04703

[7] CGAP. (2019). Cash is king in merchant payments. October. https://www.cgap.org/research/publication/cash-is-king-in-merchant-payments

[8] Morris, T. P., White, I. R., & Crowther, M. J. (2019). Using simulation studies to evaluate statistical methods. Statistics in Medicine, 38(11), 2074–2102. doi: 10.1002/sim.8086

[9] World Bank. (2025). The Global Findex Database 2025: Data download and documentation. Country-level database [CSV dataset]. Retrieved 6 September 2026. https://www.worldbank.org/en/publication/globalfindex/download-data

[10] World Bank. (2025). Survey methodology. Appendix A of The Global Findex Database 2025: Connectivity and Financial Inclusion in the Digital Economy.

Copyright & License:

© Authors retain the copyright of this article. This work is published under the Creative Commons Attribution 4.0 International License (CC BY 4.0), permitting unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.