IJBITA

International Journal of Business Integrity and Technology Advancements (IJBITA)

Volume – 02
Issue – 01
(2026)

AI-Driven Customer Intelligence and Predictive Analytics for Personalized Marketing and Customer Retention in U.S. Financial Services

*Md. Rezaul Hasan, Department of Management, Richards College of Business, University of West Georgia, Carrollton, USA
mhasan1@my.westga.edu
Masrul Hasan, College of Graduate and Professional Studies, Trine University, Angola, USA
mhasan23@my.trine.edu
Md. Saiful Islam, College of Graduate and Professional Studies, Trine University, Angola, USA
saifulcu105@gmail.com
Kazi Rakib Hasan Saurav, College of Graduate Studies, Central Michigan University, Mount Pleasant, USA
university-saura1k@cmich.edu
Md. Fazle Alahi Bhuiyan, College of Business Administration, Central Michigan University, Mount Pleasant, USA
fazleelahi.bhuiyan@gmail.com
Md. Arifur Rahman, College of Graduate and Professional Studies, Trine University, Angola, USA
rahman.arifur11@gmail.com

Article Info:
Received: 30 Jan, 2026
Processed: 03 Mar, 2026
Published: 02 Apr, 2026

Keywords:
Customer intelligence; financial services; personalized marketing; incremental effects; customer privacy; model risk.

CC BY SA

ABSTRACT

Financial institutions use customer signals to select messages, offers, and service actions. A predicted response does not show whether a contact changes the outcome. This structured narrative review uses evidence available through February 28, 2026. FDIC household data show rising primary mobile access between 2019 and 2023. A separate 2025 survey found mobile apps were the most-used banking method for 54% of surveyed bank customers. The surveys describe channel use in different populations (FDIC, 2024; American Bankers Association, 2025).

The proposed framework checks permission and eligibility before ranking, estimates the incremental value of a contact, and retains a no-contact option. It uses chronological validation, randomized holdouts, contribution margin, and customer safeguards. A hypothetical campaign illustrates how an apparent response advantage can fail to cover costs. The framework keeps marketing, service, and lending decisions distinct and draws on the February 2026 financial-sector AI risk framework for governance (U.S. Department of the Treasury, 2026).

Mobile access informs channel choice, while local experiments would establish whether a particular offer benefits eligible customers. The review uses no private records and estimates no new institution-level effects.

INTRODUCTION

A bank can see that a customer recently used its mobile app, received a direct deposit, and opened a product information page. Those observations can help the bank understand what information might be useful. They cannot establish that the customer wants another message, that the product is appropriate, or that a message will create any additional value. The analytical problem is therefore a decision under uncertainty: determine what action, if any, should be taken for an eligible customer at a particular moment.

The distinction matters because financial relationships are persistent. A short campaign may raise clicks while degrading trust, increasing complaints, or shifting an account opening that would have happened without the offer. In this setting customer intelligence means the organized use of permitted information to understand needs, channel preferences, and potential actions. It includes

Corresponding email: mhasan1@my.westga.edu

measurement and governance as well as prediction. Marketing analytics research describes how organizations can connect these decisions across channels and over the customer journey (Wedel & Kannan, 2016; Lemon & Verhoef, 2016).

This article asks three questions. First, what do available U.S. access data and earlier marketing studies establish about the context for personalization? Second, how can an institution separate response prediction from the expected incremental benefit of an action? Third, what controls are needed to make the resulting decisions reviewable? The answer is a decision framework with an explicit no-contact option, treatment assignment before measurement, and independent permission and eligibility checks.

The scope is consumer financial services in the United States, with banking examples where public access measures are strongest. Application to insurance or investment services would require different outcomes and product controls. The 2026 revision includes evidence that was publicly available by February 28, 2026; subsequent releases and updates are outside its scope.

EVIDENCE AND RESEARCH METHOD

This structured narrative review used targeted searches for banking channel use, personalization, privacy and trust, causal marketing measurement, and model governance. It prioritized original studies, official U.S. publications, and dated public releases. February 28, 2026 was the eligibility cutoff for the cited version. The search was purposive and does not estimate the prevalence of findings across the literature.

Three evidence categories were kept separate: published statistics for defined populations, empirical findings within a study's design, and proposed workflows or hypothetical arithmetic. A randomized retail experiment can guide a bank's experiment design even when it cannot predict the bank's lift (Gordon et al., 2019; Fang et al., 2026).

The quantitative comparison transcribes the three survey-year distributions in FDIC Table 2.1 and computes contrasts of the rounded percentages. Its denominator is banked households that accessed an account in the previous twelve months. The national FDIC unbanked estimate instead uses all households. The appendix documents the arithmetic (FDIC, 2024).

Two newer surveys provide context for the historical FDIC series. The ABA and Morning Consult report is based on a weighted online survey fielded in October 2025. The Federal Reserve's May 2025 report covers adults surveyed in 2024. The analysis uses no microdata, pools no samples, and computes no significance tests for survey contrasts (American Bankers Association, 2025; Federal Reserve Board, 2025).

PUBLISHED EVIDENCE ON BANKING ACCESS

Table 1 reports the primary method of account access among banked U.S. households that accessed an account in the prior twelve months. The estimates come from one FDIC report, allowing a comparison of the same named categories over three survey years (FDIC, 2024). Each entry is a percentage of the stated eligible household group. Multiple channels may have been used even when one channel was identified as primary.

Primary method	2019	2021	2023	Change, 2019–23
ATM	19.5%	16.0%	13.8%	−5.7 pp
Bank teller	21.0%	14.9%	15.1%	−5.9 pp
Mobile banking	34.0%	43.5%	48.3%	+14.3 pp
Online banking	22.8%	22.0%	19.8%	−3.0 pp
Telephone banking	2.4%	2.9%	2.3%	−0.1 pp
Other	0.3%	0.7%	0.7%	+0.4 pp

Table 1: Primary method of bank account access

Source: FDIC 2023 household survey, Table 2.1, released November 2024 (FDIC, 2024). Denominator: banked U.S. households that accessed an account in the prior 12 months. pp = percentage points.

Mobile banking's reported share rose from 34.0% in 2019 to 48.3% in 2023, a difference of 14.3 percentage points. Bank tellers declined by 5.9 points and ATMs by 5.7 points over the same interval. Online banking, separately classified from mobile banking, declined by 3.0 points. These changes are arithmetic contrasts of published estimates, with no inference about the significance of small year-to-year movements.

Corresponding email: mhasan1@my.westga.edu

A household preference does not count people, app sessions, transactions, or opportunities to contact a customer. The shares cannot be applied to an institution's customer count. Keep several accessible routes available, then use local data to estimate how many eligible customers can be reached on each channel.

A. CHANNEL MOVEMENT AND CHOICE

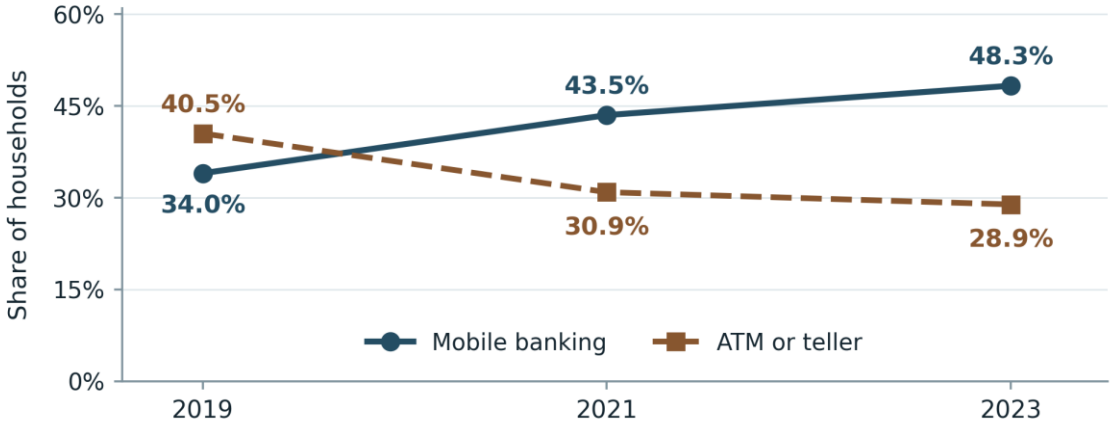


Fig 1: Mobile and in-person primary account access, 2019–2023

Source: FDIC 2023 household survey, Table 2.1 (FDIC, 2024). In-person = ATM + bank teller. Denominator: banked U.S. households that accessed an account in the previous 12 months.

In the FDIC series, 28.9% of eligible households used a teller or ATM as their primary method in 2023. In a separate October 2025 online survey, 54% of bank customers reported using mobile apps most often, 22% used desktop or laptop banking, 9% used branches, and 6% used ATMs. The ABA and FDIC percentages measure different populations and do not form a trend (FDIC, 2024; American Bankers Association, 2025).

The FDIC estimated that 4.2% of all U.S. households, about 5.6 million, were unbanked in 2023. The Federal Reserve's 2024 survey reported 6% of adults as unbanked when neither they nor a spouse or partner held a checking, savings, or money market account. These are separate measures of access with different units and definitions (FDIC, 2024; Federal Reserve Board, 2025).

Let existing customers update their channel preference; a lack of app activity can reflect an access barrier or a strong branch relationship. Reaching unbanked prospective customers requires a separate recruitment or community-partnership design and measures of access as well as sales.

B. PERSONALIZATION TIMING AND RELEVANCE

Research on online advertising indicates that a detailed message can be useful at one stage of a decision and intrusive or poorly timed at another. Lambrecht and Tucker's (2013) study of retargeting showed why information specificity should be related to the state of a consumer's search rather than assumed to improve every impression. Bleier and Eisenbeiss (2015) likewise examined the interaction of content, timing, and placement. Both motivate an action-specific hypothesis, not a universal effect size for banking.

Consider a person who reads information about checking accounts but has not requested help. One policy offers a neutral comparison guide after a permitted interval. A second policy sends a product-specific invitation. A third makes no contact. The first two policies are different treatments: their promises, effort, and possible perceived intrusiveness differ. Comparing them requires consistent eligibility and an outcome that captures sustained use or informed choice rather than the click alone.

The same logic applies to a service event. An unexpected payment problem may increase activity in digital logs, but the appropriate next step may be to resolve the problem. A sales message during the incident could reduce rather than increase customer value. The institution should record the service state separately from the marketing state, and use an independent gate to prevent conflicting actions.

Huang and Rust (2021) and Davenport et al. (2020) place AI within a wider marketing process. One recent example underscores why the workflow matters: in a February 2026 preprint, Fang et al. (2026) report that randomized AI push messages and advertising titles did not generate statistically significant sales gains in their online retail tests. A bank should test its own messages against an appropriate baseline.

C. CUSTOMER TRUST AND PERMISSION

Corresponding email: mhasan1@my.westga.edu

Personalized content exposes an institution's use of information to the customer. A person who receives a helpful reminder based on a recent explicit request may see relevance. The same person may regard a message derived from an undisclosed inference about household stress as intrusive. Aguirre et al. (2015) examine how information collection and trust-building interact in online advertising; Tucker (2014) studies how privacy controls affect responses to personalized ads. Those studies motivate a testable trust mechanism, but their results should not be treated as banking-specific estimates.

The framework records permission as structured state. A permission record contains the purpose, channel, source, effective time, expiration where applicable, and any withdrawal. A feature may be technically available yet unsuitable for a proposed purpose. The system should resolve that difference before scoring rather than depend on a model to learn which uses are appropriate. A change in preference should take effect across channels before queued messages are released.

Trust also needs observable indicators. Complaints, opt-outs, repeated contacts about the same topic, and failed handoffs each reveal a different potential problem. Low complaint counts can coexist with harm when customers lack a clear path to report it. Pairing behavioral measures with a structured customer feedback route gives reviewers more useful information than relying on engagement alone.

The permission gate does not require that all customers be observed equally. A safe no-contact decision can be taken with limited data. Evaluation should still examine whether the resulting distribution of useful information systematically favors people who already interact most often through digital channels.

D. DISTINCT FINANCIAL DECISIONS

Table 2 separates three uses of similar software: marketing, service, and credit. Each has different outcomes and approval requirements. Marketing personnel should not be able to alter eligibility rules or turn an audience score into a pricing decision.

Decision	Suitable output	Boundary for the proposed workflow
Marketing	Approved information or outreach, or no contact	Check permission and independent product eligibility. Keep pricing and underwriting separate.
Service	Account help or referral to an employee	Authenticate before showing account details. Escalate disputes and unresolved issues.
Credit	Credit decision or applicable explanation	Use the institution's credit process and review. A marketing score is not a reason for a credit action.

Table 2: Three financial decision types and their boundaries

Design proposal. The legal source for the credit notice is Regulation B; the 2023 circular is historical (CFPB, 2023a; Regulation B, 2025).

Studies of machine learning in lending show why prediction gains alone do not establish fair distributional outcomes (Fuster et al., 2022; Bartlett et al., 2022). They do not estimate campaign effects. They motivate checking who learns about products and who applies. The marketing audience gate can shape access even when the underwriting model is separate.

The CFPB withdrew Circular 2023-03 in May 2025 (CFPB, 2023a; CFPB, 2025). Separately, Regulation B requires the statement of reasons to identify the specific principal reasons for an adverse credit action (Regulation B, 2025). An explanation of a marketing score cannot substitute for a credit notice. Keep their decision owners, eligibility versions, and access logs distinct.

At the boundary, uncertainty should trigger referral rather than automation. A borderline case can receive product-neutral information and a route to a qualified employee without inferring an approval decision.

PROPOSED DECISION FRAMEWORK

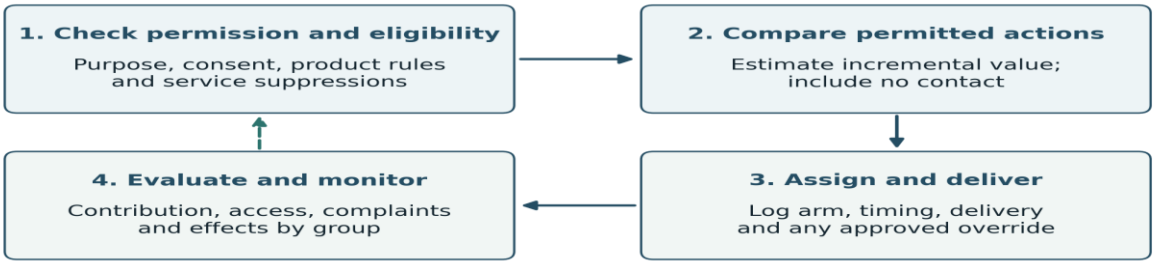


Fig 2: Proposed workflow for governed financial marketing

Corresponding email: mhasan1@my.westga.edu

The gate blocks ineligible or unauthorized contact. Assignment and delivery are recorded separately; later outcomes inform monitoring and approved policy changes.

Figure 2 shows how the framework processes one decision opportunity at a time. It first checks purpose, contact permission, product rules, and any active service suppression. It then evaluates the remaining actions, including no contact, against customer and business objectives. Assignment and delivery are recorded separately so that later analysis can distinguish the effect of the policy from delivery failures.

Each decision has a record linking the customer unit, observation time, available actions, estimated value, assignment probability, actual action, and applicable model and rule versions. Outcomes are recorded against prespecified windows. This arrangement allows an analyst to reproduce the decision without reconstructing it from a dashboard screenshot or treating a later customer interaction as if it had been known earlier.

The design incorporates two routes for change. A routine model refresh may update scores within an approved purpose and action set. A proposed new data source, purpose, or material product rule returns to governance review. The distinction follows the model-risk principle that implementation and use determine consequences, not the algorithm alone (Federal Reserve & OCC, 2011; Tabassi, 2023). A human owner can suspend an action, preserve the log, and revert to the approved baseline policy while the issue is examined.

A. CUSTOMER AND HOUSEHOLD UNITS

The unit of a model, a treatment, and an outcome need not be identical. One person can have several deposit accounts, a joint credit product, and multiple devices. Treating each account as an independent marketing opportunity can duplicate contact and inflate apparent sample size. At the same time, merging people into a household without reliable evidence can expose one person's activity to another.

The proposed system creates a documented decision unit for each campaign. A sole-owner onboarding program can use the customer relationship as its unit. A joint-account program may require permissions and communications defined at the account level. A referral experiment should consider whether messages or incentives spill over to other members of the same household. Randomization and uncertainty calculations should use the unit at which interference is likely, with an explicit statement of any remaining dependence (Kohavi et al., 2009).

Use verified relationships to link records, and log the evidence for each link. An inferred address or shared device does not authorize a disclosure. In a separate business banking application, distinguish an owner, a firm, and an account signer. The SBA's February 2026 FAQ reports about 36.2 million U.S. small businesses, 82.3% of them without paid employees; that business count cannot establish who can receive account information (SBA Office of Advocacy, 2026).

Population reporting should distinguish all customers, eligible decision units, reachable units, assigned units, and delivered messages. Each is useful for a different question. Incremental outcomes belong to the assigned eligible population; deliverability belongs to the operational queue. Reporting only delivered contacts can remove precisely the failures an institution needs to understand and can distort comparison between treatment and control.

B. DATA CONTRACTS AND OBSERVATION TIME

A trustworthy training record contains only information available when a decision could have been made. Payment settlement, call disposition, a later application decision, and a subsequent product closure can all enter a warehouse after the decision date. If joins use a current customer snapshot to reconstruct an old decision, future information can leak into training even when all event dates appear plausible.

The design retains event time and availability time. Event time identifies when the underlying activity occurred; availability time records when the system could first have used it. The feature builder requires both to precede the decision snapshot. Product versions and permissions are joined using their effective dates and the state known at the time. Data corrections are versioned rather than silently applied to an old training run.

The unit-level table should distinguish missing, absent, and suppressed observations. An app event that never occurred is different from an event stream that failed, and both differ from information deliberately excluded for privacy. A missingness flag can reveal process quality, but it can also act as a proxy for uneven access. Its inclusion should have a documented purpose and a plan for review.

Corresponding email: mhasan1@my.westga.edu

The bank telemarketing work of Moro et al. (2014) illustrates how customer and campaign records can be used to predict response. Its institutional setting is Portuguese, so it does not establish performance in the United States. A bank adapting such a design should inspect whether any candidate field, such as call duration, exists only after the action. A feature that is useful for describing a completed call is unavailable for deciding whom to call.

C. BASELINES BEFORE COMPLEX MODELS

A model comparison should begin with an executable simple policy. For onboarding, that could be a fixed sequence of permitted reminders based on the customer's selected channel. For retention, it might be a recent unresolved service request with a clear response owner. The baseline must operate under the same budget, contact frequency, product constraints, and outcome definition as the candidate.

An initial predictive model can be regularized logistic regression using a short list of justified features. A tree ensemble can be tested if interactions appear plausible and data volume is sufficient. Recency, frequency, and monetary history provide a disciplined starting point for customer value analysis, but the variables require interpretation in context (Fader et al., 2005). For example, lower frequency may indicate satisfaction with a savings product, not disengagement. Historical balance or payment volume is not itself a measure of net incremental value.

Performance comparisons should use chronological development, calibration, and evaluation periods. Repeated customers should be handled consistently with the intended deployment population. An evaluation on later observations of the same customers addresses time transfer; a holdout of whole relationships tests transfer to new customers. Neither split alone answers both questions. The choice should be stated before measuring model performance.

The institution should retain a baseline it can run without a model provider. A complex model earns deployment through an observable improvement in policy value or customer outcomes under controlled evaluation. If a simpler model achieves comparable results and is easier to inspect, its lower governance and maintenance cost is part of the comparison.

D. INCREMENTAL RESPONSE AND POLICY VALUE

Response prediction estimates the chance of an outcome under the historical action process. The decision problem instead asks how the outcome would change if a particular action were taken. Let Y(1) denote an outcome under the proposed contact and Y(0) the outcome under no contact, both within a fixed window. The conditional treatment effect is the expected difference between these potential outcomes given information available before assignment. Only one outcome is observed for a decision unit.

Random assignment identifies an average effect for the eligible study population when assignment is implemented as planned and interference is addressed. Heterogeneous-effect estimators can explore where effects differ, but they do not solve confounding merely by using a flexible algorithm (Künzel et al., 2019). Strong observational fit is not evidence that the model has estimated the causal response. Advertising measurement studies demonstrate why nonrandom targeting can give different answers from controlled comparisons (Gordon et al., 2019).

Define V(a, x) as expected contribution under action a at customer state x, net of variable cost and material expected harm. The relevant comparison is V(a, x) minus V(no contact, x). When several actions are eligible, the system chooses among them subject to capacity, fairness, and customer safeguards. The no-contact option provides a reference and a legitimate decision, particularly when uncertainty is large or expected benefit is small.

The proposed order is estimate the overall effect first, inspect preplanned subgroup contrasts second, and consider an individualized policy only when held-out value supports it. This sequence limits the chance that an unstable pattern discovered during exploration becomes an automated decision rule.

E. EXPERIMENTAL DESIGN AND PRECISION

A defensible experiment specifies the population, intervention, comparison, assignment unit, outcome window, and analysis before contacts begin. A bank testing an onboarding reminder could assign newly eligible relationships within enrollment-week and preferred-channel blocks. Every assigned relationship remains in the analysis, including contacts that fail to deliver. This intention-to-treat approach evaluates the effectiveness of the policy the institution can actually implement (Kohavi et al., 2009).

Sample size depends on the baseline outcome, the minimum effect worth acting on, error probabilities, and any dependence. As an illustration, comparing an 8% baseline rate with a 9% rate at two-sided 5% significance and 80% power under independent equal groups requires about 12,200 units per arm using the normal approximation for two proportions. This is planning arithmetic, not a study

Corresponding email: mhasan1@my.westga.edu

result or an institution-specific requirement. Clustered assignment, missing outcomes, sequential monitoring, or multiple comparisons can change the needed sample.

A review plan should also identify possible interference. If one household member receives a message and another opens an account, customer-level assignments may contaminate the control group. Branch staff may change behavior after seeing a campaign, affecting both arms. The design can cluster, isolate the intervention, or redefine the estimand to include spillovers. Where isolation is impossible, the uncertainty and limitations should be reported.

Every trial should have a harm rule independent of conversion. A material increase in complaints, inaccurate disclosures, unauthorized contact, or service delays can justify suspension even before the primary outcome matures. The record should distinguish a planned safety review from an unplanned decision made after looking at favorable sales results.

F. ACTION CHOICE AND OPPORTUNITY COST

Personalization often defaults to choosing the highest-scoring product. That is insufficient when products have different costs, delivery constraints, or customer implications. An action with the highest predicted response can have low incremental contribution because many customers would act without it. A lower-response intervention can create more value if it resolves a specific friction for customers who would otherwise abandon an appropriate process.

The decision score should therefore use the expected difference in contribution between a candidate action and no contact. Contribution may include realized margin over a specified horizon, minus acquisition incentives, delivery costs, and expected incremental service costs. When the outcome involves account usage, count balances or activity only under a definition compatible with business and customer objectives. A transfer between accounts within the institution may not create the same value as a new relationship.

Capacity creates an additional opportunity cost. A relationship manager's time used for a marginal prospect is unavailable for a customer with an unresolved issue. The candidate policy can reserve capacity for required service and rank only the remaining eligible marketing opportunities. When two channels are viable, compare their incremental contributions at the same horizon and include expected handling time.

An institution can start by using uncertainty as an eligibility condition for automation. A policy whose estimate changes sign under plausible cost assumptions should remain a limited experiment. If projected value is near zero, no contact is often the more informative baseline because it avoids spending resources on an unproven action and preserves a clear counterfactual (Gordon et al., 2019).

ACQUISITION AND ONBOARDING

Customer acquisition begins before an account is funded. An ad impression, a visit to an information page, an application start, an approval, and an active relationship are separate events. Optimizing a model for the earliest event is easier because there are more observations, but it can steer the organization toward people who are easy to engage and unlikely to benefit from the product. A financial institution should choose a success measure that reflects both a usable account and an informed decision.

A suitable first experiment tests a permitted educational prompt, not a change in the terms offered. The prompt might explain how to complete a step the customer has already chosen to begin. The comparison group receives the normal service experience. The outcome could be a completed, funded account with qualifying activity within a fixed period, net of reversals and duplicate applications. In parallel, the bank measures complaints, requests for assistance, and opt-outs.

The design should preserve access for customers who do not receive or engage with the message. A branch visit, telephone call, or accessible digital route remains available. Channel preferences and language needs should be recorded as customer choices rather than inferred from protected or sensitive characteristics. The option to get help is valuable independently of campaign conversion.

Interpret the results at each stage of the funnel, but do not make stage selection dependent on which result is favorable. An increase in starts without completed accounts may signal friction, not campaign success. A useful implementation connects marketing logs, product records, and customer service outcomes while preserving the separation between promotional and eligibility decisions (Wedel & Kannan, 2016; Lemon & Verhoef, 2016).

A. RETENTION AND SERVICE RECOVERY

The highest predicted churn risk does not necessarily identify customers whose behavior can be changed by a retention offer. Some high-risk customers have already decided to leave; others are responding to a problem that requires service recovery rather than an incentive. Ascarza's (2018) work on retention targeting gives an empirical reason to distinguish risk from responsiveness to treatment.

Corresponding email: mhasan1@my.westga.edu

A financial retention design should define the event carefully. Formal closure is observable, while inactivity can have several meanings. A savings customer may use an account infrequently by design. A temporary balance decline can reflect a planned purchase rather than dissatisfaction. The proposed approach defines a relationship-specific outcome and a maturity period, then uses pre-action events to identify whether there is a service issue, a product mismatch, or a plausible opportunity to provide useful information.

Service recovery should remain available to everyone who needs it. A test may compare permitted ways of delivering optional help, but it should not withhold required service or relief. A no-offer control can receive the standard service process. Discount or fee-waiver experiments need a complete account of cost and an assessment of whether similar customers are treated consistently.

Evaluate retention beyond the first response. A temporary return to activity can be followed by delayed closure, a complaint, or an incentive cost that exceeds the value of retained activity. Use a mature outcome window and report a longer-term check when feasible. Any model trained on historical offers should account for the possibility that the existing contact policy created part of the pattern it is trying to predict.

B. CONVERSATIONAL ASSISTANCE AND HANDOFFS

An AI assistant can help a customer find general product information or identify the next step in a routine process. The same interface becomes risky when it fabricates a term, cannot recognize a dispute, or prevents the customer from reaching a person. The CFPB's 2023 issue spotlight describes limitations of chatbots in consumer finance and the consequences of inadequate escalation (CFPB, 2023b).

The framework restricts automated content to a versioned source of approved material. A response that depends on a customer's specific account state must use authenticated, authorized retrieval. The assistant should avoid guessing at fees, eligibility, or the status of a complaint. Where information is missing, the appropriate output is a limited answer and a clear route to a qualified employee.

Service and marketing evaluation must be separated. Useful service indicators include correct resolution, transfer completion, repeated contact for the same problem, and time to reach a human. A sales conversion attributed to a chatbot session says little about whether the service problem was resolved. If an assistant presents optional marketing content, the permission and eligibility gates apply before the content is shown.

Deployment should begin with bounded intents and representative test cases, including ambiguous requests and users who need an accessible alternative. Human review should record examples of unsupported claims, missed escalation, and misleading explanations. If the assistant cannot preserve the required record or give the customer a reliable exit, restricting the task is preferable to increasing automation. NIST's AI risk framework supports an ongoing measurement and response process rather than a one-time approval (Tabassi, 2023).

METRICS AND MODEL ASSESSMENT

Table 3 distinguishes ranking, probability quality, and policy value. A model can rank responses well and still be poorly calibrated or prioritize customers who would respond without contact.

Question	Primary checks	Population or denominator
Does the score rank?	Precision at contact capacity; average precision when appropriate	Later, eligible decision opportunities
Are probabilities reliable?	Brier score; observed vs. predicted rates	Matured outcomes by horizon and group
Does contact help?	Assigned-arm difference and interval	All eligible units assigned to each arm
Does value cover cost?	Incremental contribution after cost	Assigned units, same outcome horizon
Are safeguards working?	Suppression failures; opt-outs; complaints	Eligible, sent, or assigned units, as defined

Table 3: Evaluate prediction, policy, and customer outcomes separately

Proposed measures. Disclose counts, uncertainty, and any unresolved outcome matching.

Calibration is assessed with predicted probabilities and observed frequencies over a declared horizon. A proper scoring rule, such as mean squared probability error, penalizes inaccurate probability estimates (Gneiting & Raftery, 2007). Calibration displays should

Corresponding email: mhasan1@my.westga.edu

include group counts and uncertainty where feasible. An apparently flat pattern in a small segment may reflect a lack of observations rather than a dependable fit.

Policy value belongs to the experimental comparison, using outcomes for all assigned eligible units. A decision dashboard should pair incremental contribution with contact burden, complaint rates, and reach by channel. A single composite index can hide a worsening customer outcome behind a favorable revenue estimate. Reporting the components allows an owner to decide whether the trade-off meets the institution's standards.

Reconcile duplicate accounts, cancellations, and reversals before comparison. Preserve the decision snapshot, rule version, and outcome maturity date so the analysis can be rerun.

A. FAIR ACCESS AND SUBGROUP UNCERTAINTY

Fairness evaluation starts with who can become eligible and who is reached. A policy that performs similarly among delivered messages can still systematically exclude customers whose preferred channel is not supported. Measure the progression from eligible to assigned to delivered and then to the final outcome. Interpret group differences in light of customer choice, data coverage, product rules, and any barriers created by the institution.

Predictive fairness criteria can conflict when groups have different base rates and the model is imperfect. Chouldechova (2017) demonstrates a mathematical version of this issue in another domain. It does not prescribe a banking marketing policy. Its practical implication is that an institution should state which risks matter, measure them directly, and avoid claiming that one favorable statistic establishes overall fair treatment.

Use group analysis only with an appropriate basis and controlled access. Protected attributes used for authorized auditing should be separated from the campaign scoring pipeline. A small subgroup estimate needs uncertainty, a record of its denominator, and care about disclosure. Where group data are missing or unreliable, the appropriate conclusion is limited visibility rather than presumed equality. Channel accessibility can still be examined with customer-selected preferences and process tests.

Remediation can involve changing the action set, adding a service route, or simplifying eligibility before changing a model. Credit-market studies show why the distribution of benefits can matter even when aggregate predictive performance improves (Fuster et al., 2022; Bartlett et al., 2022). For marketing, the central question is whether the policy expands useful access without making sensitive inferences, concentrating burdens, or quietly restricting information to favored groups.

B. FEATURE EXPLANATIONS AND DECISION REASONS

Explainability should answer a concrete question for a specific audience. A model developer needs to detect leakage, instability, and reliance on questionable proxies. A reviewer needs to understand why an action was available and whether the policy operated as approved. A customer may need a clear account of why a message was sent and how to change preferences. These needs are related but cannot all be met by a feature-importance plot.

An attribution method can describe how a fitted score differs from a reference. The Shapley-based framework of Lundberg and Lee (2017) is one method for allocating prediction contributions. Its outputs depend on the model and the reference distribution; correlated variables and omitted information complicate interpretation. An explanation that a score was influenced by app activity does not mean that creating more app activity will cause better retention.

The decision record should include both model and policy reasons. If a customer did not receive an offer because of a contact limit, the relevant reason is the rule, even if the model score was high. If two offers were eligible and one was chosen because of its experimentally estimated value, record that comparison separately from the predicted response. This prevents an attribution from being mistaken for the actual decision logic.

Customer-facing explanations should be accurate, limited to relevant information, and paired with a way to correct records or change preferences. They should not reveal sensitive inferences about other people in a shared relationship. The institution should test whether reviewers can use explanations to detect errors, not simply whether a tool can produce them.

C. PRIVACY ARCHITECTURE AND VENDOR RISK

The most effective protection for an unnecessary field is to avoid collecting or repurposing it. The customer intelligence platform should maintain a feature register that identifies purpose, data source, retention period, access roles, and permitted recipients. Training data, production features, and marketing exports should each have separate access controls. A vendor's ability to process a field does not make the transfer appropriate for the institution's intended use.

Corresponding email: mhasan1@my.westga.edu

Pseudonymous identifiers help separate direct contact details from analytical records, but linkage can still reveal a person. Encryption, role-based access, deletion procedures, and logged exports address different points of exposure. An institution should test whether a customer's withdrawal or correction reaches the feature store, campaign queue, and relevant service provider. A control described only in a policy document does not establish that a queued contact will stop.

For external model services, specify what data leave the institution, how the provider handles prompts and outputs, and whether any information can be retained for provider training. If a service can be performed with aggregate counts or approved public content, do not send account-level context. A minimally scoped interface also simplifies testing when the vendor changes its model or configuration.

Governance should address availability as well as confidentiality. A service failure must not result in a default release of a suppressed campaign. Fail closed for permission checks and route ambiguous cases to human review. The AI RMF's emphasis on mapping risks, measuring controls, and managing changes supports this lifecycle approach (Tabassi, 2023).

D. HISTORICAL U.S. GOVERNANCE CONTEXT

The 2011 Federal Reserve and OCC model risk guidance describes development, validation, and governance practices (Federal Reserve & OCC, 2011). The CFPB's October 2024 final rule addressed consumer-authorized data access (CFPB, 2024). In October 2025, the district court in Forcht Bank enjoined the CFPB from enforcing that rule until it completes reconsideration. At the February 2026 cutoff, the original rule should not be read as an enforceable implementation timetable (Forcht Bank v. CFPB, 2025).

A practical design separates authorization to obtain data from permission to use it for a specific marketing purpose. A customer's request to share account information for a budgeting service does not by itself define the uses available for an unrelated sales campaign. Record purpose and source as decision inputs, and require separate review when a use changes. This is an institutional design principle rather than a legal determination.

On February 19, 2026, Treasury released a financial-services AI risk framework and lexicon (U.S. Department of the Treasury, 2026). For this proposed marketing workflow, an implementation would record the use case and owner, review data and vendor dependencies, validate permissions and customer outcomes, and test the fallback before approval. Changes to a model, prompt, or action rule should trigger a review proportionate to the affected customers and potential harm. These are proposed implementation steps, informed by the sector framework.

An implementation team should obtain institution-specific review before operational use, especially where marketing intersects with credit, pricing, consumer data sharing, or automated service decisions. The article's contribution is an auditable structure for those decisions. It makes no claim of regulatory approval or compliance certification.

WORKED CAMPAIGN ECONOMICS

Consider an illustrative trial with 5,000 eligible customers assigned to a message and 5,000 to no contact. Assume 450 treated customers and 400 control customers reach the same verified outcome within the prespecified window. The observed arm rates in this constructed example are 9% and 8%. The difference is one percentage point, corresponding to 50 incremental outcomes in the treated group if the comparison is valid.

Assume that each incremental outcome yields $90 in contribution before campaign costs. The estimated incremental contribution is $4,500. At $0.30 per assigned message, delivery costs $1,500. If all 450 treated converters receive a $10 incentive, the cost is $4,500, including incentives paid to customers who would have converted without the message. The estimated net value is therefore negative $1,500. The calculation shows why counting every conversion as incremental can reverse the decision (Gordon et al., 2019).

The same assumptions yield a break-even requirement of about 67 incremental outcomes because $6,000 in observed campaign costs divided by $90 is 66.7. At this treatment-arm size, that corresponds to about 1.33 percentage points of lift, holding the observed incentive payout fixed. A forward plan should instead allow incentive payout to vary with conversion. If baseline conversion remains 8% and every treated converter receives $10, the variable-cost break-even equation is $90d - 0.30 - 10(0.08 + d) = 0$, giving $d = 0.01375$.

These are constructed values to illustrate an accounting method. They are neither survey results nor expected returns for a U.S. financial institution.

Corresponding email: mhasan1@my.westga.edu

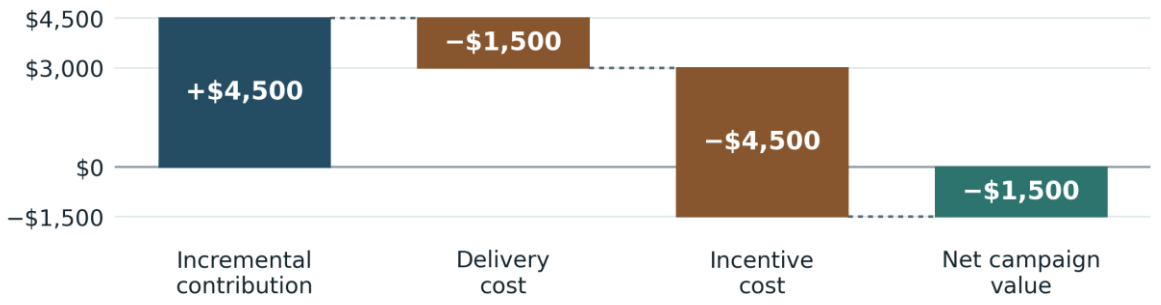


Fig 3: Contribution bridge for the hypothetical campaign

A. DECISION UNCERTAINTY AND SENSITIVITY

In the illustrative trial, a one-point effect estimate is compatible with a wide range of true effects. Under the simplifying assumption of independent binomial groups, the standard error of the observed difference is approximately 0.56 percentage points. A normal 95% interval is roughly −0.09 to 2.09 percentage points. This approximation is used only to demonstrate uncertainty; clustered customers or nonstandard assignment require different calculations (Kohavi et al., 2009).

The example supports a decision distinction. The point estimate is positive for conversion and negative for net contribution. The confidence interval spans no effect and includes effects large enough to break even. A decision maker cannot responsibly describe the trial as proof of profitability or proof that no useful effect exists. An appropriate next step would depend on the expected value of more information, the risk of further contact, and the opportunity cost of another trial.

Sensitivity analysis should vary assumptions that can change the recommendation. The contribution per new account, incentive payout, cancellation rate, and service cost may matter more than a minor change in model discrimination. Report a break-even condition so that finance can see which assumption would need to be true. If margins differ by product or customer behavior, use realized contribution at the unit level rather than multiplying an average value by an attributed count.

The policy can also include a conservative decision rule: expand only when estimated value clears a prespecified hurdle and customer safeguards remain within limits. That rule is a managerial choice, not a statistical law. The threshold should be recorded before results are examined and revisited when costs, channel policies, or customer needs change.

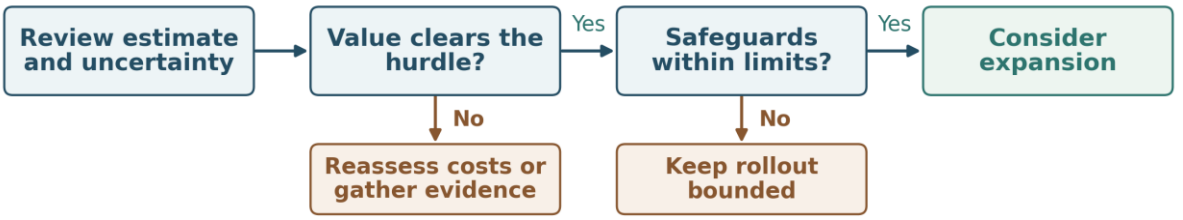


Fig 4: Decision gates for expanding a campaign

B. IMPLEMENTATION AND ACCOUNTABLE OWNERSHIP

An initial pilot should have one accountable business owner, an independent reviewer, a data steward, and an operational contact for customer issues. Roles can be combined in a small institution, but approval and execution responsibilities should be distinguishable (Tabassi, 2023). A model change must have a version, purpose, and release record. A small institution can use a simple permission and assignment ledger; a larger one may need a shared cross-product contact budget. The reviewer should be able to trace a decision from the customer unit and eligibility state to the approved action and observed outcome.

Begin with a narrow treatment such as an approved onboarding reminder. Use a documented baseline and a randomized holdout, then confirm that assignments, suppressions, and deliveries reconcile. A shadow run can reveal whether the policy would violate contact limits, use stale records, or produce an unmanageable queue. Because shadow outputs do not change customer exposure, they cannot estimate the treatment effect.

Expansion should be staged by risk and readiness. Extending a well-tested message to another channel changes timing, presentation, and customer expectations, so it may require a new controlled comparison. Expanding to a different product can introduce new eligibility or suitability concerns. The owner should distinguish changes that can be monitored from changes that need prior review.

Corresponding email: mhasan1@my.westga.edu

The fallback should be explicit. If data quality drops or a vendor becomes unavailable, stop the affected automated contact, preserve the audit log, and continue standard service. An effective rollback restores a known safe policy rather than merely reverting software. The system should support a later reconciliation of messages already queued or sent and a process for handling any resulting customer complaint.



Fig 5: Pilot roles and the traceable decision record

C. MONITORING CHANGES OVER TIME

Monitoring needs to detect several different changes. Feature distributions can shift because of seasonality, data collection, or channel redesign. Outcome rates can change because of economic conditions, eligibility, or customer needs. The effect of a message can change even when both features and outcomes appear stable. A monitoring dashboard should not treat a shift in one of these quantities as evidence that the others changed.

Assign a cadence to each measure. Permission and delivery failures can be monitored promptly. Contribution and retention require an outcome window long enough for reversals and closures to mature. Keep a small randomized comparison where permissible and useful so that the organization can distinguish declining campaign incrementality from a change in its overall customer base (Gordon et al., 2019). Record any intervention made during a monitoring period.

Use segmentation carefully. A new feature can improve aggregate predictions while degrading performance for customers with less complete digital histories. Review channel reach, calibration, opt-outs, and complaints by meaningful and adequately supported groups. For very small groups, avoid a false sense of precision. Combine quantitative review with targeted process testing and customer feedback.

When deterioration appears, diagnose the source before retraining. Retraining on mislabeled outcomes can amplify a data issue. Adding more messages can worsen a fatigue problem. A rule change may alter who was eligible and invalidate the comparison with earlier cohorts. The review record should identify the change, possible causes, corrective action, and the evidence needed before the policy is restarted or expanded.

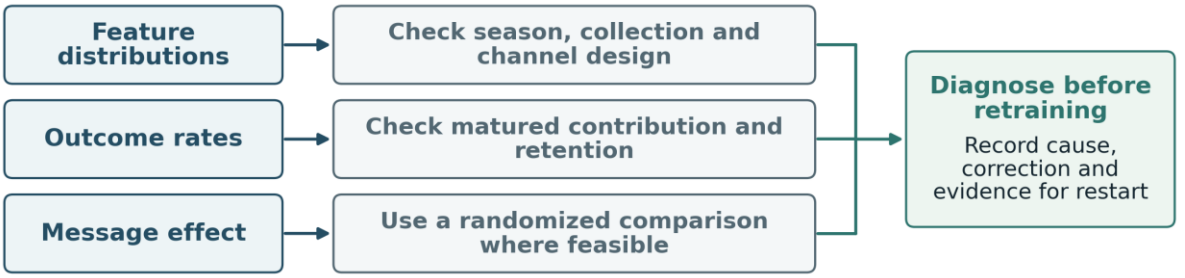


Fig 6: Separate monitoring signals before choosing a response

DISCUSSION

The FDIC and ABA evidence supports keeping several ways to reach customers. It does not show whether a policy changes their behavior (FDIC, 2024; American Bankers Association, 2025). The December 2025 Census release put e-commerce at 16.4% of sales among covered U.S. retail employers in the third quarter of 2025. That denominator describes retail sales, so it cannot benchmark financial contact or account access (U.S. Census Bureau, 2025).

Marketing studies help formulate hypotheses about timing, context, and trust (Bleier & Eisenbeiss, 2015; Aguirre et al., 2015; Tucker, 2014). Adoption alone also leaves an impact question. In the February 2026 working paper, Yotzov et al. (2026) report 69% AI use in their four-country firm survey and 89% reporting no effect on labor productivity over the prior three years. These self-reports do not measure U.S. bank marketing returns.

The original design contribution is the separation of permission, eligibility, effect estimation, and action execution. This separation changes what the organization can audit. It can identify customers excluded by a rule, distinguish a predicted response from

Corresponding email: mhasan1@my.westga.edu

expected lift, and evaluate whether the action delivered matched the approved policy. It also makes a no-contact decision a measurable part of the system rather than a missing record.

One implication is that the best initial AI project may use less data than technically available. A narrow decision with a clear outcome can produce better evidence than an enterprise-wide targeting engine whose interventions and denominators are not recorded. More elaborate models become useful when the underlying decision and measurement system are strong enough to test their added contribution.

A. LIMITS OF THE EVIDENCE

This article's quantitative comparison uses rounded, published FDIC estimates (FDIC, 2024). It does not analyze household microdata, adjust for differences in household composition, or test whether observed changes are statistically distinguishable. The results should not be interpreted as a change in the same households, because the survey-year distributions do not constitute a matched panel for this analysis.

The review is targeted and narrative. It can connect relevant findings to a proposed design, but it does not estimate an average effect of personalization across studies. Much of the experimental literature concerns advertising or other products outside U.S. banking. Differences in trust, product complexity, oversight, channel reach, and customer stakes limit transfer. Credit-market studies provide useful context for fairness and access but cannot estimate the return to marketing.

The proposed system also assumes an institution can identify eligibility and capture outcomes with sufficient quality. Those conditions may not hold when customer records are fragmented, service is outsourced, or consent states are inconsistent. The model-risk controls are design requirements to test, not assurances that an institution has met them. Worked numerical examples are hypothetical and should not be treated as evidence of profitability or harm.

The governance discussion reflects the February 28, 2026 evidence cutoff. It distinguishes the historical model risk guidance, the litigation status of the personal financial data rights rule, and the new sector AI risk framework (Federal Reserve & OCC, 2011; U.S. Department of the Treasury, 2026; Forcht Bank v. CFPB, 2025). Operational decisions depend on the institution, product, data use, and rules that apply when the decision is made.

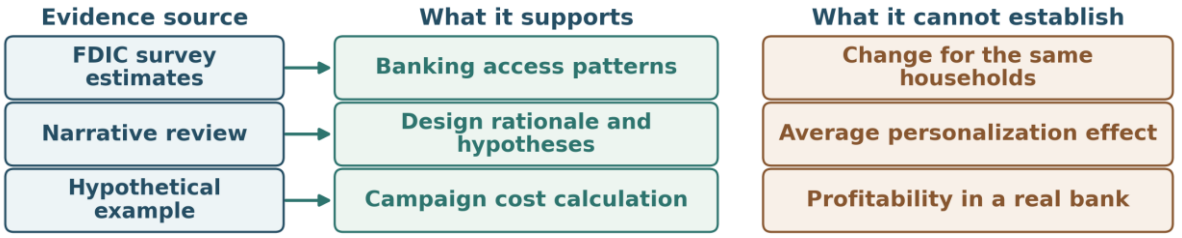


Fig 7: Evidence types and the limits of interpretation

B. RESEARCH AGENDA FOR VALIDATION

A first empirical study could compare a permitted reminder, a neutral information guide, and the standard experience for newly eligible customers. Define the primary endpoint as an account activation with meaningful use at a fixed horizon. Include customer complaints, opt-outs, assistance requests, and contribution as secondary measures. A multi-institution replication would help distinguish effects of the treatment from effects of a particular delivery platform or customer mix.

A second study could compare risk-only retention targeting with a policy that estimates incremental response (see Ascarza, 2018). Both would use the same contact capacity and approved treatment menu. Evaluation would use a randomized comparison and an untouched test period so that exploratory subgroup choices do not determine the result. Report uncertainty around policy value and any differences in service burden.

A third study could test whether a transparent permission and explanation interface improves customer understanding. Measure whether people can identify why a message appeared, change a preference, and access service after declining contact. A usability experiment should include people who rely on in-person, phone, or accessible channels. It should assess comprehension and control rather than rely on satisfaction alone.

Each study needs a data retention plan and a realistic approach to repeated customers and delayed outcomes. Negative and null results should be retained. A framework that consistently records when an intervention fails can help an institution avoid repeating

Corresponding email: mhasan1@my.westga.edu

costly campaigns and can generate more useful evidence than selectively reporting successful examples. No such studies were conducted for the present article.

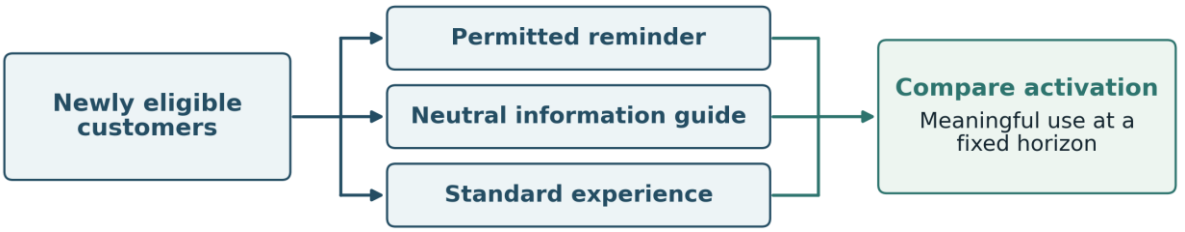


Fig 8: Proposed comparison of onboarding options

CONCLUSION

The available evidence supports a channel strategy that preserves alternatives for customers who do not primarily bank through an app. The FDIC and ABA measures show why mobile access matters; their differing populations preclude a direct time comparison. Neither survey measures gains from AI personalization (FDIC, 2024; American Bankers Association, 2025).

The proposed approach puts permission and eligibility ahead of ranking, treats no contact as a valid decision, and evaluates incremental value with an appropriate comparison. It also separates marketing from service and credit decisions and requires an audit trail that explains the action actually taken. These features reduce the risk of confusing engagement with customer benefit or attributing existing demand to a campaign.

The most useful first step is a bounded, measurable experiment with an approved action and a clearly defined customer outcome. A simple policy that can be tested and monitored is more informative than a complex policy without a credible comparison. Scaling should follow evidence of contribution and acceptable customer safeguards rather than the strength of a predictive score alone.

The article establishes a method and an evidence boundary. Its public statistics and research citations support the design rationale. The operational recommendations, economic example, and proposed studies remain to be validated in the institutions that would use them.

APPENDIX A REPRODUCTION AND REVIEW RECORD

Public data extraction. Use Table 2.1 of the 2023 FDIC household survey report. Transcribe the 2019, 2021, and 2023 rows for the specified banked-household population. Subtract the 2019 mobile estimate from the 2023 estimate: 48.3 − 34.0 = 14.3 percentage points. The comparable teller change is 15.1 − 21.0 = −5.9 points. The in-person primary share in 2023 is 13.8 + 15.1 = 28.9%. The unbanked estimate comes from the national findings and is not combined with the channel shares (FDIC, 2024).

Evidence search topics. Banking access, digital personalization, causal measurement, model risk, adverse-action notices, and consumer data rights. Freeze sources at February 28, 2026. Use the October 2025 ABA survey and the Federal Reserve's May 2025 report as separate series. For the two February research papers, use the February working paper and the February 10, version 3 preprint (American Bankers Association, 2025; Federal Reserve Board, 2025; Yotzov et al., 2026; Fang et al., 2026).

Study record to retain. Specify the decision unit and eligibility date, treatment and baseline, assignment probability, outcome horizon, contact and incentive costs, human review path, and safety measures. Save the versioned feature schema and approved rules. Preserve an analysis plan and all allocation and outcome counts needed to reproduce a difference in assigned-arm rates.

Interpretation. Public statistics are descriptive. The arithmetic example is hypothetical. The workflow is proposed. No private financial data, live interventions, or new institution-level effect estimates are reported.

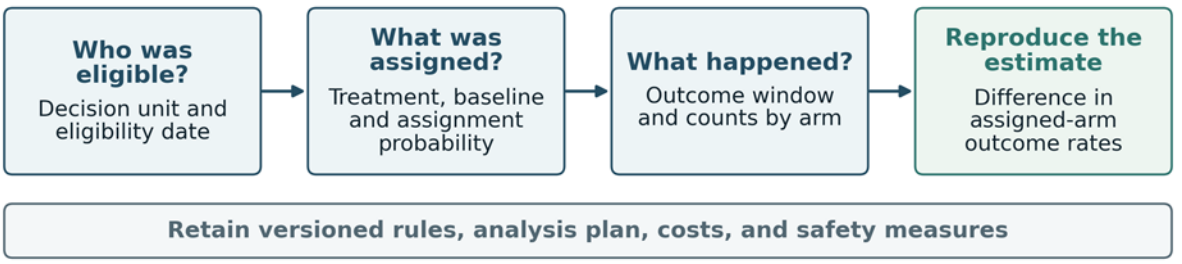


Fig 9: Records needed to reproduce an assigned-arm comparison

Corresponding email: mhasan1@my.westga.edu

REFERENCE

1. Federal Deposit Insurance Corporation (2024). 2023 FDIC National Survey of Unbanked and Underbanked Households. Table 2.1 and national findings. FDIC report and tables
2. Wedel, M., & Kannan, P. K. (2016). Marketing analytics for data-rich environments. Journal of Marketing, 80(6), 97–121. https://doi.org/10.1509/jm.15.0413
3. Lemon, K. N., & Verhoef, P. C. (2016). Understanding customer experience throughout the customer journey. Journal of Marketing, 80(6), 69–96. https://doi.org/10.1509/jm.15.0420
4. Lambrecht, A., & Tucker, C. (2013). When does retargeting work? Information specificity in online advertising. Journal of Marketing Research, 50(5), 561–576. https://doi.org/10.1509/jmr.11.0503
5. Bleier, A., & Eisenbeiss, M. (2015). Personalized online advertising effectiveness: The interplay of what, when, and where. Marketing Science, 34(5), 669–688. https://doi.org/10.1287/mksc.2015.0930
6. Huang, M.-H., & Rust, R. T. (2021). A strategic framework for artificial intelligence in marketing. Journal of the Academy of Marketing Science, 49(1), 30–50. https://doi.org/10.1007/s11747-020-00749-9
7. Davenport, T., Guha, A., Grewal, D., & Bressgott, T. (2020). How artificial intelligence will change the future of marketing. Journal of the Academy of Marketing Science, 48(1), 24–42. https://doi.org/10.1007/s11747-019-00696-0
8. Aguirre, E., Mahr, D., Grewal, D., de Ruyter, K., & Wetzels, M. (2015). Unraveling the personalization paradox: The effect of information collection and trust-building strategies on online advertisement effectiveness. Journal of Retailing, 91(1), 34–49. https://doi.org/10.1016/j.jretai.2014.09.005
9. Tucker, C. E. (2014). Social networks, personalized advertising, and privacy controls. Journal of Marketing Research, 51(5), 546–562. https://doi.org/10.1509/jmr.10.0355
10. Consumer Financial Protection Bureau (2023a). Adverse action notification requirements and the proper use of the CFPB's sample forms provided in Regulation B. Consumer Financial Protection Circular 2023-03; withdrawn May 12, 2025; historical source. Official publication
11. Fuster, A., Goldsmith-Pinkham, P., Ramadorai, T., & Walther, A. (2022). Predictably unequal? The effects of machine learning on credit markets. The Journal of Finance, 77(1), 5–47. https://doi.org/10.1111/jofi.13090
12. Bartlett, R., Morse, A., Stanton, R., & Wallace, N. (2022). Consumer-lending discrimination in the FinTech era. Journal of Financial Economics, 143(1), 30–56. https://doi.org/10.1016/j.jfineco.2021.05.047
13. Board of Governors of the Federal Reserve System, & Office of the Comptroller of the Currency (2011). Supervisory guidance on model risk management. Attachment to SR 11-7. Official publication
14. Tabassi, E. (2023). Artificial Intelligence Risk Management Framework (AI RMF 1.0). NIST AI 100-1. National Institute of Standards and Technology. https://doi.org/10.6028/NIST.AI.100-1
15. Kohavi, R., Longbotham, R., Sommerfield, D., & Henne, R. M. (2009). Controlled experiments on the web: Survey and practical guide. Data Mining and Knowledge Discovery, 18(1), 140–181. https://doi.org/10.1007/s10618-008-0114-1
16. Moro, S., Cortez, P., & Rita, P. (2014). A data-driven approach to predict the success of bank telemarketing. Decision Support Systems, 62, 22–31. https://doi.org/10.1016/j.dss.2014.03.001
17. Fader, P. S., Hardie, B. G. S., & Lee, K. L. (2005). RFM and CLV: Using iso-value curves for customer base analysis. Journal of Marketing Research, 42(4), 415–430. https://doi.org/10.1509/jmkr.2005.42.4.415
18. Künzel, S. R., Sekhon, J. S., Bickel, P. J., & Yu, B. (2019). Metalearners for estimating heterogeneous treatment effects using machine learning. Proceedings of the National Academy of Sciences, 116(10), 4156–4165. https://doi.org/10.1073/pnas.1804597116
19. Gordon, B. R., Zettelmeyer, F., Bhargava, N., & Chapsky, D. (2019). A comparison of approaches to advertising measurement: Evidence from big field experiments at Facebook. Marketing Science, 38(2), 193–225. https://doi.org/10.1287/mksc.2018.1135

Corresponding email: mhasan1@my.westga.edu

20. Ascarza, E. (2018). Retention futility: Targeting high-risk customers might be ineffective. Journal of Marketing Research, 55(1), 80–98. https://doi.org/10.1509/jmr.16.0163
21. Consumer Financial Protection Bureau (2023b). Chatbots in consumer finance. Issue spotlight. Official publication
22. Gneiting, T., & Raftery, A. E. (2007). Strictly proper scoring rules, prediction, and estimation. Journal of the American Statistical Association, 102(477), 359–378. https://doi.org/10.1198/016214506000001437
23. Chouldechova, A. (2017). Fair prediction with disparate impact: A study of bias in recidivism prediction instruments. Big Data, 5(2), 153–163. https://doi.org/10.1089/big.2016.0047
24. M. N. Amjad, "Strategic Engineering Management for Sustainable and Resilient US Industrial Infrastructure," International Journal of Business Integrity and Technology Advancements, vol. 1, no. 2, pp. 8–20, 2025.
25. Consumer Financial Protection Bureau (2024). Required rulemaking on personal financial data rights. Original final rule released October 22, 2024. Original rule and supporting documents
26. M. Imran, "Explainable AI for cybersecurity risk assessment in US financial institutions," American Journal of Multidisciplinary Knowledge Insights, vol. 2, no. 1, pp. 11–28, 2023.
27. Board of Governors of the Federal Reserve System (2025). Economic well-being of U.S. households in 2024: Banking and credit. Survey of Household Economics and Decision making, May 2025. Report, Figure 24
28. U.S. Department of the Treasury (2026). Treasury releases two new resources to guide AI use in the financial sector. February 19 release accompanying the Financial Services AI Risk Management Framework and AI Lexicon. Treasury release | Framework and lexicon
29. M. N. Amjad, "AI-Enabled Engineering Project Management: Evaluating Productivity and Risk Reduction in the United States," Journal of Computational Systems & Engineering Insights, vol. 3, no. 1, pp. 25–41, 2025.
30. Regulation B (2025). Equal Credit Opportunity Act: Notifications. 12 C.F.R. § 1002.9(b)(2), 2025 annual edition. Official CFR text
31. Forcht Bank, N.A. v. Consumer Financial Protection Bureau (2025). Memorandum opinion and order. No. 5:24-cv-00304-DCR, docket entry 90 (E.D. Ky. Oct. 29, 2025). Court order
32. A. Kumar and K. K. S. Nair, "The Impact of Cloud-Based Business Analytics on Operational Efficiency in US Small and Medium-Sized Firms," Spanish Journal of Innovation and Integrity, vol. 49, pp. 303–318, 2025.
33. M. Imran, "Integrating deep learning and big data analytics for real-time cyberattack detection," Journal of Computational Systems & Engineering Insights, vol. 1, no. 1, pp. 37–50, 2023.
34. Yotzov, I., Barrero, J. M., Bloom, N., Bunn, P., Davis, S. J., Foster, K. M., Jalca, A., Meyer, B. H., Mizen, P., Navarrete, M. A., Smietanka, P., Thwaites, G., & Wang, B. Z. (2026). Firm data on AI. NBER Working Paper 34836 (February 2026 version). https://doi.org/10.3386/w34836 | February 2026 paper
35. Fang, L., Yuan, Z., Zhang, K., Donati, D., & Sarvary, M. (2026). Generative AI and firm productivity: Field experiments in online retail. arXiv preprint 2510.12049, version 3, February 10, 2026. https://doi.org/10.48550/arXiv.2510.12049](https://doi.org/10.48550/arXiv.2510.12049) | February 10 version

Corresponding email: mhasan1@my.westga.edu