AMERICAN JOURNALS
Open Access, Peer Reviewed Journals

American Journal of Business Management, Economics and Banking
ISSN (E): 2832-8078 Volume 48, May - 2026

FROM WALLET REGISTRATION TO SUSTAINED USE IN RETAIL CENTRAL BANK DIGITAL CURRENCY

Tursunova Ibodat Xo‘jayor qizi
Senior Economist, Central Bank of the Republic of Uzbekistan

ABSTRACT

Retail central bank digital currency (CBDC) initiatives are often assessed using wallet registrations, transaction totals and merchant enrolment. These indicators measure different stages of adoption and cannot, by themselves, establish sustained use or additional social value. This article develops a measurement framework that connects two-sided payment adoption to the evaluation of retail CBDC pilots. A critical synthesis of structural models, survey evidence and official implementation reports is combined with accounting bounds for activity and retention. The analysis shows why transaction totals do not identify the number of active users and why repeated cross-sectional activity rates do not identify cohort retention. Recalculation of published model scenarios also illustrates how higher adoption can coexist with greater substitution away from bank deposits. The proposed evaluation design separates access, payment activity, post-incentive retention, merchant acceptance and benefits relative to existing payment services. It links these outcomes to a factorial pilot that varies consumer incentives and merchant integration while holding the underlying CBDC design constant. The contribution is a transparent framework for specifying estimands and preventing unsupported inference from aggregate adoption statistics. The article does not estimate realised CBDC demand or a causal effect of issuance. Its implications concern the evidence required before a successful pilot can be interpreted as a case for wider deployment.

KEYWORDS

Central bank digital currency; payment adoption; network effects; financial inclusion; retention; programme evaluation.
JEL classification: E42; E58; G21; O33.

Introduction

A wallet can be opened without being used. A payment can be induced by a temporary reward without establishing a lasting preference. A merchant can be enrolled without offering a reliable checkout option. These distinctions matter because an apparent increase in CBDC adoption can reflect administrative reach, subsidised experimentation or repeated transactions by a small group rather than durable demand.

The relevant policy question is therefore more demanding than whether people register. Under what conditions does a retail CBDC become a regularly used payment option that improves outcomes relative to the services already available? Answering this question requires separate evidence about

consumers, merchants, persistence and the counterfactual payment environment. A digital public payment instrument may have value, but its public status does not remove the need to demonstrate that value at the margin.

This article makes three connected contributions. First, it translates the two-sided adoption problem into distinct evaluation outcomes. Second, it derives simple bounds that expose the limits of aggregate wallet and transaction statistics. Third, it specifies a pilot design capable of distinguishing temporary incentive effects from improvements in merchant acceptance and sustained use. The bounds are accounting results, not a new behavioural model. Their usefulness lies in showing precisely which conclusions require linked longitudinal data.

The analysis is organised as a conceptual and methodological contribution. It draws on published research and documented implementation episodes rather than confidential central bank records. No observations from an Uzbek CBDC pilot are claimed. The framework is applicable to Uzbekistan and other jurisdictions considering a pilot, but local adoption levels and welfare effects remain empirical questions.

Economic mechanisms and the interpretation of existing evidence

Consumer adoption depends partly on whether merchants accept the instrument; merchant acceptance depends partly on expected consumer demand. Tan (2023) formalises this feedback in a dynamic two-sided model. A consumer-only intervention can therefore encounter an acceptance constraint, while a merchant-only intervention can leave acceptance capacity underused. For evaluation, this mechanism implies that household and merchant outcomes must be measured together, and that treatment spillovers are economically meaningful rather than merely a statistical nuisance.

Design choices also affect the composition of demand. Agur, Ari and Dell'Ariccia (2022) analyse how preferences over payment characteristics and network effects interact with financial intermediation. Their framework cautions against treating adoption as a sufficient welfare criterion. A design can attract users while changing the allocation of funds between cash, bank deposits and central bank liabilities. Whether that change is desirable depends on the underlying frictions and the resulting adjustment by banks and payment providers.

Li (2023) provides a complementary lesson about units of measurement. The structural analysis predicts CBDC demand equivalent to 4–52% of household liquid assets across specifications and design assumptions. This is a model-dependent range for asset allocation, not a confidence interval for the share of people who would use CBDC. Portfolio demand, occasional payment use and daily transaction choice are different outcomes. A pilot questionnaire should therefore avoid asking a single adoption question to stand in for all three.

Privacy evidence requires similar care. Choi et al. (2025) use a randomised survey experiment to study how privacy protection and information affect willingness to use CBDC. The reported increase of up to 64% in privacy-sensitive purchases concerns stated willingness under the experiment; it is not a 64-percentage-point increase in observed payment use. Such evidence supports treating privacy as a potentially important demand determinant. It does not establish how frequently participants would transact after onboarding and incentives end.

The implementation literature reinforces the need for a broader evaluation structure. Koonprasert et al. (2024) organise adoption strategies around regulation, education, design and deployment, and incentives. These categories concern different constraints. A communication campaign may improve understanding but leave merchant integration unchanged; a financial reward may increase trial without

resolving accessibility problems. Separating mechanisms makes both intervention design and interpretation more informative.

Finally, the prevalence of institutional exploration is not evidence of household demand. Illes, Kosse and Wierts (2025) report that 85 of 93 surveyed central banks were exploring retail CBDC, wholesale CBDC or both in 2024. This describes central bank activity across heterogeneous projects. It cannot be used as a denominator or success rate for retail adoption.

Analytical approach and scope

The evidence base combines peer-reviewed economic studies, IMF analytical work and official central bank implementation reports. Sources were selected purposively for their relevance to adoption mechanisms, measurement and pilot evaluation. This is a critical synthesis, not a systematic review or a meta-analysis. The documentary evidence concerns publications from 2022–2025 and the earlier implementation periods that those publications describe; it is not a claim about the current operational status of every CBDC initiative.

Each source is classified by the inference it supports. Structural models generate conditional predictions. Randomised surveys identify effects on survey responses within the experimental setting. Implementation reports document activity under the reporting definitions used by the issuer. None of these categories is interchangeable with a causal estimate of sustained, economy-wide payment use. The numerical comparison below reproduces published scenario values and computes differences; it does not recalibrate or independently validate the underlying model.

The measurement analysis then asks what can be inferred when only aggregate counts are available. Its assumptions are stated explicitly so that the results can be reproduced without access to personal data. The proposed pilot design follows from the resulting identification gaps. It is a research protocol for future evaluation, not an account of an experiment already conducted.

Why aggregate adoption statistics are insufficient

Transactions do not identify active users

Let W denote the number of wallets in a fixed set of eligible registered wallets and T the number of eligible, completed outgoing purchase transactions during one reporting window. Let A be the number of distinct payer wallets making at least one such transaction. If T is positive and each transaction has one eligible payer wallet, then:

$$1 \leq A \leq \min(W, T). \quad (1)$$

The lower bound is attained when one wallet makes all purchases. The upper bound is attained when purchases are distributed across as many distinct eligible payer wallets as possible. Consequently, T/W is the average number of purchases per registered wallet, not the proportion of wallets that are active. For an active share, the corresponding bounds are 1/W and min(1, T/W). A count of transaction endpoints would require a different definition; it cannot be substituted for outgoing purchases without changing the interpretation.

Consider an illustrative registry containing 100,000 eligible wallets and 10,000 outgoing purchases in a month. The same totals are compatible with 100 wallets making 100 purchases each or 10,000 wallets making one purchase each. The implied active-wallet shares are 0.1% and 10%, respectively. Neither is an estimate for a real CBDC. The example establishes that aggregate totals alone cannot distinguish concentrated use from broad participation.

Wallets are also not necessarily people. Multiple wallets per person, shared business accounts and changes in registration rules can break that correspondence. Person-level adoption requires an authorised, privacy-preserving method for linking eligible accounts to a consistent unit of observation. Where such linkage is unavailable, results should remain explicitly wallet-level.

Ree's (2023) account of the eNaira's first year illustrates why activity is a substantive evaluation concern. The report documents limited transaction activity relative to registered wallets and discusses approximately 14,000 weekly transactions. Its implementation context is informative, but a transaction-to-wallet ratio should not be repurposed as a direct estimate of distinct active people. The counting rules, denominator date and reporting window must first be reconciled.

Stable activity rates do not establish retention

Suppose a fixed eligible wallet universe of size W is observed in two consecutive windows. The active-wallet counts are A_0 and A_1, and R is the number active in both. Basic set identities imply:

$$\max(0, A_0 + A_1 - W) \leq R \leq \min(A_0, A_1). \quad (2)$$

When A_0 is positive, dividing these bounds by A_0 gives bounds for retention among wallets active in the first window. The bounds follow from the largest and smallest feasible overlap of two subsets of the same universe. They require consistent activity definitions and eligibility across windows.

For example, if W equals 100,000 and 20,000 wallets are active in each month, retention can range from zero to 100%. Every initial user could have stopped and been replaced, or exactly the same users could have continued. Identical monthly active totals are consistent with both outcomes. A post-incentive retention claim therefore requires cohort linkage rather than repeated headline totals.

New entrants, account closures and unequal exposure periods complicate this comparison. A pilot should fix the cohort and eligibility rules in advance, report entry dates, and allow each included participant the same follow-up horizon. Otherwise, recently onboarded users may appear less active simply because they have had fewer opportunities to transact.

Enrolment and effective acceptance are different margins

The Central Bank of The Bahamas (2023) reported more than 100,000 mobile wallet accounts and more than 1,500 businesses with access to merchant services. These are useful measures of reach. They do not identify unique consumers, the number of outlets reliably accepting payments on a given day, or the distribution of purchase activity. Because both figures are lower bounds, dividing them would not produce a reliable consumer-to-merchant ratio.

Effective acceptance should instead be measured at eligible outlets and relevant purchase occasions. A payment option that is technically enabled but unavailable at checkout provides little practical network value. Merchant acceptance indicators should capture successful payments, rejection reasons, downtime and whether staff can complete refunds. Geographic and sectoral coverage matter because aggregate merchant counts can conceal gaps in everyday purchasing opportunities.

Adoption gains and financial intermediation

Table 1 recalculates selected comparisons from Tan's (2023) baseline and remunerated-CBDC scenarios. In the remunerated scenario, the CBDC rate is 1.5%, compared with a 3% bank deposit rate. The household adoption share rises from 39% to 63%, and firm acceptance from 55% to 79%. At the same time, the reported changes in bank account ownership and the share of savings held in banks become more negative relative to the pre-CBDC benchmark.

Table 1. Selected outcomes in Tan's published CBDC scenarios

Outcome and unit	Baseline	Remunerated	Difference in pp
Households adopting CBDC (%)	39	63	+24
Firms accepting CBDC (%)	55	79	+24
Change in bank account ownership (pp)	−4	−27	−23
Change in share of savings held in banks (pp)	−1	−14	−13

Source: Tan (2023), Section 4.4. Differences are author calculations, remunerated minus baseline.

The final two rows are changes relative to the pre-CBDC benchmark; the first two are adoption levels. pp = percentage points. These are model outputs, not observed treatment effects.

The comparison has a narrow but important interpretation. Within this model, an intervention that improves the adoption margin also changes the intermediation margin. It does not show that remunerated CBDC necessarily reduces lending in every country. Nor does it identify an optimal rate for Uzbekistan. The model's assumptions and banking responses govern the result. The evaluation implication is that a pilot scorecard should report payment use alongside funding substitution, rather than maximise one observed adoption measure.

An additional distinction concerns the counterfactual. A CBDC purchase may replace cash, a bank transfer, a card payment or an otherwise impossible transaction. These substitutions can have different resource and inclusion consequences. Counting every CBDC payment as an additional digital payment would overstate incremental use whenever existing digital instruments are displaced. Surveys linked to a consented subsample, payment diaries and aggregate comparison data can help identify the displaced method, subject to recall and selection limitations.

A pilot design that identifies mechanisms

Outcomes and denominators

Table 2 specifies complementary outcomes. The primary sustained-use outcome should be defined before recruitment, with a fixed window after temporary incentives have ended. One workable definition is the proportion of an eligible onboarding cohort making at least one qualifying purchase in each of two consecutive 30-day follow-up windows. Other definitions are possible, but changing the threshold after examining results would undermine comparability.

Table 2. Outcomes for a retail CBDC evaluation

Dimension	Operational measure	Required comparison or linkage
Access	Eligible people completing onboarding	People offered access and all eligible people
Activity	Distinct wallets making qualifying purchases within a fixed window	Consistent wallet universe and transaction filters
Persistence	Cohort members active in both post-incentive windows	Linked cohort and equal follow-up exposure
Acceptance	Eligible outlets completing payments when requested	Outlet availability and rejected payment attempts
Additional value	Fees, completion time, reliability and accessibility	Existing payment options under comparable conditions

Source: Author framework. Wallet-level measures should not be described as person-level measures without a valid linkage procedure.

Purchase counts should exclude cash-in, cash-out, reversed transactions, system tests and self-transfers under rules specified before analysis. The registration denominator should distinguish all eligible people from those offered access and those completing onboarding. A treatment can improve payment activity among existing users while having little effect on access; conditioning only on completed onboarding would conceal that distinction.

Alongside the primary outcome, the evaluation should report the distribution of purchase frequency and value, the concentration of activity among heavy users, failed transactions and accessibility by relevant population groups. Inclusion should concern previously underserved users and usable services, rather than the mere creation of another account. Where subgroup samples are small, uncertainty should remain visible rather than being replaced by rankings.

Consumer incentives and merchant integration

A factorial design can separate two practical constraints. Eligible geographic or merchant-network clusters can be assigned to four conditions: standard onboarding; a temporary consumer incentive; additional merchant integration support; or both interventions. The underlying CBDC features, safety information and basic service availability should be held constant. Merchant support should be specified concretely, for example as checkout integration, staff training and refund handling, so that the treatment has a reproducible meaning.

Randomisation at cluster level is preferable where consumer and merchant interactions are predominantly local. It does not eliminate spillovers: users can shop across boundaries and merchants can serve untreated users. The protocol should map such exposure, avoid overstating independence and interpret effects as those of the assigned market environment. The primary intention-to-treat estimates should include everyone assigned to an intervention, irrespective of subsequent onboarding or compliance.

The interaction between the two treatments tests whether their joint effect differs from the sum of their separate effects on the chosen outcome scale. A positive interaction would be consistent with complementary constraints in this setting, but would not validate every parameter of a two-sided structural model. Inference should account for clustering and multiple prespecified primary comparisons. Sample size should be determined from a minimum detectable effect, expected cluster correlation and anticipated attrition; no numerical power claim is warranted without those inputs.

Persistence and additional value

Both immediate and post-incentive outcomes are necessary. The intervention may generate trial during the reward period and little subsequent use. Alternatively, it may overcome a learning barrier and produce continued activity. An unchanged aggregate transaction count cannot distinguish these mechanisms without cohort data. Attrition and missing linkage should be reported by assigned group, with sensitivity analysis where their relationship to treatment is plausible.

Additional value requires comparison with existing payment options. Relevant outcomes include total user fees, checkout completion time, failed payments, merchant reconciliation effort and access for previously underserved groups. A fast-payment service may achieve similar improvements at lower implementation cost. Conversely, a CBDC may provide an accessibility or resilience function that transaction averages fail to capture. Such claims require targeted tests of the relevant function and cannot be inferred from registration growth.

Discussion and limitations

The framework changes the interpretation of a successful pilot. Technical availability, broad registration, frequent use, retention and incremental welfare are separate achievements. Progress on one does not establish the others. A defensible expansion decision should combine a prespecified sustained-use criterion with evidence on cost, accessibility, privacy and the consequences for intermediaries.

For Uzbekistan, the immediate implication is to establish the baseline payment problem and the comparison service before choosing an adoption target. Existing retail payment habits, merchant coverage, identification requirements and intermediary incentives may differ materially from those in published models and foreign implementations. The proposed design provides a way to test local mechanisms; it supplies neither a calibrated demand forecast nor a recommendation to issue CBDC.

Several limitations remain. Public reports use inconsistent definitions and do not provide the linked microdata needed to estimate retention. Structural results depend on assumptions that are not validated by this synthesis. Survey intentions may not translate into behaviour, while a small pilot may not reproduce nationwide network effects or bank funding responses. The accounting bounds are deliberately conservative and become wide when information is sparse. Their purpose is to prevent overinterpretation, not to replace better data.

Conclusion

Retail CBDC adoption should be evaluated as a sequence from access to effective acceptance and sustained use. Wallet registrations and transaction totals are valuable operational indicators, but they cannot identify distinct active users, cohort retention or benefits relative to existing payment services without additional information. Published economic research also shows why greater adoption must be considered alongside changes in financial intermediation.

The article offers explicit activity and retention bounds, a set of aligned outcome definitions and a pilot design linking interventions to measurable mechanisms. Its central implication is procedural and testable: determine the unit of observation, comparison service and follow-up horizon before interpreting adoption statistics. Wider deployment should rest on demonstrated additional value and persistent use under locally relevant conditions.

References

1. Agur, I., Ari, A., & Dell'Ariccia, G. (2022). Designing central bank digital currencies. Journal of Monetary Economics, 125, 62–79. https://doi.org/10.1016/j.jmoneco.2021.05.002
2. Central Bank of The Bahamas. (2023, April 3). Public update on The Bahamas digital currency SandDollar. https://www.centralbankbahamas.com/news/press-releases/press-release-public-update-on-the-bahamas-digital-currency-sanddollar-1
3. Choi, S., Kim, B., Kim, Y. S., & Kwon, O. (2025). Central bank digital currency and privacy: A randomized survey experiment. International Economic Review, 66(2), 823–847. https://doi.org/10.1111/iere.12746
4. Illes, A., Kosse, A., & Wierts, P. (2025). Advancing in tandem—Results of the 2024 BIS survey on central bank digital currencies and crypto. BIS Papers No. 159. Bank for International Settlements. https://www.bis.org/publ/bppdf/bispap159.pdf

5. Koonprasert, T. T., Kanada, S., Tsuda, N., & Reshidi, E. (2024). Central bank digital currency adoption: Inclusive strategies for intermediaries and users. Fintech Notes, 2024/005. International Monetary Fund. https://doi.org/10.5089/9798400289422.063
6. Li, J. (2023). Predicting the demand for central bank digital currency: A structural analysis with survey data. Journal of Monetary Economics, 134, 73–85. https://doi.org/10.1016/j.jmoneco.2022.11.007
7. Ree, J. (2023). Nigeria's eNaira, one year after. IMF Working Paper No. 2023/104. International Monetary Fund. https://www.elibrary.imf.org/view/journals/001/2023/104/article-A001-en.xml
8. Tan, B. J. (2023). Central bank digital currency adoption: A two-sided model. IMF Working Paper No. 2023/127. International Monetary Fund. https://doi.org/10.5089/9798400244858.001