Central Bank Digital Currency (CBDC), Bank Deposits, and Digital Payments: First Evidence from a Live Retail CBDC*

Marco Di Maggio Pulak Ghosh Soumya Kanti Ghosh

Nishant Vats Andrew Wu

Abstract

We provide the first evidence on how a live retail CBDC affects bank funding, lending, and payments, using India's digital rupee rollout. The RBI's staggered bank authorization and coordinated marketing campaigns show that CBDC eligibility lowers retail deposits by 2.7%, driven by liquid savings accounts; term deposits are unaffected. Credit declines briefly before recovering. CBDC use displaces deposit-backed digital payments, not cash, and adoption is concentrated among wealthier, uninsured, and more educated depositors. A sufficient-statistics framework implies short-run welfare losses but long-run gains, as household convenience and safety benefits outweigh lost deposit franchise value, temporary credit contraction, and higher funding costs.

Keywords: CBDC, bank deposits, credit, digital payments, financial inclusion
JEL Codes: E42, E51, G21, G28, G51

*Di Maggio: Imperial Business School (m.di-maggio@imperial.ac.uk); Ghosh: Indian Institute of Management Bangalore (pulak.ghosh@iimb.ac.in); Kanti Ghosh: State Bank of India (soumya.ghosh@sbi.co.in); Vats: Olin Business School, Washington University in St. Louis (vats@wustl.edu); Wu: University of California San Diego (anw059@ucsd.edu). We thank Viral Acharya, Sumit Agarwal, David Andolfatto, Ehsan Azarmsa, Lin William Cong, Phil Dybvig, Andreas Fuster, Todd Gormley, Brett Green, Song Ma, Amiyatosh Purnanandam, Uday Rajan, Linda Schilling, Kelly Shue, Alp Simsek, Anjan Thakor, and participants at WashU Brownbag for helpful feedback. We take full responsibility for any errors. We thank the State Bank of India for providing access to de-identified administrative data. The views expressed are those of the authors and do not necessarily reflect those of the State Bank of India. SBI did not provide financial support for this research and had no role in the research design, empirical analysis, interpretation of the findings, or conclusions. SBI had no right to review, approve, or pre-clear the manuscript, the results, or the conclusions before submission or circulation. The authors retained full independence over the analysis and publication decisions. This paper subsumes the results of the paper circulated under the title "Impact of Retail CBDC on Digital Payments, and Bank Deposits: Evidence from India."

1 Introduction

Central banks around the world are actively debating the design and implications of retail central bank digital currencies (CBDCs), even as many have moved from discussion to large-scale experimentation.[1] The stated policy objectives for introducing CBDC are broad, encompassing the modernization of payment systems, increased competition in retail finance, and enhanced financial stability and monetary control. Whether CBDC ultimately improves welfare, however, hinges on its effects along three core margins of the financial system. First, on the banking side, policymakers are concerned that the perceived safety of central bank money could trigger sizable deposit outflows from commercial banks, undermining their role in intermediation and potentially destabilizing the system. Second, on the payments side, it remains unclear whether CBDC will complement or displace existing digital payment systems and change the mix of retail transactions between digital channels and cash. Third, on the household side, CBDC adoption may broaden access to safe and liquid assets, or instead remain concentrated among already-advantaged users, leaving existing financial inclusion gaps largely intact. These concerns are especially salient in emerging markets, where banking systems rely heavily on retail deposit funding, cash remains central to retail transactions, and a substantial share of the population remains unbanked or underbanked.

These considerations have generated a substantial theoretical literature, but its conclusions for deposits, lending, and inclusion are sign-ambiguous and highly sensitive to assumptions about CBDC design and household behavior (see Carapella and Flemming, 2020; Niepelt, 2021; Auer et al., 2022b; Chapman et al., 2023; Duffie and Economy, 2022; Ahnert et al., 2024; Infante et al., 2022, 2023, 2024, for a review of the literature). Yet, despite extensive theoretical and policy debate, there is still little direct evidence from large-scale, live CBDC deployments on these core questions. In particular, we lack evidence from a large-scale retail CBDC in a setting with both an advanced digital payments rail and a cash-intensive, partially excluded population.

This paper attempts to address this gap by providing the first empirical evidence on the real-world effects of a retail CBDC using the introduction of India's retail digital rupee (e-rupee) in December 2022 in a setting that closely matches the environments studied in recent quantitative models. Our central finding is that a non-interest-bearing, capped retail CBDC generates sizable changes in deposits and digital payments, but only a transitory effect on bank lending. Banks experiencing CBDC exposure see a decline in retail deposits; however, the resulting contraction in lending is modest and short-lived. This pattern is consistent with the prediction of models such as Whited, Wu, and Xiao (2023) in which banks actively rebalance their funding structure to absorb deposit outflows,

[1] As of May 2026, 41 jurisdictions have launched active CBDC pilot projects, including 14 G20 economies. Auer, Cornelli, and Frost (2023) present a comprehensive database of CBDC projects across the globe, and investigate the economic and institutional factors that correlate with CBDC project efforts. For an in-depth overview of country-specific CBDC initiatives, refer to the Atlantic Council Project.

thereby insulating credit supply. On the payments margin, CBDC adoption primarily substitutes away from existing deposit-backed digital payment methods, with little detectable effect on cash usage suggesting that CBDC competes more directly with bank-based digital payments than with physical currency. On the household side, early adoption is concentrated among male, higher-income, highly educated, and uninsured depositors, consisting largely of households that are already well integrated into the formal banking system. Consequently, the initial expansion in access to central bank money accrues to relatively advantaged, digitally connected households, rather than to those at the margins of financial inclusion. Finally, embedding our empirical estimates in a simple sufficient-statistics welfare framework, we find that the policy is welfare-improving on net in the long-run. Gains to households outweigh costs to the banking sector, and, importantly, the credit contraction that features prominently in policy discussions appears transitory rather than persistent.

To organize these findings, we develop a simple framework of household portfolio choice and payment adoption. Households allocate wealth across savings deposits, term deposits, cash, digital payments (UPI), and CBDC. They trade off liquidity, returns, payment convenience, and safety when making these choices. A key distinction is that bank deposits are only insured up to a limit, whereas CBDC provides a direct, default-free claim on the central bank. The framework highlights three margins that map directly to our empirical design. First, households are likely to fund CBDC by reducing liquid instruments, especially savings deposits, rather than higher-return, illiquid instruments such as term deposits. CBDC should therefore affect both the level and composition of bank funding. Second, CBDC is most likely to compete with existing digital instruments that offer similar transactional features, suggesting substitution away from UPI rather than from cash. Third, adoption incentives are stronger for households with uninsured deposit balances and greater financial sophistication, who value the safety and transactional flexibility of central bank money. We use these predictions to discipline our empirical strategy and to interpret the resulting patterns in deposits, payments, and adoption.

Our setting is the Reserve Bank of India's (RBI), India's central bank, retail digital rupee. The e-rupee (e₹) is a token-based, non-interest-bearing CBDC issued by the RBI and distributed through commercial banks under a two-tier, "intermediated" model. Retail users can open CBDC wallets only at participating banks and must fund them by converting balances from their bank deposit accounts, subject to a holding cap of ₹100,000 per user. The e₹ therefore operates alongside a large deposit-funded banking system, a cash-intensive ecosystem, and a widely used digital payments platform in UPI, yet it is meaningfully distinct from each: unlike bank deposits, it bears no interest; unlike bank deposits and UPI wallets, it carries no default risk; and unlike cash, it cannot be held anonymously. From the perspective of the theoretical literature, this is a conservative retail CBDC design implemented in a mature financial system.

Empirically, our strategy exploits two sources of quasi-experimental variation and a combination of aggregate and micro data. At the district level, we use the interaction of India's long-standing lead bank system with the RBI's bank-level authorization policy to identify the effects of CBDC access on aggregate deposits and credit. At the individual level, we exploit centrally scheduled CBDC marketing waves within a single large bank to trace out household portfolio, payment, and adoption responses. In both cases, the identifying variation comes from the timing of CBDC access or promotion, and we validate the parallel trends assumption.

At the district level, our empirical design exploits the interaction of India's lead bank system with the RBI's bank-level authorization policy. Since 1969, each district has had a designated lead bank, typically a public-sector bank (Burgess and Pande, 2005). These lead banks dominate retail intermediation and are responsible for branch expansion and credit planning in that district (Reserve Bank of India, 2018). The RBI granted CBDC authorization to banks in a staggered sequence as they demonstrated operational readiness, independently of short-run district conditions. As a result, the timing of CBDC access for households varies across districts according to historical lead-bank assignments and the RBI's bank-level sequencing, rather than recent local trends in deposits or digital payments. Our difference-in-differences specification compares districts whose lead banks become CBDC-eligible at different dates, controlling for district and time fixed effects. We employ the Callaway and Sant'Anna (2021) estimator to address concerns related to the staggered roll-out.[2] The key identifying assumption is that, absent CBDC, districts led by early- and late-authorized banks would have followed parallel trends in deposits and credit. We assess the validity of this assumption by examining pre-treatment event-time coefficients for all our analyses. The plausibility of this assumption is strengthened by the fact that lead-bank assignments were fixed decades before CBDC was conceived and by institutional accounts indicating that the RBI's authorization sequence reflected banks' internal technical readiness rather than granular district-level deposit or payments trends.

At the individual level, we exploit the staggered timing of intensive CBDC marketing campaigns run by one of India's largest state-owned banks participating in the pilot. In response to RBI guidance, the bank's central digital and marketing teams designed standardized campaign content and rolled it out across districts in several discrete waves. Internal documentation and RBI communications indicate that the sequencing and targeting of these district waves were determined centrally as part of a nationwide roll-out plan, rather than by local branch managers reacting to short-run changes in CBDC interest or economic conditions. By the end of our sample window, all districts served by the bank had received both technological access and marketing. We combine this staggered timing with granular depositor-level data, including monthly CBDC and UPI transactions, ATM

[2]Throughout our empirical analysis, we show that the results are robust to using alternative estimators, such as Borusyak, Jaravel, and Spiess (2024) and de Chaisemartin, D'Haultfœuille, and Vazquez-Bare (2024). Moreover, the results are robust to employing long differences to make the visual presentation more comparable with the traditional TWFE event-study regressions.

withdrawals, and deposit balances, matched to rich demographics and geocoded branches, to estimate event-study difference-in-differences models of portfolio and payment behavior and Cox proportional hazard models of CBDC adoption. Our specifications include individual and time fixed effects, so identification comes from within-person changes in outcomes around the campaign date, netting out all time-invariant heterogeneity in demographics, baseline financial behavior, and branch characteristics. The identifying assumption is that, conditional on fixed effects, the exact month in which a district is assigned to the intensive campaign is orthogonal to the short-run evolution of its residents' deposit and payment behavior. This assumption implies that the treated and the control group (not-yet-treated group) would have evolved according to parallel trends in the absence of the treatment. We provide evidence supporting this assumption by presenting an assessment of pre-trends for all our analyses.

Our first set of results concerns bank funding and lending. Using the lead-bank variation, we find that the introduction of CBDC eligibility is associated with a persistent decline in total retail deposits and a smaller, transitory decline in credit. The average treatment effect across post-treatment periods implies a 2.7% decline in retail deposits, which at the median district level of deposits in 2022 Q3 corresponds to a reduction of roughly ₹1.8 billion (≈ $88.5 million in PPP terms). Total credit falls by 2% on average, but this effect reverts toward zero within about six to seven quarters. Credit initially co-moves with deposits, consistent with short-run funding pressure, but then partially recovers even though deposits remain lower. This pattern suggests that banks adjust on several margins, for instance, raising deposit rates, tapping wholesale markets, or reallocating liquid assets, to support lending in the medium run. Given that the e₹ is non-interest-bearing and subject to a ₹100,000 holding cap, the deposit effect we estimate is likely a lower bound effect in magnitude relative to CBDC designs that are either uncapped or interest-paying.

Our second set of results centers on the composition of households' balance sheets and payments. Exploiting the staggered marketing campaigns at our sample bank, we show that CBDC-related outflows come almost entirely from liquid low-yield savings accounts, with no economically meaningful changes in illiquid and high-yield deposit instruments, such as fixed or recurring deposits. This result suggests that households fund CBDC from the liquid, transactions-oriented portion of their portfolio.

On the payments side, we find that CBDC primarily displaces existing digital transaction methods, UPI in the Indian context, rather than cash. After the campaign, UPI volume falls by about 6.3%, and the number of UPI transactions declines by 0.85%. ATM withdrawal volume, in contrast, does not change significantly. Decomposing the UPI effect into person-to-person (P2P) and person-to-merchant (P2M) transactions, we find that P2P usage falls nearly twice as much as P2M usage. Overall, these results are consistent with a key model proposition that CBDC is a strong competitor to deposit-backed digital payments but a weak substitute for cash.

Our third set of results characterizes CBDC adoption patterns. These estimates are descriptive

rather than causal. Cox proportional hazards models with ZIP-code fixed effects reveal a steep and economically large adoption gradient. Male customers adopt at roughly 2.3 times the hazard rate of female customers. Relative to the lowest income quartile, individuals in the top income quartile exhibit more than a fivefold increase in adoption hazards, while those with deposit levels above the deposit insurance threshold of ₹500,000 adopt at more than four times the rate of those with deposits below the threshold. Educational attainment also matters: individuals with degrees in fields such as computer science, finance, or engineering adopt at twice the hazard rate of others. While prior UPI usage is positively correlated with adoption in univariate specifications, this relationship reverses once we control for income, education, and other covariates. Conditional on these characteristics, prior UPI users adopt more slowly than comparable non-users, consistent with lower incremental benefits when an effective digital payment option is already in place.

Overall, early adoption appears to be concentrated among men, higher-income individuals, depositors with balances above the deposit insurance limit, and financially sophisticated individuals. This concentration directly maps to the household margin in our framework: households with uninsured balances and greater financial sophistication derive the greatest safety and portfolio benefits from holding central bank money. These adoption patterns suggest two key policy implications. First, on the financial inclusion margin, they indicate that a conservatively designed retail CBDC primarily expands access to central bank money for already well-banked, digitally connected users, rather than immediately reaching households at the margins of the formal financial system. Second, they suggest that CBDC offers uninsured depositors an additional safe asset and portfolio instrument, in the spirit of Ghosh, Limodio, and Vats (2025), which can potentially alter how these households manage risk and liquidity in the banking system.

Taken together, our reduced-form results suggest that the welfare effect of a retail CBDC is theoretically ambiguous. On one side, CBDC eligibility reduces bank deposits and initially tightens credit, and it shifts banks' funding away from cheap, insured retail deposits toward more expensive wholesale sources. On the other side, adopters reveal substantial gains in convenience, safety, and segmentation by voluntarily moving into a capped, non-interest-bearing wallet, and the credit effects revert as banks adjust their funding mix. This pattern motivates a welfare exercise that takes the reduced-form estimates as inputs rather than treating them as ancillary. We therefore extend our framework to incorporate the CBDC holding cap and a simple banking sector that links retail deposit funding to lending, and we use a sufficient-statistics approach to express the first-order welfare impact in terms of elasticities that our research design recovers directly. The resulting welfare expression decomposes aggregate effects into six components: household surplus from CBDC adoption, the deposit franchise that banks lose when balances migrate, a deadweight loss from moving up the wholesale funding supply curve, a deadweight loss from banks cutting loans, a credit-contraction

externality, and a net seigniorage which captures the fiscal value of replacing commercial bank liabilities with non-interest-bearing central bank money. Evaluated at the median treated district, the model delivers a net annual welfare gain that is negative in the short-term but economically meaningful and positive in the long-term, with household convenience and safety accounting for most of the gross gains and net seigniorage playing a supporting role. On the cost side, the credit-contraction externality that dominates policy discussions is sizable on impact but fades as banks expand wholesale funding and credit recovers, while the remaining permanent costs come from the loss of deposit franchise value. These bank-side costs are quantitatively smaller than the user-side gains, so in our setting, credit disintermediation is a transitional adjustment rather than a permanent tax on the banking system, and the convenience and safety value revealed through adoption dominates in the long run.

Related Literature: Our paper contributes to the literature on CBDC, which remains largely theoretical due to the lack of data availability. Under a variety of assumptions about CBDC design and household behavior these papers focus on: (1) the effect of CBDC on deposit funding and intermediation (García et al., 2020; Andolfatto, 2021; Adalid et al., 2022; Piazzesi and Schneider, 2022; Williamson, 2022b,a; Chang et al., 2023; Chiu et al., 2023; Keister and Sanches, 2023; Whited, Wu, and Xiao, 2023; Burlon, Muñoz, and Smets, 2024; Niepelt, 2024; Li, Usher, and Zhu, 2025), (2) the impact of CBDC on financial stability and broader macroeconomic outcomes, such as monetary policy transmission, (Bordo and Levin, 2017; Meaning et al., 2018; Brunnermeier, James, and Landau, 2019; Brunnermeier and Niepelt, 2019; Duffie et al., 2019; Barrdear and Kumhof, 2022; Davoodalhosseini, 2022; Garratt, Yu, and Zhu, 2022; Kim and Kwon, 2022; Das et al., 2023; Gust, Kim, and Ruprecht, 2023; Carapella et al., 2024; Murakami, Shchapov, and Viswanath-Natraj, 2024; Ahnert et al., 2025; Abad, Nuño, and Thomas, 2025; Paul, Ulate, and Wu, 2026), and (3) the design of CBDC to ensure financial stability or promote financial inclusion (Bindseil, 2020; Maniff, 2020; Assenmacher et al., 2021; Duffie, 2021; Fernández-Villaverde et al., 2021; Kumhof and Noone, 2021; Schilling, Fernández-Villaverde, and Uhlig, 2021, 2024; Shy, 2021; Agur, Ari, and Dell'Ariccia, 2022; Auer et al., 2022a; Keister and Monnet, 2022; Minesso, Mehl, and Stracca, 2022; Li, 2023; Tan, 2024; Wang and Hu, 2025). We add to this literature by presenting the first-ever empirical evidence from a large developing economy, with an established digital payments rail and a large cash economy, that introduced a non-interest-bearing, capped retail CBDC.

Our work contributes to the policy and academic debate on CBDC along three dimensions. First, we show that a non-interest-bearing, capped retail CBDC is associated with a persistent decline in deposits, but the resulting contraction in credit is modest and short-lived. Moreover, we find that CBDC mainly displaces liquid, low-yield savings deposits, with little effect on illiquid, high-yield instruments such as term and recurring deposits, implying changes in both the level and the composition of bank funding. Together, these results speak directly to the theoretical literature on CBDC that discusses its

effect on bank funding and intermediation. Specifically, our results are consistent with models such as Whited, Wu, and Xiao (2023), in which banks offset deposit outflows by adjusting their balance sheets, so that CBDC primarily affects deposits rather than credit supply per se.

Second, we show that CBDC introduction coincides with a sharp decline in the use of existing digital payment instruments, but only a modest reduction in cash usage. This joint pattern for liquid deposits and digital transaction volumes suggests that retail CBDC functions primarily as a means of payment that competes with incumbent digital rails. This mechanism is consistent with recent work by Berg et al. (2024), who find that positive news about a potential digital euro lowers the stock prices of U.S. payment firms, which they interpret as evidence that CBDC acts as a means of payment and competes with established payment platforms. We build on their findings by providing direct, individual-level evidence on how digital payment behavior adjusts following the roll-out of a live retail CBDC. Moreover, our results are informative about the effect of CBDC on cash-based payment methods.

Third, we show that early CBDC adoption is heavily skewed toward households with large deposits, higher incomes, and specialized training in fields such as computer science, finance, or engineering. This pattern implies that the initial diffusion of CBDC primarily reallocates safety and convenience among already-advantaged users rather than immediately narrowing financial inclusion gaps. It is consistent with Auer et al. (2022a), who argue that broad financial and digital literacy efforts are likely needed for inclusive CBDC take-up, and with Bijlsma et al. (2024), who find that financial knowledge predicts stated CBDC adoption.

Finally, our work informs the CBDC debate on the classic question of narrow banking, which asks whether the safekeeping of liquid money-like claims and the origination of risky assets should, in principle, be separated across different institutions. In our setting, a non-interest-bearing, capped retail CBDC effectively carves out part of households' liquid deposit base into a public, default-free transactions asset, while leaving credit intermediation on commercial banks' balance sheets. This institutional structure parallels narrow-banking proposals surveyed by Pennacchi (2012), in which safe transaction liabilities are backed by central bank or government assets and credit risk is borne elsewhere in the system. Viewed through this lens, our evidence provides an empirical benchmark for how far a conservatively designed retail CBDC can push the system toward a narrower banking structure by reducing reliance on deposit-funded payments while leaving the core of bank intermediation largely intact.

The rest of the paper is organized as follows. Section 2 describes the institutional features of the digital rupee, the lead bank system, and the CBDC marketing campaigns. Section 3 introduces the data. Section 4 presents the conceptual framework and derives the main propositions. Section 5 outlines our empirical strategies at the district and individual levels. Section 6 documents the effects of CBDC

on bank deposits and credit. Section 7 studies substitution between CBDC, UPI, and cash. Section 8 analyzes the determinants of CBDC adoption. Section 9 adopts a sufficient-statistics approach to quantify the overall welfare impact of CBDC introduction and decompose the contribution of the underlying channels. Section 10 concludes with policy implications and directions for future research.

2 Institutional Details

India's payments economy is dominated by cash. In 2017, cash transactions represented 90% of the total number of payment transactions (PIB Thiruvananthpuram, 2023). Furthermore, the Reserve Bank of India (RBI), India's central bank, found that cash was the preferred method for regular expenses, while digital methods were the second-preferred method. However, the costs of running a cash-reliant economy are high. For instance, the RBI spent 49 billion Rupees printing cash, not including environmental, social, governance, and operational costs (Reserve Bank of India, 2022a). One of the main reasons India is exploring a digitized economy fueled by a digital payment system and a central bank digital currency (CBDC) is the high cost and inefficiency of its physical cash system. Thus, the RBI spearheaded the introduction of two revolutionary digital payment systems in the country: the Unified Payments Interface (UPI) and the digital rupee CBDC (e₹). We direct readers to Dubey and Purnanandam (2025), Alok et al. (2025), Cramer et al. (2025), and Appendix A.1 for a detailed discussion on UPI and present a detailed discussion of the digital rupee in this section.

The Digital Rupee: The RBI started the process of launching its CBDC by first proposing an expansion of the definition of a "bank note" in the RBI Act, 1934 to include digital currencies in October 2021. The Cryptocurrency and Regulation of Official Digital Currency Bill, 2021 was subsequently introduced in the 2021 winter session of Parliament to create a facilitative framework for creation of the official digital currency to be issued by the Reserve Bank of India (RBI). In a speech in July 2021, T. Rabi Shankar, the deputy governor of the RBI, defined a CBDC as a "legal tender issued by a central bank in a digital form," which "is the same as fiat currency and is exchangeable one-to-one with the fiat currency" (Kaul, 2021). To further raise awareness and understanding of CBDCs, the RBI issued its CBDC concept note on October 7, 2022 announcing the phased introduction of its CBDC, the digital rupee, and detailing the motivations, objectives, design, technology and policy choices, and the benefits and risks of the e₹ (Reserve Bank of India, 2022a).

The RBI and the Ministry of Finance stated several objectives for launching CBDC: reducing the operational costs of physical cash management, fostering financial inclusion, and enhancing the resilience, efficiency, and innovation of the payments system. Additional stated aims include improving settlement efficiency in financial markets, boosting innovation in cross-border payments, and offering the public a sovereign digital alternative to private virtual currencies without their associated risks.

The RBI launched the retail CBDC pilot on December 1, 2022. After the initial launch, the

RBI implemented a staggered roll-out of the CBDC to assess technology and system robustness and generate evidence on the CBDC's "impact on users, on monetary policy, on the financial system and on the economy" (Reserve Bank of India, 2024, 2026).

The retail digital rupee is legal tender issued under a two-tier, "intermediated" model: the RBI creates the currency, while banks and other service providers distribute it via dedicated mobile applications and link it to existing digital payment rails such as the UPI (Reserve Bank of India, 2026). Moreover, the RBI imposes a cap on the amount that users can hold in CBDC wallets. Specifically, retail users load funds into their wallets from their bank deposit accounts, up to a holding limit of ₹100,000 (about $4,900 in PPP terms). The CBDC is non-remunerative, meaning that it does not bear interest, and it is token-based, with each token functioning like a banknote whose holder is the owner. Users transact with other users (P2P) or with merchants (P2M) using QR codes (Reserve Bank of India, 2022b). In essence, the retail digital rupee is a non-interest-bearing, cash-like token issued by the central bank and distributed by commercial banks.

Although the digital rupee is distributed through commercial banks, it is not a liability of those banks. Instead, it is issued and backed by the Reserve Bank of India, so it appears as a direct liability on the RBI's balance sheet. This makes it fundamentally different from a bank deposit, which is a claim on a commercial bank. In that sense, the digital rupee functions more like digital cash than like money held in a regular bank account.

Appendix Figure A.4 shows the home screen of a CBDC wallet application, displaying the transaction functions that correspond to the pay and collect transaction types we observe in our CBDC transaction data. Appendix Figure A.5 shows a digital rupee acceptance display at a merchant location, illustrating that CBDC payments are made via QR codes that are visually similar to existing UPI QR codes, some of which are shown in Appendix Figure A.2.

Since retail users can only open CBDC wallets and convert bank deposits into e₹ if their bank has received RBI authorization and has deployed a functioning CBDC application, eligibility depends jointly on the RBI's bank-level authorization timing and each bank's internal technology roll-out capacity. Importantly, the RBI's authorization decisions were made at the bank level and implemented in stages as banks demonstrated operational and technical readiness to participate in the pilot (Reuters, 2022; Reserve Bank of India, 2026).

Table 1 documents the staggered roll-out of CBDC authorization across India's major public sector banks. State Bank of India (SBI) was authorized in 2022 Q4, followed by Bank of Baroda, Canara Bank, and Union Bank of India in 2023 Q1, Punjab National Bank in 2023 Q2, and other public sector banks in subsequent quarters.

2.1 Lead Bank Assignments and Geographic Variation in CBDC Access

The RBI's bank-level authorization policy generates district-level variation in CBDC access through India's *lead bank* system. Since the introduction of the Lead Bank Scheme in December 1969, the RBI has assigned one public sector bank as the designated lead bank for each district (Burgess and Pande, 2005). The lead banks have a mandate to coordinate financial inclusion, expand branch networks, and oversee credit planning in that district (Reserve Bank of India, 2018). These assignments were made on historical and administrative grounds, such as legacy branch networks and administrative jurisdiction, and have remained remarkably stable over time (Dubey and Purnanandam, 2025). As a consequence, lead banks maintain a dominant branch presence in their assigned districts, particularly in rural and semi-urban areas, and intermediate the majority of retail deposits and payment relationships in those districts.

Because the RBI granted CBDC authorization at the bank level rather than the regional level, a district's effective access to CBDC is largely determined by when its lead bank becomes CBDC-eligible. When a lead bank receives authorization and deploys its CBDC app, households in that district gain access to CBDC through the bank that already serves as their primary financial intermediary. This institutional structure converts bank-level authorization into a staggered roll-out of CBDC *treatment* across districts. Districts led by the State Bank of India (SBI) gain access as early as 2022 Q4. Districts led by other authorized banks gain access in later quarters, as their lead banks enter the pilot. Districts led by never-authorized banks remain untreated throughout the sample window. Figure 1 presents the geographic distribution of CBDC roll-out across lead bank districts, illustrating the spatial variation in treatment timing determined by lead bank assignments.

The combination of historical lead bank assignments and the RBI's bank-level CBDC authorization decisions yields plausibly exogenous variation in the timing of CBDC access across districts. First, the identity of a district's lead bank was fixed decades before CBDC was conceived and is driven by administrative considerations unrelated to recent shifts in digital adoption or deposit trends at the district level. Second, the RBI's sequencing of CBDC authorization reflects banks' internal technical readiness, rather than granular district-level economic conditions or contemporaneous trends in retail deposits or digital payment usage at the district level. Third, the mapping from bank-level authorization to district-level treatment operates mechanically through the lead bank system: once a lead bank is authorized, its entire district portfolio transitions from non-eligible to eligible for CBDC, regardless of local economic conditions or the district's appetite for digital innovation.

2.2 Staggered Roll-Out of CBDC Marketing Campaigns Within Banks

Our second source of identification comes from the CBDC marketing strategy of our sample bank, one of the largest public-sector banks participating in the RBI's retail CBDC pilot. In response to

the RBI's guidance to scale up e₹ usage, the sample bank rolled out a series of centrally designed CBDC marketing campaigns targeted at selected districts. These campaigns included targeted in-app notifications, SMS messages, and branch-level onboarding drives. As a result, these campaigns were implemented in a staggered manner across districts. Some districts received marketing interventions earlier while others received them later. By the end of our sample window, all districts served by our sample bank were covered by these campaigns. Figure 2 presents the geographic distribution of CBDC roll-out across districts, illustrating the spatial variation in treatment timing determined by our sample bank. Appendix Figure A.3 presents examples of CBDC marketing materials distributed by participating banks.

While the campaign content and design were coordinated by the bank's central digital banking and marketing teams, the sequencing of the staggered roll-out of the campaign across districts was decided by the Reserve Bank of India. This staggered roll-out of CBDC marketing campaigns across districts provides a second, conceptually distinct experiment that we use to identify the effect of CBDC promotional exposure on individual-level deposits, usage of digital payment systems, and cash transactions. The key institutional feature is that campaign timing and targeting were determined by the RBI as part of a nationwide roll-out plan, rather than by branch-level managers responding to local variation in CBDC take-up or economic conditions. This design implies that, conditional on the district's baseline characteristics, the exact month in which a district is assigned to the intensive CBDC marketing campaign is plausibly orthogonal to short-run changes in local economic outcomes. Moreover, because all districts served by our sample bank ultimately become technologically eligible for CBDC, receive marketing campaigns, and are served by the same bank, our identification of the marketing effect relies on within-bank, within-district timing variation in promotional intensity, rather than on cross-bank differences in technology or product design.

We exploit the two sources of variation to understand the effect of the CBDC at different levels of aggregation. The lead bank variation identifies the effect of the introduction of CBDC on aggregate bank deposits and credit at the district level. The marketing campaign variation identifies the effect of CBDC promotional exposure on individual households' use of deposits, UPI, and cash.

2.3 Aggregate CBDC Usage Patterns in India

Since the launch of the retail digital rupee pilot, CBDC usage in India has expanded rapidly along both the extensive and intensive margins. Between February 2023 and March 2026, the number of registered CBDC users rose from about 80 thousand to over 10.6 million, implying more than a hundred-fold increase in the user base in just over three years. Over the same period, aggregate monthly CBDC transaction volume rose to roughly ₹340 billion, with particularly pronounced month-on-month increases during 2024 as the pilot scaled nationally and participating banks intensified onboarding

efforts. Figure 3a shows the joint evolution of the number of CBDC users and aggregate monthly transaction volumes over this period.

The time path of adoption was strongly convex early on and then gradually flattened. In the first phase of the pilot (February 2023–December 2023), the number of users reached roughly 3.8 million and monthly transaction volume reached about ₹1.2 billion by the end of 2023. The second phase (January 2024–December 2024) was characterized by especially rapid growth: monthly transaction volume increased from around ₹6.0 billion in January 2024 to roughly ₹135 billion by December 2024, and the user base rose from about 4.2 million to almost 6.0 million over the same period. In the most recent phase (January 2025–March 2026), growth remained strong but became more gradual: the number of users expanded from roughly 6.1 million to 10.6 million, while monthly transaction volume climbed from around ₹153 billion to ₹340 billion, indicating continued deepening of usage on both the extensive and intensive margins.

Viewed relative to incumbent payment rails, CBDC became a secondary but non-negligible component of the overall payments landscape. Figure 3b shows the evolution of CBDC transaction volumes relative to UPI and cash and illustrates the gradual increase in CBDC's share of aggregate payment activity. Measured as a share of UPI transaction volume, CBDC usage increased from essentially zero in early 2023 to around 0.1 percent by March 2024, and then rose further to slightly above 1 percent by early 2025, fluctuating in the neighborhood of 1–1.2 percent through March 2026. Relative to cash, the increase was more pronounced: CBDC transaction volume reached about 0.65 percent of cash volume by March 2024 and approximately 12 percent of cash volume by March 2026. Over this horizon, the digital rupee evolved from a newly introduced instrument to one that accounted for a noticeable share of aggregate payment volumes, especially when benchmarked against cash.

Two broad patterns emerge from these aggregate series. First, CBDC adoption and usage evolve in a nonlinear fashion, with a concentrated period of rapid scale-up around the nationwide roll-out followed by a more gradual expansion as the user base becomes larger. This timing is consistent with the institutional roll-out described above, in which the Reserve Bank of India staggered bank authorizations and marketing campaigns across districts. Second, CBDC appears primarily as an additional digital rail layered on top of an already large digital and cash payments infrastructure: by 2026, its volume remains well below the scale of UPI and total cash usage, but is nonetheless substantial enough to warrant examining how it interacts with bank funding and existing digital payment rails.

3 Data

This section describes the data used in the analysis. The empirical work draws on three primary sources: the timing of CBDC roll-out, district-level data on bank deposits and credit, and individual-level data

on retail CBDC usage, bank deposits, UPI transactions, and ATM withdrawals. District-level deposits and credit are obtained from the Reserve Bank of India's publicly available *Quarterly Statistics on Deposits and Credit of Scheduled Commercial Banks*. The individual-level data come from one of India's largest state-owned commercial banks, which, as of September 2025, served approximately 500 million customers and accounted for 23% of total banking assets and 25% of aggregate deposits and lending activity in India.

3.1 CBDC Roll-Out

First, we assemble bank-level and district-level data on the staggered roll-out of CBDC access for lead banks and the staggered roll-out of the CBDC marketing campaign for the sample bank, respectively. Table 1 reports the timing of CBDC authorization across bank-level treatment cohorts.

For our sample bank, we also observe the districts targeted at each expansion stage: December 2022, January 2023, February 2023, May 2023, June 2023, and November 2023. These districts were not selected for access per se, but rather for targeted outreach and widespread marketing campaigns designed to educate residents about the CBDC and its usage. By November 2023, all districts served by our sample bank had been reached by these efforts. Figure 2 presents a map of India documenting the geographic distribution of districts targeted in each cohort.

3.2 District-Level Bank Deposits and Credit

We collect data on district-level bank deposits and credit from the Reserve Bank of India's publicly available *Quarterly Statistics on Deposits and Credit of Scheduled Commercial Banks*. Specifically, Statement No. 4B reports the quarterly outstanding stock of deposits at the district level for all scheduled commercial banks, and Statement No. 4A reports the corresponding stock of bank credit. The sample spans 2018 Q1 through 2025 Q3. We focus on total retail deposits, defined as the sum of savings deposits in interest-bearing retail accounts subject to regulated rates of return and term deposits placed with the bank for a specified maturity at a predetermined interest rate. Panel A of Table 2 reports summary statistics for quarterly retail deposits and credit at the district level.

3.3 Individual-Level Data

The individual-level data come from one of the largest state-owned commercial banks in India and include customer demographics, granular retail CBDC and UPI transaction records, monthly ATM withdrawal volumes, and monthly snapshots of deposit balances.

We observe self-reported demographics for all depositors in our sample, including gender, annual income, age, education, ZIP code, and the branch at which the customer primarily banks. Using the location of the home branch, we identify the date on which each account was exposed to the CBDC marketing campaign. Panel D of Table 2 reports summary statistics for the demographics we

observe. The sample contains 153,875 individual accounts. Among these, 4.4% have adopted CBDC and 49.8% have adopted UPI. The sample is 59.2% male, with a mean age of 39.4 years. Mean log income is 11.97, corresponding to an average annual income of approximately ₹159,000 (≈ $7,800 in PPP terms), and less than 1% of the sample holds a sophisticated education credential. Mean deposit holdings are ₹15.7 million, with a standard deviation of ₹53.4 million. The large standard deviation relative to the mean, together with a median deposit of ₹352,390 that is well below the mean, indicates substantial right skewness in the distribution of deposits, reflecting considerable wealth inequality within the sample.

Our granular CBDC data contain more than 630,000 transactions from December 2022 through March 2024, organized at the user-ID level. We also have access to granular UPI data for our sample depositors, which contain more than 48 million transactions from April 2022 through March 2024. Lastly, we observe monthly ATM withdrawal volumes at the depositor level from April 2022 through March 2024, which we use as a proxy for cash usage.

Our dataset also provides individual-level monthly snapshots of savings, fixed-deposit, and recurring-deposit balances over the same period. Savings deposits are interest-bearing accounts subject to regulated rates of return and can be withdrawn on demand. Fixed deposits are lump-sum investments at a predetermined rate for a fixed maturity, whereas recurring deposits involve periodic contributions over a fixed term.[3] The deposit data include month-end savings balances, month-average savings balances, and the balances in fixed and recurring deposits. For each month, we define total deposits as the sum of the month-average savings balance, the fixed-deposit balance, and the recurring-deposit balance.

We merge the CBDC transaction data with the UPI, ATM withdrawal, and deposit-balance data for each depositor to construct the analysis panel with depositor-month as the unit of analysis. Panel C of Table 2 reports summary statistics for monthly CBDC and UPI usage, ATM withdrawals, and deposit balances. The sample contains 1,591,968 depositor-month observations. Overall, CBDC adoption remains modest: the mean log CBDC volume is 0.034 and the mean log number of transactions per month is 0.007. By contrast, UPI usage is far more widespread: the mean log UPI volume is 10.021 and the mean log number of transactions per month is 2.355. These patterns indicate that UPI is the dominant digital payment method, while CBDC remains in an early stage of adoption during the sample period.

Cash remains an important payment instrument: the mean log ATM volume is 3.031. Turning to deposit balances documented in Panel B, mean log savings balance is 13.207, mean log fixed deposit balance is 2.645, and mean log recurring deposit balance is 0.638, implying that savings accounts are the primary repository of deposits for most customers, with fixed and recurring deposits playing a

[3]For additional details on deposit account types in India, see here.

smaller role. Mean log total deposit balance is 13.641, with substantial cross-sectional dispersion reflected in a standard deviation of 3.043.

4 Conceptual Framework

This section develops a simple portfolio and payments model to organize the effects of retail CBDC on bank deposits and payment behavior. The model delivers three main implications, which we test in Sections 6–8: (i) CBDC primarily crowds out liquid savings deposits, (ii) CBDC substitutes for deposit-backed digital payments but not cash, and (iii) adoption is concentrated among households that tend to be wealthier, more sophisticated, and hold uninsured deposit balance.

In this simple framework, households allocate wealth across four instruments: savings deposits, term deposits, cash, and CBDC wallets. Savings deposits are liquid, interest-bearing claims on commercial banks and are typically linked to digital payment rails such as UPI. Term deposits offer higher returns but are illiquid within the period. Cash is default-free and anonymous but is physical and less convenient for remote or larger-value transactions. CBDC is a non-interest-bearing, default-free, digital claim on the central bank that can be used directly for person-to-person and person-to-merchant payments through a wallet app integrated with existing QR-code infrastructure.

4.1 Environment

Each household i starts the period with liquid wealth $W_i > 0$ and chooses nonnegative holdings

$$D_i^S \geq 0, \quad D_i^T \geq 0, \quad H_i \geq 0, \quad C_i \geq 0,$$

representing savings deposits, term deposits, cash, and CBDC, respectively, subject to the budget constraint

$$D_i^S + D_i^T + H_i + C_i = W_i. \tag{1}$$

Returns and deposit insurance. Savings deposits earn gross return $1 + r^S$ and are fully liquid. Term deposits earn gross return $1 + r^T$, with $r^T > r^S$, and are illiquid within the period. CBDC and cash are liquid and pay gross return 1.

Deposit insurance applies jointly to the sum of savings and term deposits. Let total deposits be

$$D_i^{\text{tot}} = D_i^S + D_i^T.$$

An insurance limit $\bar{D} > 0$ covers deposits up to $\bar{D}$ per household. Balances above $\bar{D}$ are uninsured and subject to expected loss. Let $\delta \in (0, 1)$ denote the expected loss rate per rupee of uninsured deposits.

The expected monetary payoff from deposits is

$$\Pi(D_i^S, D_i^T) = (1 + r^S)\ \min\{D_i^{\text{tot}}, \bar{D}\} + (1 + r^S)(1 - \delta)\ \max\{D_i^{\text{tot}} - \bar{D}, 0\}. \tag{2}$$

Replacing one rupee of uninsured deposits with CBDC saves expected loss $\delta(1 + r^S)$ until deposits fall back to the insurance limit.

In principle, households may also hold illiquid term deposits D^T that earn return r^T. We assume that these long-term positions are chosen before CBDC is introduced and remain fixed over the period in which households decide whether to adopt CBDC and how to reallocate their liquid balances. As a result, D^T and r^T do not appear in the expected monetary payoff in equation (2), and all portfolio adjustment in response to CBDC operates through liquid savings deposits, cash, and CBDC. This assumption fits our empirical setting, where the digital rupee is non-interest-bearing and therefore competes mainly on liquidity and payment convenience rather than as a high-yield savings instrument, and more broadly any environment in which the CBDC remuneration rate is well below the term-deposit rate r^T so that incentives to unwind illiquid positions are limited. In Appendix E.7, we relax this assumption and allow households to choose term deposits and liquid assets jointly in a simple intertemporal portfolio problem; the extension shows that when CBDC remains low-yield relative to r^T and early withdrawal from term deposits is sufficiently costly, the optimal response to CBDC is to adjust only liquid savings, cash, and CBDC, so treating D^T as predetermined in the baseline model provides a reasonably accurate reduced-form description of portfolio behavior over the horizon we study.

Digital payment technologies. Household i faces a flow of digital payment needs with intensity $\mu_i \geq 0$. These payments can be executed either via UPI, funded from savings deposits, or via CBDC, funded from the CBDC wallet. UPI is available if $D_i^S > 0$ and yields convenience γ^{UPI} per unit of digital payment. CBDC is available if $C_i > 0$ and yields convenience γ^{CBDC} per unit plus an additional safety/segmentation benefit from holding CBDC balances. Specifically, if the household adopts CBDC and holds $C_i > 0$, it receives a benefit $\phi_i S(C_i)$, where $\phi_i \geq 0$ captures the household-specific valuation of segmentation and $S : \mathbb{R}_+ \to \mathbb{R}_+$ is increasing and concave. Adopting and using CBDC entails a fixed cost $k_i \geq 0$.

If the household does not adopt CBDC, all digital payments run through UPI and receive a utility

$$U_i^{\text{no CBDC}} = \mu_i \gamma^{\text{UPI}}. \tag{3}$$

If the household adopts CBDC and uses the wallet as its main digital channel, then it receives a utility

$$U_i^{\text{CBDC}} = \mu_i \gamma^{\text{CBDC}} + \phi_i S(C_i) - k_i. \tag{4}$$

Total utility. Total utility for household i is

$$U_i = M_i(D_i^S, D_i^T, H_i, C_i) + \max\{U_i^{\text{no CBDC}}, U_i^{\text{CBDC}}\}, \tag{5}$$

where expected monetary payoffs are

$$M_i(D_i^S, D_i^T, H_i, C_i) = \Pi(D_i^S, D_i^T) + H_i + C_i. \tag{6}$$

We solve the household's problem in two steps. For a given adoption choice, the household chooses (D_i^S, D_i^T, H_i, C_i) to maximize expected monetary payoffs (6) subject to the budget constraint (1). It then compares the value of adopting and not adopting and chooses the higher-utility option.

4.2 Deposits: Savings Versus Term Deposits

Our first empirical result is that CBDC introduction reduces deposits primarily along the savings margin, with little change in fixed and recurring deposits. We now show how this arises in the model. To simplify notation, we suppress the index i in this subsection. Consider first the portfolio problem without CBDC ($C = 0$). Term deposits strictly dominate excess liquid balances for non-transactional wealth because they offer a higher return and illiquidity is not costly once liquid needs are met.

Lemma 4.1 (Baseline portfolio without CBDC)**.** *In the absence of CBDC, there exists an optimal allocation (D_0^S, D_0^T, H_0) such that: (i) for a given optimal level of liquid balances $L_0 = D_0^S + H_0$, any additional wealth is allocated to term deposits D_0^T, and (ii) if L_0 is uniquely optimal, then D_0^T is weakly increasing in wealth W and strictly positive for W sufficiently large.*

Assumption 1. *We impose that what matters for non-payment utility is total liquidity, not the exact split among savings, cash, and CBDC. Specifically, conditional on a given level of total liquid balances $L = D^S + H + C$, the household is indifferent over different splits of L across (D^S, H, C) for purposes other than digital payments.*[4]

[4]This assumption restricts the *non-payment* part of utility to depend on liquid resources only through their sum L (together with consumption), rather than on the full vector (D^S, H, C). In particular, holding fixed consumption and total liquid balances L, reallocating one rupee between savings deposits, cash, and CBDC does not change non-payment utility. At the same time, we do not assume that the three instruments are fully equivalent. First, they can differ in returns and safety in the budget constraint (for example, savings deposits earn interest and carry bank default risk, whereas CBDC does not bear interest but is a direct liability of the central bank). Second, they enter the *payments* component of utility differently, because CBDC and UPI facilitate digital transactions while cash is used for physical payments. Thus, composition affects behavior only through these explicitly modeled interest, safety, and payment-convenience channels, and we abstract from any additional intrinsic taste for holding a particular liquid instrument (such as a preference for cash for informality or for CBDC for peace of mind). This simplification keeps the portfolio problem focused on the trade-offs we focus on in the empirical section of the paper, rather than on unobserved heterogeneity in tastes over specific asset types.

Let (D_0^S, D_0^T, H_0) be an optimal portfolio in the no-CBDC case. Once CBDC is available, let (D_1^S, D_1^T, H_1, C_1) denote an optimal portfolio if the household adopts CBDC and chooses $C_1 > 0$. Henceforth, we will index by 0 in the no-CBDC case, and index by 1 in the CBDC case.

Assumption 2. *For every household i, $\phi_i S'(D_0^S + H_0) \leq \min\{r^T,\ (1 + r^T)(1 - \delta) - 1\}$.* [5]

Proposition 4.1 (CBDC funded from savings, not term deposits)**.** *Suppose (i) $r^T > r^S$ and term deposits are illiquid within the period, (ii) CBDC is liquid and pays gross return* 1*, (iii) Assumption 1 holds, and (iv) Assumption 2 holds. Then there exists an optimal portfolio with CBDC such that $D_1^T = D_0^T$ and*

$$D_1^S + H_1 + C_1 = D_0^S + H_0.$$

In particular, increases in C_1 relative to $C_0 = 0$ are financed by reductions in savings deposits or cash, rather than by reductions in term deposits.

The proof, given in Appendix E.1, shows that any portfolio where CBDC is funded by cutting term deposits can be dominated by increasing term deposits back to their baseline level and reducing liquid assets while keeping total liquidity fixed. Proposition 4.1 implies that CBDC disintermediates banks primarily along the savings margin. We test this implication in Section 6 using district- and individual-level data on deposits and credit.

4.3 Digital payments: Substitution between CBDC and UPI

Our second set of theoretical predictions is that CBDC is associated with a decline in usage of pre-existing digital payments and an insignificant decline in cash. In the model, this pattern follows from the close substitutability between CBDC and UPI as digital payment technologies and from the limited substitutability between cash and digital channels for certain transactions.

For a CBDC adopter with $C_i > 0$, let x_i^{UPI} and x_i^{CBDC} denote the fractions of the household's digital payment needs that are routed through UPI and CBDC, respectively, with

$$x_i^{\text{UPI}} + x_i^{\text{CBDC}} = 1.$$

For each transaction, the household compares per-unit net utility across channels. Institutional details suggest that $\gamma^{\text{CBDC}} \approx \gamma^{\text{UPI}}$ as the CBDC app offers similar convenience to UPI. Moreover, CBDC

[5] Assumption 2 says that a household whose entire liquid position sat in the CBDC wallet would value the marginal rupee of safety and segmentation at less than the term premium it gives up. The value of safety is concentrated in the first rupees moved. A household concerned about the safety of its bank holds that concern over a fixed amount of exposure, and once enough has been shifted into central bank money to cover it, further rupees buy protection against a risk that has already been neutralized. Payment convenience is bounded in the same way, since a wallet need only be large enough to fund the transactions a household actually makes. The assumption therefore requires only that safety and convenience be satiable in this sense, so that a household deep enough into its wallet no longer values the next rupee more than the interest forgone by not placing it in a term deposit.

provides additional segmentation benefits that are increasing in C_i. Therefore, CBDC becomes more attractive as the wallet balance and ϕ_i increase.

Lemma 4.2 (Payment-channel choice)**.** *For a CBDC adopter with $C_i > 0$ and $\gamma^{CBDC} \approx \gamma^{UPI}$, the optimal share x_i^{CBDC} is weakly increasing in C_i and in ϕ_i, and weakly decreasing in k_i. For sufficiently large C_i and ϕ_i, we have $x_i^{CBDC} \to 1$.*

We assume that a subset of transactions is intrinsically cash-intensive and remains in cash even after CBDC adoption.

Assumption 3. *There is a subset of transactions for which digital channels are poor substitutes for cash, so these transactions continue to be executed in cash regardless of CBDC adoption.*

Proposition 4.2 (Substitution from UPI to CBDC, limited effect on cash)**.** *Suppose (i) adopters fund CBDC from savings as in Proposition 4.1, (ii) for adopters, x_i^{CBDC} increases with C_i as in Lemma 4.2, and (iii) Assumption 3 holds. Then as CBDC adoption and wallet balances increase, aggregate UPI volume and transaction counts fall sharply, while aggregate cash usage falls by much less.*

Section 7 tests Proposition 4.2 by examining how CBDC affects UPI and cash transactions.

4.4 Adoption: Role of total deposits and sophistication

Our third set of results concerns who adopts CBDC. The model predicts that adoption hazards are much higher among customers with uninsured deposit balances, higher income and education. In the model, this pattern arises because the safety and segmentation benefits of CBDC are increasing in total uninsured deposits and because adoption costs are lower for more financially sophisticated households.

Let $V^0(W_i)$ denote maximal utility when household i does not adopt CBDC and $V^1(W_i)$ maximal utility when it adopts. Adoption is optimal if $V^1(W_i) \geq V^0(W_i)$. Conditional on adoption, let $C_i^*(W_i)$ be the optimal CBDC balance and $(D_{1i}^{S,*}, D_{1i}^{T,*}, H_{1i}^*)$ the associated deposits and cash.

Consider a household for which, in the no-CBDC benchmark, total deposits $D_{0i}^{\text{tot}} = D_{0i}^S + D_{0i}^T$ exceed the insurance limit $\bar{D}$. Replacing one rupee of uninsured deposits with CBDC yields an expected gain of size $\delta(1 + r^S)$, lowering the effective interest cost of CBDC relative to deposits. Let $\tilde{r}_i(W_i) \in [0, r^S]$ denote this effective net cost per rupee of CBDC for household i.

Lemma 4.3 (Adoption condition)**.** *For a given household (ϕ_i, μ_i, k_i) with $D_{0i}^{tot} > \bar{D}$, there exists an optimal CBDC balance $C_i^*(W_i)$ that solves*

$$\max_{0 \leq C \leq W_i} \{\phi_i S(C) - \tilde{r}_i(W_i)\, C\}\,.$$

Adoption is optimal if and only if

$$\phi_i S(C_i^*(W_i)) + \mu_i \left(\gamma^{CBDC} - \gamma^{UPI}\right) \geq \tilde{r}_i(W_i) C_i^*(W_i) + k_i. \tag{7}$$

Households with larger wealth and higher total deposits are more likely to have $D_{0i}^{\text{tot}} > \bar{D}$ and therefore face a lower effective cost $\tilde{r}_i(W_i)$ of holding CBDC instead of deposits.

Proposition 4.3 (Total deposits and adoption)**.** *Consider two households i and j with the same (ϕ, μ, k) but different wealth, $W_i > W_j$. Suppose that in the no-CBDC benchmark, total deposits satisfy $D_{0j}^{tot} \leq \bar{D}$ (fully insured) for household j and $D_{0i}^{tot} > \bar{D}$ (partially uninsured) for household i. Then:*

(i) If it is optimal for j to adopt CBDC, it is strictly optimal for i to adopt CBDC.

(ii) There exist parameter values $(\phi, \mu, k, r^S, \delta)$ such that i adopts CBDC while j does not.

Proposition 4.3 formalizes the idea that households with greater uninsured exposure have stronger incentives to adopt CBDC, even though the funds they move into CBDC are drawn from savings rather than term deposits. Heterogeneity in ϕ_i and k_i can be interpreted as capturing financial and digital sophistication, so that more sophisticated households are more likely to satisfy the adoption condition (7). Section 8 studies these adoption patterns using individual-level micro data that correlate CBDC adoption with these characteristics.

4.5 Discussion

Taken together, Propositions 4.1, 4.2, and 4.3 deliver three central implications that we directly test in the data. First, CBDC crowds out liquid savings deposits rather than term deposits. Second, CBDC competes primarily with deposit-backed digital payments via UPI and only weakly with cash. Third, adoption is concentrated among households with uninsured deposit balances and higher financial sophistication. The extension in Appendix E.4 makes this adoption pattern more explicit by introducing heterogeneity in households' valuation of safety and in their technology adoption costs: holding income and deposits fixed, CBDC take-up is higher for households that place greater value on the segmentation benefits of CBDC and face lower adoption costs, helping to explain why individuals with sophisticated education and stronger digital engagement adopt earlier. Overall, the model provides a unified interpretation of the effects on deposits, digital payments, and adoption that we document in Sections 6, 7, and 8.

We also use the framework to study how alternative CBDC designs would modify these effects. The design extension in Appendix E.5 shows that paying interest on CBDC materially strengthens adoption incentives and amplifies substitution away from bank deposits, potentially affecting term deposits as well as savings when the CBDC rate is sufficiently high. In this sense, our empirical setting

with a zero-interest CBDC can be interpreted as a lower bound on the potential impact of CBDC on both deposits and digital payments.

Finally, a key difference between our baseline model and the actual design of the digital rupee is that the Reserve Bank of India caps individual CBDC balances at ₹100,000. To capture this feature, the extension in Appendix E.6 introduces an exogenous holding limit on the CBDC wallet and studies how it alters portfolio and payment choices. The cap does not change which households find CBDC attractive. Adoption still concentrates among wealthier, more sophisticated depositors. However, it truncates their CBDC balances and mechanically limits how much they can draw down savings deposits and re-route digital payments from UPI to CBDC. In this sense, the holding cap acts as a quantitative brake on deposit disintermediation and payment substitution: it compresses the magnitude of the effects we document, without overturning the qualitative patterns of who adopts CBDC and which margins of the balance sheet and payment behavior adjust. More broadly, stock-based limits on CBDC holdings and related flow-based limits on transaction volumes are under active consideration in several jurisdictions (e.g. Kosanović, 2025; European Central Bank, 2020). These tools are viewed as a way to manage the impact of retail CBDC on bank intermediation and financial stability (Keister and Sanches, 2023; Carapella et al., 2024). The implications of this extension therefore speak not only to the Indian case, but also to ongoing international policy discussions about capped or otherwise restricted CBDC designs.

5 Empirical Strategy

We now describe the empirical strategies used to test the model's predictions. The empirical analysis combines two staggered roll-out designs: district-level variation in CBDC eligibility across lead-bank districts and within-bank variation in the timing of CBDC marketing campaigns across districts. We estimate the effect of CBDC introduction on aggregate bank deposits and credit, as well as on individual-level deposit balances. We use a difference-in-differences (DiD) design that exploits the staggered roll-out of the CBDC across units.

Recent work in the DiD literature shows that the traditional two-way fixed effects estimator can be biased when treatment effects vary across cohorts or over time. In our setting, the staggered roll-out of the CBDC makes such heterogeneity particularly likely. First, early-treated cohorts may systematically differ from later-treated cohorts, which can lead them to respond differently to CBDC introduction. Second, the CBDC ecosystem evolved during the roll-out, with changes in features, merchant acceptance, and public awareness, which could again induce some heterogeneity in cohorts treated at different times. Third, the macroeconomic and policy environment at the time of adoption may differ across cohorts, meaning that the same CBDC intervention can have different effects depending on concurrent business-cycle conditions and regulatory changes.

We address these concerns by using the estimator proposed by Callaway and Sant'Anna (2021) to construct event-time DiD estimates. This estimator allows for arbitrary treatment effect heterogeneity across cohorts and over time. It also makes the choice of control group explicit (Roth et al., 2023). We complement these results with estimates based on Borusyak, Jaravel, and Spiess (2024) and de Chaisemartin, D'Haultfœuille, and Vazquez-Bare (2024) to assess robustness to alternative estimators.

Following Callaway and Sant'Anna (2021), we estimate group-time average treatment effects. Let $G = g$ denote the cohort of units first treated in period g. The group-time average treatment effect is defined as $ATT(g,t) = \mathbb{E}[Y_t(g) - Y_t(0) \mid G = g]$. Here, $Y_t(g)$ denotes the potential outcome at time t for units first treated in period g, and $Y_t(0)$ denotes the potential outcome without treatment. We aggregate these group-time effects to obtain event-time estimates that summarize dynamic treatment effects relative to the time of treatment. We also report three aggregate measures. The simple aggregation computes a weighted average of all group-time treatment effects. The group-specific aggregation averages treatment effects within each cohort and then across cohorts. The event-time aggregation averages treatment effects across periods with the same exposure length.

The key identification assumption is the parallel trends assumption. In the absence of treatment, outcomes for treated and control units would have followed parallel trends within each cohort. This assumption allows for systematic differences in levels across cohorts, as long as these differences are time-invariant or evolve similarly in the absence of treatment. In particular, treatment does not need to be randomly assigned across units. All our specifications include unit and time fixed effects. Unit fixed effects absorb time-invariant differences across units, such as baseline levels of financial development or banking access, while time fixed effects absorb aggregate shocks that affect all units simultaneously. The identification therefore comes from differences in the timing of treatment across units. While this assumption is inherently untestable, we conduct pre-trends assessments across all specifications to support its validity.

5.1 Aggregate Bank Deposits and Credit

We first implement this framework at the district level by exploiting variation in the timing of CBDC authorization across lead banks and the districts they serve. The unit of observation is district d in quarter t. We define treatment timing g as the quarter in which the lead bank serving district d receives CBDC authorization. For each district-quarter, we examine the effect of the introduction of CBDC on the natural logarithm of total retail deposits and total credit.

This design maps bank-level treatment to district-level outcomes through the lead bank structure. Each district is assigned a lead bank that plays a central role in local banking activity. When a lead bank receives CBDC authorization, districts linked to that bank become exposed to the treatment.

We assume no anticipation, i.e., districts do not systematically adjust their deposit or credit

behavior prior to CBDC authorization in anticipation of CBDC introduction. This assumption is plausible given that authorization decisions are made at the bank level based on internal technical readiness. These bank-level decisions are filtered down to districts based on their lead banks. Therefore, districts are unlikely to observe or predict the exact timing of these decisions.

We use districts whose lead banks are not yet authorized as the control group and include district and time (quarter) fixed effects in our analysis. Therefore, the identification of the effect relies on the assumption that the timing of CBDC authorization across lead banks is not correlated with differential trends in district-level outcomes. This assumption is reasonable because lead bank assignments were fixed decades before CBDC was conceived and are driven by administrative considerations unrelated to recent deposit trends or the adoption of digital payments at the district level. We also present pre-trend assessment for all analyses to provide further supporting evidence for our identifying assumption.

In our district level setting, a natural concern is that some comparison districts may still experience CBDC penetration even when their designated lead bank has not yet been authorized. As a result, the binary treatment indicator based on lead bank authorization is an imperfect proxy for the underlying intensity of CBDC usage. Households in these districts can open CBDC wallets through non-lead banks that maintain branches in the district. Therefore, districts that are treated later can begin to accumulate CBDC exposure through banks that are authorized earlier but are not the lead bank, which implies that control districts are not literally CBDC free.

Appendix F formalizes this concern in a simple district level setting where CBDC exposure is allowed to vary continuously across treated and comparison districts, and where outcomes respond to the intensity of CBDC penetration rather than just to a binary treatment indicator. The key implication of this framework is that our estimated empirical object, based on lead bank authorization, is an intention-to-treat estimate under partial CBDC exposure and is therefore naturally interpreted as a lower bound in magnitude on the effect of moving from zero CBDC penetration to the treated districts' level of CBDC penetration, under the assumptions of our framework. We then use this framework to derive a transparent mapping from our reported intention-to-treat estimate to an implied *zero-to-treated* average effect of introducing CBDC, from no penetration to the penetration level observed in treated districts. Implementing this mapping requires combining our estimate with data on the geographic footprint of banks' deposit and credit activity across districts, as detailed in Appendix F.

5.2 Individual-Level Deposit Balances

We next implement a more granular version of the design at the individual level by exploiting variation in the timing of CBDC marketing campaigns across districts. The unit of observation is individual i in month t. We define treatment timing g as the month in which individual i's district receives the marketing intervention. For each individual-month, we examine savings, fixed deposit, recurring deposit, and total deposit balances.

This design captures variation in exposure to CBDC promotion and information at the individual level. Unlike the district-level analysis, which relies on bank-level authorization, this approach uses within-bank and within-district timing of marketing intensity. As a result, the design isolates how individuals adjust their deposit choices in response to changes in awareness and salience of CBDC, holding constant the underlying banking environment.

This individual-level design offers two key advantages compared to the district-level analysis. First, individual fixed effects allow us to control for unobserved heterogeneity at a much finer level, including baseline wealth, preferences, and banking behavior. This reduces concerns about compositional differences across treated and control units. Second, the use of monthly data allows us to trace treatment dynamics at a higher frequency. The key shortcoming of this analysis is that we do not observe credit at this unit of analysis.

We assume no anticipation of marketing exposure, i.e., individuals do not systematically adjust deposit balances prior to the arrival of marketing campaigns. This assumption is plausible because campaign timing is determined centrally by the RBI and is not communicated in advance at a granular level, which limits individuals' ability to adjust behavior prior to treatment.

Our estimation employs individual and time (month) fixed effects. Identification therefore relies on the assumption that the timing of marketing campaigns is not correlated with differential trends in individual deposit outcomes. Put differently, in the absence of the marketing campaign, early- and late-treated individuals would have followed parallel trends in deposit balances. This assumption is supported by two features of the setting. First, campaign timing follows a centralized roll-out strategy rather than local economic conditions or individual behavior. Second, all individuals in the sample are served by the same bank, which removes cross-bank differences in technology, product design, and implementation. We present pre-trend assessment to provide evidence in support of the identifying assumption.

6 CBDC and Bank Deposit Disintermediation

A central question in the theoretical literature on CBDC is whether its introduction reduces commercial banks' deposit base and, in turn, their capacity to extend credit. Existing models show, under a range of assumptions, that even a non-interest-bearing retail CBDC can lead to deposit disintermediation and, in some cases, a contraction in bank lending (see Whited, Wu, and Xiao, 2023; Chang et al., 2023; Gross and Schiller, 2021; Federal Reserve Board of Governors, 2022; Fernández-Villaverde et al., 2021, among others). The conceptual framework in Section 4 captures one central channel emphasized in this work by modeling CBDC as a safe, liquid instrument that competes primarily with households' savings deposits rather than with illiquid term deposits and by characterizing how this substitution can affect banks' funding conditions.

In this section, we bring these ideas to the data at two complementary levels of aggregation and quantify the associated changes in bank credit. At the district level, we use the staggered roll-out of CBDC access across lead bank districts to estimate the effect of CBDC introduction on aggregate deposits and credit. This analysis provides direct evidence on whether CBDC erodes the funding base of commercial banks and how much this translates into a reduction in lending. At the individual level, we exploit within-bank, within-district variation in CBDC marketing campaigns to study how households adjust the composition of their deposits. We show that the decline in deposits is concentrated in liquid savings balances rather than in fixed and recurring deposits, which is consistent with the mechanism in Section 4 that links CBDC adoption to a reallocation of transaction-related balances.

6.1 Effect of CBDC on Aggregate Bank Deposits and Credit

We first present district-level evidence on the effect of CBDC introduction on aggregate retail deposits and credit. We use the staggered authorization of lead banks as plausibly exogenous variation in the timing of CBDC access across districts, following the empirical design in Section 5.1. Figure 4 reports event-time estimates based on the Callaway and Sant'Anna (2021) estimator. The sample consists of a balanced panel of districts from 2018 Q1 to 2025 Q3. Panel (a) uses the logarithm of total retail deposits as the outcome, and panel (b) uses the logarithm of total credit.

The estimates in Figure 4 yield two key takeaways. First, the event-time coefficients do not show differential pre-trends in deposits or credit prior to CBDC eligibility, which provides suggestive support for the parallel trends assumption. Second, following eligibility, deposits decline on impact and remain persistently below pre-treatment levels, whereas credit also falls on impact but the decline is short-lived and quickly reverses.

The average treatment effect across post-treatment periods implies a 2.7% decline in total retail deposits ($ATT = -0.0273$), which is statistically significant at the 1% level. At the median level of deposits in 2022 Q3, this effect corresponds to a reduction of approximately ₹1.8 billion per district. The corresponding estimate for credit implies a 2% decline ($ATT = -0.0201$). This effect, however, appears to revert toward zero after six to seven quarters.

We view the effect of CBDC on retail deposits as a moderate thinning of the retail funding base rather than a run-like episode. To put this magnitude in context, we consider a simple adoption scenario. Suppose a district has a population of about 2 million people. Following the individual sample (Table 2), suppose that 4.4% of its residents adopt CBDC and move balances of ₹22,968.[6] This corresponds to 88,000 individuals, who are a small share of the district, each shifting ₹22,968, for a total of about ₹2 billion. At the median district in 2022 Q3, retail deposits are approximately

[6] In our data, the average monthly CBDC volume conditional on CBDC adoption is ₹22,968. We use this as a proxy for the average balance an adopter funds out of deposits, since we observe CBDC transaction flows rather than wallet balances.

₹67.5 billion, so this back-of-the-envelope calculation suggests that a modest adoption rate of 4.4% among residents at observed wallet sizes would reduce the local deposit base by roughly 3% of its pre-treatment level. This rough magnitude is close to, and slightly larger than, our estimated 2.7% decline in deposits, which suggests that the district-level effect we find is well within the range implied by observed adoption rates and wallet sizes among higher-balance households.

The dynamic pattern of credit is informative about banks' adjustment margins. Credit initially co-moves with deposits but partially recovers within roughly six to seven quarters, even though the decline in deposits appears persistent. This pattern suggests that banks initially face tighter funding conditions and reduce lending in proportion to the deposit outflows, but subsequently substitute toward alternative funding sources or adjust their balance sheets to support lending in the medium run. In the short run, the fall in deposits increases the marginal cost of funds and leads to a reduction in credit that mirrors the deposit decline. As the horizon extends, banks can respond by raising deposit rates, tapping wholesale or interbank markets, or reallocating liquid assets, which attenuates the effect on lending and generates the partial recovery in credit.

We want the readers to note two key points while interpreting the presented estimates. First, as discussed earlier and in Appendix F, the estimated effects on deposits and credit should be interpreted as lower bound intention-to-treat effects under partial CBDC exposure, given the assumptions made in our framework. While we are comfortable making this interpretation for the estimated effects on deposits, the limited effects on credit might be driven by the same partial-exposure attenuation bias. In principle, partial CBDC penetration in comparison districts could shrink the estimated difference between treated and comparison districts and thereby reduce the apparent impact of CBDC eligibility on credit. In our setting, however, this concern primarily affects the level of the estimated effect and does not generate the dynamic pattern that we document. The credit event study shows an initial decline in credit followed by a gradual recovery toward zero, which is more naturally interpreted as banks adjusting their funding mix and lending policies over time rather than as a mechanical consequence of partial CBDC exposure.[7]

Second, since the Reserve Bank of India caps individual CBDC balances at ₹100,000, our estimates are likely to represent a lower bound in magnitude on the effect of CBDC on deposits and credit relative to settings in which CBDC holdings are unconstrained. This distinction matters for external validity. The model extension in Appendix E.6 introduces an explicit holding limit into the portfolio problem and shows that the cap does not change which households find CBDC

[7]For partial CBDC exposure to generate a time varying attenuation of the credit effect, comparison districts must converge toward treated districts in CBDC penetration in a way that systematically changes the treated-versus-comparison contrast over event time. This convergence requires the ratio of CBDC usage in comparison versus treated districts to vary sharply across post-treatment quarters, a pattern that would reflect rapid catch-up of CBDC usage in comparison districts. While we cannot directly observe district-level CBDC penetration, the institutional design of the rollout, with authorization concentrated in lead banks that dominate local intermediation, makes such rapid convergence less plausible, so we view the partial exposure bias as a source of level attenuation rather than as an explanation for the time varying credit pattern.

attractive or which margins of their balance sheets adjust, but it truncates their desired CBDC balances and mechanically limits how much they can draw down savings deposits and re-route digital payments from UPI. In that extension, the holding cap therefore acts as a quantitative brake on deposit disintermediation and payment substitution: it scales down the magnitude of the effects we document while leaving their qualitative patterns intact.

The district-level results are robust to using alternative estimators. Appendix Figure C.3 presents long-difference estimates based on Callaway and Sant'Anna (2021), and Appendix Figure C.4 reports results using the estimators in Borusyak, Jaravel, and Spiess (2024) and de Chaisemartin, D'Haultfœuille, and Vazquez-Bare (2024). The results are similar across all three estimators, which indicates that our findings are not driven by the choice of difference-in-differences methodology.

The district-level results are also robust to relaxing the parallel trends assumption. Because early-authorized cohorts may differ from later cohorts along observable characteristics, we re-estimate the effects on deposits and credit under conditional parallel trends, which requires only that trends be parallel among districts with similar baseline characteristics. Appendix Figures C.1 and C.2 report event-time estimates based on Callaway and Sant'Anna (2021), conditioning on baseline night light intensity, deposits per capita, and credit per capita, both individually and jointly. Across all specifications, we find no evidence of differential pre-trends and recover effects on both deposits and credit that are similar to our unconditional estimates. Conditioning on baseline district characteristics therefore does not change our conclusions.

Overall, the district-level evidence is consistent with recent theoretical work that emphasizes the deposit channel as the primary way in which retail CBDC affects banks (see Infante et al., 2022, 2023, 2024, for reviews). A reduction in deposits raises the marginal cost of funds, but the impact on credit depends on how easily banks can substitute across funding sources and adjust portfolio choices. We observe a clear and persistent decline in deposits, while the effect on credit is more modest and fades over the medium term, which is consistent with a setting in which banks face initial funding pressure but retain some flexibility to adjust their liability structure.

6.2 Effect of CBDC on Individual-Level Deposit Balances

The district-level results help establish that CBDC introduction is associated with a decline in aggregate deposits and a transient decline in credit. These aggregate patterns, however, do not reveal which parts of households' balance sheets adjust. Aggregate data cannot distinguish between a reallocation out of liquid savings and a more general reduction in deposit funding, including term deposits. This distinction is important because the two margins have different implications for households and for banks.

At the household level, a shift out of liquid savings affects the resources that households use for day-to-day transactions and short-term shocks, whereas a shift out of term deposits affects long-term

savings and portfolio risk. The composition of the outflows also helps distinguish whether CBDC operates primarily as a transactions instrument or, despite paying no interest, as a savings vehicle. At the bank level, liquid deposits and term deposits play different roles in funding and liquidity management. Liquid deposits provide low-cost, on-demand funding, while term deposits are more expensive but more predictable over the contract horizon. Identifying which margin adjusts therefore helps interpret the implications of CBDC for households' financial resilience and for banks' funding structures, and it directly tests the prediction in Proposition 4.1 that CBDC displaces savings deposits rather than term deposits.

We use individual-level data to trace the aggregate decline in deposits to specific components of household balance sheets. We study four outcomes: savings balances, fixed deposit balances, recurring deposit balances, and total deposit balances. The specification includes individual fixed effects, so the estimates absorb all time-invariant heterogeneity in demographics, baseline financial behavior, and branch characteristics. The identifying variation comes from differences in the timing of the CBDC marketing campaign across districts for otherwise comparable customers within the same bank, as described in Section 5.2.

Figure 5 reports event-time estimates based on the Callaway and Sant'Anna (2021) estimator. Across all four outcomes, we do not observe significant pre-trends in the months before the intensive CBDC marketing campaign, which supports the parallel trends assumption. Following the campaign, we document a 5.6% decline in individual-level total deposit balances ($ATT = -0.0575$). Moreover, the decline in total deposit balances is concentrated in savings accounts, whereas we do not observe economically meaningful or statistically significant changes in illiquid high-yield deposits, such as fixed or recurring deposits.

Appendix C.2 presents robustness checks based on the long-difference approach of Callaway and Sant'Anna (2021) (Appendix Figure C.5) and on the alternative DiD estimators in Borusyak, Jaravel, and Spiess (2024) and de Chaisemartin, D'Haultfœuille, and Vazquez-Bare (2024) (Appendix Figure C.6). The results are qualitatively similar across these methods, which indicates that our findings are not sensitive to the choice of estimator.

These results are consistent with the framework in Section 4. In that framework, CBDC is a close substitute for low-yield, highly liquid savings deposits. Households therefore fund CBDC by shifting balances out of savings when the perceived safety and segmentation benefits of central bank money exceed the foregone interest. Illiquid, higher-yield term deposits are not the natural funding source. Proposition 4.1 therefore predicts that CBDC should compete primarily with liquid savings deposits rather than with term deposits.

The individual-level evidence matches this prediction. CBDC-related outflows arise mainly from savings accounts, whereas fixed and recurring deposits remain largely unchanged. This pattern

indicates that households adjust the liquid portion of their balance sheets that supports everyday transactions, rather than drawing down long-term savings, and it aligns with the view in Section 4 that CBDC and savings deposits provide similar transaction services while term deposits remain a distinct, long-horizon savings instrument.

7 CBDC as a Substitute for Digital & Cash Payments

Having documented that CBDC introduction is associated with a decline in liquid deposit balances, we next examine whether CBDC also substitutes for the payment methods through which those balances were previously used for transactions. The framework in Section 4 predicts that CBDC should compete more directly with deposit-backed digital payments than with cash (Proposition 4.2). We test this prediction by estimating its effect on UPI use and ATM withdrawals, using the latter as a proxy for cash use.[8] We combine the individual-level transaction data with the staggered geographic variation in CBDC marketing intensity at our sample bank and apply the same event-study difference-in-differences design as in Section 5.2 to test these predictions.

7.1 Effect of CBDC on Digital Transactions

We first examine the effect of CBDC on digital payments. We focus on UPI, which is the main digital rail in our setting. Figure 6 reports event-time estimates for aggregate UPI usage. We estimate an approximately 6.3% decline in UPI volume, and a 0.85% decline in the number of UPI transactions. We do not observe significant pre-trends in the months before the CBDC marketing campaign for any of these outcomes, which supports the parallel trends assumption. After the campaign, UPI usage declines.

We next interpret the size of these effects. A 0.85% decline in the number of UPI transactions is much smaller than the change in volume, indicating that most customers continue to use UPI and that many still make a substantial number of UPI payments after the CBDC campaign. The roughly 6.3% decrease in UPI volume is therefore best interpreted as an intensive-margin response in which a subset of users shifts part of their digital payment flows to CBDC. The fact that volume falls much more than transaction counts suggests that medium- and higher-value payments are migrating to CBDC, while many small-ticket payments remain on UPI. Because these estimates are local to customers of our sample bank in districts exposed to targeted CBDC promotion, and because baseline UPI penetration and volumes are very high, we interpret the effects as indicating a noticeable rebalancing of digital transactions among early adopters rather than a broad retreat from UPI in the wider economy.

This interpretation is also in line with the adopter profile we document in Section 8, where early CBDC users are disproportionately high-deposit, high-income, and digitally sophisticated households

[8] In India, the dominant digital payment method is UPI, and UPI transactions are primarily funded by liquid deposits such as savings accounts.

who plausibly account for a large share of UPI volume even though they represent a modest share of all potential users. These adopter characteristics help rationalize the magnitude and pattern of the UPI effects. If a relatively small group of heavy and affluent users moves part of its payment activity to CBDC, UPI volume can fall by a sizeable amount even when most customers continue to use UPI and many still make a large number of UPI payments. In this setting, a 0.85% decline in the number of transactions is consistent with the continued broad reliance on UPI by the average depositor. The much larger decline in UPI volume can then be understood as suggesting a reallocation of medium- and higher-value transactions by this concentrated group of intensive users. Seen through this lens, the UPI magnitudes indicate a pronounced adjustment in the behavior of early, high-intensity adopters rather than a general collapse of UPI usage across the wider customer base.

Appendix Figures C.7 and C.8 decompose the aggregate UPI decline into person-to-person (P2P) and person-to-merchant (P2M) transactions. Both channels decline with P2P falling nearly twice as much as P2M. We estimate an approximately 4.5% decline in P2P volume ($ATT = -0.0462$) and an approximately 2.5% decline in P2M volume ($ATT = -0.0253$). This pattern indicates that substitution away from UPI is broad-based across transaction types rather than concentrated in either interpersonal transfers or merchant payments.

7.2 Effect of CBDC on ATM Withdrawals

We next examine the effect of CBDC on cash usage, as proxied by ATM withdrawals. Figure 7 reports the corresponding event-time estimates. We find a statistically insignificant and an economically small decline in ATM withdrawal volume. This result stands in contrast with the result on UPI, suggesting that CBDC substitutes primarily for digital payments rather than for cash.

We assess robustness to alternative specifications in Appendix C.4. We re-estimate the effects using the long-difference approach of Callaway and Sant'Anna (2021) and the alternative difference-in-differences estimators in Borusyak, Jaravel, and Spiess (2024) and de Chaisemartin, D'Haultfœuille, and Vazquez-Bare (2024). We also replicate the UPI analysis using weekly instead of monthly outcomes. Across these specifications, we do not find evidence of pre-trends and we obtain qualitatively similar declines in UPI usage and ATM withdrawals. These tests indicate that our results are not sensitive to the choice of estimator or temporal aggregation.

Taken together, the evidence on digital payments and cash supports Proposition 4.2 and suggests that cash is a relatively weak substitute for CBDC. Cash is physical and anonymous and is used predominantly for small, face-to-face transactions where digital infrastructure is limited, so the main margin of substitution runs through deposit-backed digital payments rather than through currency.

8 Who Adopts CBDC?

So far, we have shown that CBDC introduction is associated with deposit disintermediation and substitution away from existing digital payments. We now turn to the question of which households adopt CBDC. We combine individual-level data on CBDC usage with survival analysis, which allows us to account for right-censoring when some individuals have not yet adopted by the end of our sample.[9] This analysis provides context for the results in Sections 6 and 7 and speaks to the model's prediction in Section 4 that adoption hazards are higher for wealthier, more sophisticated, and uninsured depositors.

Table 3 reports the estimates of the determinants of CBDC adoption based on a Cox proportional hazard model, expressed as hazard ratios. The dependent variable is the time until an individual first uses CBDC, measured in weeks since the CBDC marketing campaign begins in that individual's district. Columns (1)–(6) consider each characteristic separately: gender, age quartile, income quartile, prior UPI usage, sophisticated education, and deposit insurance threshold, respectively. These specifications show the raw association between each variable and the rate of CBDC adoption, without controlling for other individual attributes or location. Column (7) then includes all covariates jointly, so the hazard ratios capture the partial correlation of each characteristic with adoption, holding the others fixed but without fixed effects. Column (8) adds district fixed effects, and column (9) replaces them with ZIP code fixed effects. These last two specifications absorb time-invariant geographic heterogeneity in economic conditions and digital infrastructure, so the estimates reflect within-district or within-ZIP differences across individuals.

For completeness, Appendix Table D.1 reports the underlying Cox proportional hazards coefficients that correspond to the hazard ratios in Table 3. Appendix Table D.2 reports coefficients from a linear probability model of CBDC adoption estimated on the same sample. We also estimate a linear probability model of UPI usage using the same covariates, excluding prior UPI usage, to characterize which individuals use UPI during our sample period in Appendix Table D.3. Finally, Figure 8 presents Kaplan–Meier estimates of cumulative CBDC adoption probabilities across characteristic subgroups. The Kaplan–Meier curves show the cumulative probability of adoption over time for each subgroup without conditioning on other characteristics.

The magnitudes in the full specification with ZIP code fixed effects in column (9) of Table 3 indicate a pronounced adopter gradient. Male customers adopt CBDC at roughly 2.3 times the hazard rate of female customers, which means that in any given week men are more than twice as likely to adopt as otherwise similar women. Age effects are modest but suggest earlier adoption among middle-aged individuals and slower adoption among the oldest quartile once other characteristics are held constant. Income is strongly related to adoption: relative to the lowest income quartile, the top income quartile adopts at more than five times the hazard rate, even after controlling for

[9] Appendix Section D.1 provides a detailed discussion of the survival analysis methodology.

demographics, location, and other economic variables. Customers with deposit balances above the deposit insurance threshold adopt at roughly four times the hazard rate of those with deposit balances below the threshold.

Digital sophistication further differentiates adopters. Individuals with a sophisticated educational background adopt at nearly twice the hazard rate of those without such training, which suggests that technical or financial skills meaningfully reduce the cost of adopting a new digital instrument. Prior UPI usage is positively associated with adoption in the univariate column (4), but its hazard ratio falls below one once we control for income, deposits, education, and other covariates: in columns (7)–(9), prior UPI users adopt at only about 0.8 times the hazard rate of otherwise similar non-users. This shift indicates that the raw positive correlation between UPI usage and adoption is largely driven by the fact that UPI users are richer, hold more deposits, and are more educated; conditional on these factors, they adopt somewhat more slowly, which is consistent with weaker incremental benefits when a satisfactory digital payment option is already in place.

Overall, columns (1)-(6) document strong univariate links between CBDC adoption and gender, age, income, digital sophistication, deposits, and spending behavior, while columns (7)–(9) show that these patterns remain economically large once we control for other characteristics and for granular geographic heterogeneity. The combined evidence points to an adopter profile that is disproportionately male, higher-income, relatively sophisticated, and holds uninsured deposit balance. This profile is consistent with Proposition 4.3 in Section 4 and with the model extension in Appendix E.4, in which households with larger uninsured deposits and greater financial sophistication face stronger incentives to adopt CBDC, and it suggests that these households are the primary drivers of the aggregate disintermediation and payment-substitution effects documented earlier.

9 The Costs & Benefits of CBDC

The preceding sections document three core facts: the digital rupee reduces banks' liquid retail funding, the resulting contraction in credit is short-lived, and adoption is concentrated among higher-income, highly educated, and uninsured depositors. These findings raise a natural policy question: does a non-interest-bearing, capped retail CBDC increase or decrease aggregate welfare?

To answer this question, we extend the model in Section 4 along two key dimensions. First, we incorporate a cap on CBDC wallet holdings to reflect a key institutional feature of the e-₹. Second, we introduce a banking sector that captures how CBDC adoption affects bank funding and liquidity creation. We then adopt a sufficient-statistics approach in the spirit of Chetty (2009), which allows us to express the first-order welfare effect as a function of a set of elasticities that map directly to the reduced-form estimates in our setting. Combining these estimates with the model, we quantify the overall welfare impact of CBDC introduction and decompose the contribution of the underlying

channels.

9.1 A Minimal Bank Block

There is a representative bank sector that funds a stock of loans ℓ with three liabilities: retail deposits D, wholesale funding B, and equity, which we hold fixed within the period. The bank's within-period funding identity is

$$\ell = D + B + e, \tag{8}$$

where e represents the fixed equity net of required buffers. Retail deposits are the cheapest and stickiest funding source. The bank pays a deposit rate r^S, but enjoys a deposit franchise value, so the effective marginal cost of retail deposits is $r^S - \lambda$, where $\lambda > 0$ is the per-rupee deposit franchise spread.[10] Wholesale funding is more expensive at the margin, with an upward-sloping supply so that the total wholesale cost is $R(B) = \rho_0 B + \frac{1}{2}\psi B^2$, where $\psi \geq 0$ denotes the slope of the marginal wholesale funding cost. Loans earn r^ℓ and are subject to a convex origination/servicing cost $c(\ell)$ with $c'' > 0$. Bank profit is thus

$$\Pi^{\text{bank}} = r^\ell \ell - (r^S - \lambda)D - R(B) - c(\ell), \tag{9}$$

where the bank maximizes their objective over (ℓ, B) taking retail deposits D as determined on the household side subject to the bank's within-period funding identity (8).

Proposition 9.1 (Deposit-to-loan pass-through)**.** *When a CBDC shock withdraws $dD < 0$ of retail deposits, the bank replaces a fraction of the shortfall with wholesale funding and passes the rest through to lending:*

$$\frac{d\ell}{dD} = \frac{\psi}{\psi + c''(\ell)} \equiv \theta \in [0, 1), \quad \frac{dB}{dD} = -(1 - \theta). \tag{10}$$

The pass-through θ measures the tightness of the wholesale-substitution margin: it is the share of a deposit outflow that passes into a lending contraction rather than being absorbed by wholesale funding. Two limiting cases sharpen the interpretation. When wholesale funding is perfectly elastic ($\psi \to 0$), the bank replaces lost deposits at no additional marginal cost, so $\theta \to 0$ and deposits are irrelevant for credit. When wholesale funding is prohibitively costly at the margin ($\psi \to \infty$), the bank cannot substitute at all, so $\theta \to 1$ and the outflow passes one-for-one into reduced lending. The pass-through is therefore a property of banks' funding technology rather than of household preferences.

[10] The deposit franchise spread is the wedge between the value of a marginal retail deposit and its interest cost (Drechsler et al., 2026).

Corollary 9.1 (Declining pass-through)**.** *Suppose the slope of the marginal wholesale funding cost varies with the horizon over which the bank adjusts, with $\psi_{SR} > \psi_{LR} \geq 0$, and let $\theta_h = \frac{\psi_h}{\psi_h + c''(\ell)}$ denote the pass-through at horizon $h \in \{SR, LR\}$. Then $\theta_{SR} > \theta_{LR}$, and the share of a deposit outflow that is absorbed by wholesale substitution, $1 - \theta_h$, rises with the horizon.*

Corollary 9.1 reproduces the pattern we document in Section 6.1, where credit falls alongside deposits on impact and then reverts toward zero within six to seven quarters even though the deposit decline is persistent. The reason is that banks' funding margins open up at different horizons. Within a quarter, a bank facing an outflow has little recourse beyond drawing down liquid assets. Raising deposit rates to attract replacement balances, arranging wholesale and interbank borrowing, and rebalancing the securities portfolio all take longer to execute. Replacement funding is expensive at the margin in the short-run and cheaper once these channels open, hence $\psi_{SR} > \psi_{LR}$ and θ falls with the horizon.

9.2 Household Welfare with Holding Cap

On the household side, we retain the environment of Section 4 together with the holding cap $\bar{C}$ from Appendix E.6. An adopting household's indirect utility from the CBDC option, net of what it would have obtained on the deposit-plus-UPI bundle, is

$$\Delta v_i = \underbrace{\phi_i S(\hat{C}_i)}_{\text{safety/segmentation}} + \underbrace{\mu_i \left(\gamma^{\text{CBDC}} - \gamma^{\text{UPI}}\right)}_{\text{payment convenience}} - \underbrace{\tilde{r}_i(W_i)\hat{C}_i}_{\text{foregone interest/avoided loss}} - k_i \tag{11}$$

where $\hat{C}_i = \min\{C_i^*(W_i), \bar{C}\}$, and $\hat{C}_i, C_i^*(W_i), \tilde{r}_i(W_i)$ are defined as in Lemma 4.3 and Appendix E.6. By the envelope theorem, the marginal adopter (for whom $\Delta v_i = 0$) contributes no first-order welfare, so the aggregate private gain accrues to inframarginal adopters and can be written as

$$\Delta W^{\text{hh}} = \int_{\{i\,:\,\text{adopter}\}} \Delta v_i di = N a \overline{\Delta v}, \tag{12}$$

where N is the number of households, a is the adoption rate, and $\overline{\Delta v}$ is the average surplus among adopters.

9.3 Welfare Decomposition

Netting the private household gain against the change in bank surplus and the change in central-bank/fiscal surplus gives the aggregate welfare change. Let $dD = D_1 - D_0 < 0$ denote the equilibrium deposit outflow induced by the CBDC. The aggregate welfare change can then be decomposed into

the following six terms:

$$\Delta\mathcal{W} = \underbrace{Na\overline{\Delta v}}_{\text{(i) household surplus}} - \underbrace{\lambda|dD|}_{\text{(ii) lost deposit franchise}} - \underbrace{\tfrac{1}{2}\psi(1-\theta)^2|dD|^2}_{\text{(iii) wholesale-funding deadweight}} - \underbrace{\tfrac{1}{2}c''(\ell)\theta^2|dD|^2}_{\text{(iv) lending curvature loss}} - \underbrace{\eta\theta|dD|}_{\text{(v) credit-contraction externality}} + \underbrace{s^{\mathrm{CB}}|dD|}_{\text{(vi) seigniorage / fiscal}} \tag{13}$$

where $\eta \geq 0$ is the marginal social value of bank credit in excess of its private return, s^{CB} is the net social value per rupee of central-bank money issued, and $\bar{C}$ is the average CBDC balance per adopter. One can interpret η as a credit-wedge term capturing financial-accelerator or relationship-lending externalities, where $\eta = 0$ under the frictionless credit case.

Term (i) is the private value households place on the CBDC as derived in Section 9.2. Term (vi) is the fiscal counterpart of the same substitution, where s^{CB} is the net seigniorage margin the central bank earns on the $|dD|$ in commercial bank liabilities that are replaced by the CBDC.

Against these, term (ii) is the intermediation rent banks earn on cheap, sticky insured deposits and lose when those balances migrate, a cost the adopting household does not internalize.

Term (iii) is the additional resource cost of pushing banks onto a steeper wholesale funding curve. By Proposition 9.1, the bank closes a fraction $(1-\theta)$ of the funding gap with wholesale funding, so $|dB| = (1-\theta)|dD|$. Writing the total wholesale cost as $R(B) = \rho_0 B + \frac{1}{2}\psi B^2$, the additional cost of expanding from B_0 to $B_0 + |dB|$ is

$$R(B_0 + |dB|) - R(B_0) = \underbrace{R'(B_0)|dB|}_{\text{priced at the prevailing margin}} + \underbrace{\tfrac{1}{2}\psi|dB|^2}_{\text{convexity}}, \tag{14}$$

where the first component is offset by the resources the bank no longer devotes to deposit funding and the second is the incremental cost of moving up the supply schedule. Substituting $|dB| = (1-\theta)|dD|$ into the second term gives $\frac{1}{2}\psi(1-\theta)^2|dD|^2$, which is the wholesale-funding deadweight.

Term (iv) captures the second-order deadweight loss incurred by the banking sector as it contracts its loan portfolio by $\theta|dD|$. Governed by the curvature of the lending schedule $c''(\ell)$, this term reflects the increasing marginal frictions and operational costs associated with unwinding relationship loans or liquidating illiquid assets. Just as the bank faces a convex penalty for tapping wholesale funding (term (iii)), it faces a convex penalty for shrinking its credit book.

Finally, term (v) is the social value of the lending that does not occur when funding tightens, and it is the term policymakers have in mind when they worry that a retail CBDC will disintermediate credit.

9.4 Welfare Estimation

Equation (13) is useful because every term is the product of a quantity that our research design recovers and a price or wedge that can be calibrated or bounded from observable rates. None of the six terms requires us to identify the deep parameters of the household problem (ϕ_i, γ^{CBDC}, δ, or the distribution of k_i).

The district-level average treatment effect on total retail deposits is $ATT^{\text{dep}} = -0.0273$, which at the median district deposit stock of ₹67.5 billion in 2022 Q3 implies an outflow of $|dD| =$ ₹1.84 billion.

Proposition 9.1 shows that θ is the share of a deposit outflow that passes into reduced lending rather than being absorbed by wholesale substitution. Because our empirical ATTs are estimated from log outcomes, the ratio of the credit response to the deposit response yields an elasticity. To calculate the level pass-through required for our welfare formula, we must scale this ratio by the baseline loan-to-deposit ratio, ℓ_0/D_0. Using the Q3 2022 median credit (₹43.7 billion) and median deposits (₹67.5 billion), our baseline ratio is approximately 0.647. Therefore, the short-run level pass-through is identified as $\hat{\theta}_{SR} \approx \left(\frac{ATT^{\text{cred}}}{ATT^{\text{dep}}}\right) \times \frac{\ell_0}{D_0} = \left(\frac{-0.0201}{-0.0273}\right) \times 0.647 \approx 0.476$. The subsequent reversion of credit toward zero within six to seven quarters, even as deposits remain persistently depressed, implies that θ declines over time as banks scale wholesale funding, exactly the situation that Corollary 9.1 discusses.

The adoption rate in the individual sample is $a = 4.4\%$ (Table 2), which for a district of $N = 2$ million residents gives $Na \approx 88{,}000$ adopters. By matching the aggregate deposit outflow $|dD|$ to this adopter population, we calculate an implied average CBDC wallet balance of $\hat{C} = \frac{|dD|}{Na} \approx 20{,}909$ rupees.[11]

Revealed preference bounds the gross benefit from below: every adopter forgoes interest $\tilde{r}_i(W_i)\hat{C}_i$ and pays a fixed cost $k_i \geq 0$, meaning their valuation must at least cover the average opportunity cost of foregone interest, $r^S\hat{C}$.[12] This bound is likely to be slack for two reasons. The steep adoption gradients of Section 8 imply that adopters are drawn from well inside the margin rather than clustered at it: adopters are high-income, financially sophisticated individuals with uninsured deposits. We therefore set the gross benefit of adopting CBDC to ₹1,250, approximately 2 times the opportunity cost of foregone interest $\frac{1250}{0.03 \times 20{,}909} \approx 2$. This assumption allows average valuation to comfortably exceed the revealed-preference floor without imposing assumptions on the distribution of the gross benefit or cost among adopters. We thus assume that the average net surplus is $\overline{\Delta v} = 620$.[13]

We calibrate four key parameters that translate the quantity responses we estimate into welfare terms. First, we set the deposit franchise spread λ to 2 percentage points. This choice is informed by recent evidence in the deposit franchise literature that documents low single digit excess returns

[11] We validate this calculation empirically as the average monthly CBDC volume conditional on adopting CBDC is ₹22,968.
[12] The rate on savings deposits is observable and it is calibrated to 3% in our analysis.
[13] By construction, $\overline{\Delta v} = 623$, but we set it to 620 as a rounding choice.

on insured retail deposits relative to market funding, which we view as a conservative estimate of banks' per rupee franchise value on savings deposits (Drechsler et al., 2026). Second, we calibrate the slope of the wholesale funding marginal cost schedule ψ_{SR} so that $\psi_{SR}|dD| = 1.5$ percentage points. This means the marginal cost of wholesale funding at the CBDC-induced deposit change is about 150 basis points, consistent with a moderately convex funding curve as in Whited, Wu, and Xiao (2023). Third, we set the credit wedge η, the marginal social value of bank credit in excess of its private return, to 5 percent. This value is designed to be modest yet positive, in line with the magnitude of lending externalities used in financial accelerator and CBDC DSGE models where disruptions to bank credit generate noticeable but not dominant welfare losses (Assenmacher, Bitter, and Ristiniemi, 2023; Kumhof et al., 2023). Finally, we normalize the net seigniorage value per rupee of CBDC, s^{CB}, to 1 percentage point, treating a rupee of non interest bearing central bank money as yielding one rupee of fiscal surplus in present value terms after operating costs. This reduced form assumption follows standard seigniorage calculations and recent work on central bank balance sheets, which measure seigniorage as the spread between returns on central bank assets and the remuneration on their monetary liabilities.[14]

Finally, rather than assuming the lending friction c'', we use Proposition 9.1 to discipline our calibration. Given our empirical estimate of $\theta_{SR} = 0.476$, and an assumed wholesale funding marginal cost schedule $\psi_{SR}|dD| = 1.5$ percentage points, the model implies a lending curvature penalty of $c''(\ell)_{SR}|dD| \approx 1.65$ percentage points.

Under these inputs, we can compute the net welfare change following equation (13). Table 4 presents our calculations for each term. We estimate a short-run, district-level welfare loss of approximately ₹15 million, driven primarily by the initial contraction in bank credit. In the long run, however, when the credit effect reverses, we let $\theta_{LR} \to 0$. Following Proposition 9.1, we also assume that $\psi_{LR} \to 0$.[15] The credit reversal eliminates the associated deadweight losses and yields a net aggregate welfare gain of approximately ₹36.1 million. The composition is the economically meaningful output of the exercise. On the benefit side, household convenience and safety account for about 74.7% of gross gains and net seigniorage for the remaining 25.3%. On the cost side, the

[14] Appendix G.3 presents a sensitivity analysis to show that our conclusions are qualitatively robust to different assumptions on the four parameters $\left(\overline{\Delta v}, \lambda, \psi_{SR}, \eta\right)$.

[15] In the long run, we set $\theta_{LR} \to 0$ because the estimated effect of CBDC adoption on bank credit supply becomes statistically indistinguishable from zero. Thus, the initial contraction in lending reflects a transitional balance-sheet adjustment rather than a permanent reduction in the supply of credit. By Proposition 9.1, the disappearance of the credit-supply response also implies that $\psi_{LR} \to 0$. The term ψ captures the incremental marginal funding-cost wedge generated when a bank replaces CBDC-displaced deposits with wholesale or other non-deposit liabilities. In the short run, this replacement can be costly because banks must adjust their liability structure rapidly, access wholesale markets at the margin, and potentially bear a liquidity- or risk-related premium. As banks subsequently adjust their balance sheets, diversify funding sources, extend the maturity of their liabilities, and rebuild their desired funding mix, this transition-specific marginal wedge can dissipate. This interpretation is consistent with evidence that wholesale liabilities provide banks with greater quantity flexibility than retail deposits and are therefore used to accommodate changes in lending and balance-sheet needs (Dinger and Craig, 2015). The relevant long-run restriction is therefore not that deposits lose their franchise value. Rather, it is that the CBDC-induced displacement of deposits no longer creates an additional marginal funding distortion large enough to affect loan supply.

credit-contraction externality is the largest item in the short-run but vanishes in the long-run as credit recovers, leaving the lost deposit franchise as the only persistent cost, while the wholesale deadweight and the lending curvature losses remain strictly second-order throughout. Under this framework, the cost that dominates the policy debate over retail CBDC, a contraction in bank lending, is a temporary adjustment rather than a permanent tax. Ultimately, the long-run benefits of household convenience and seigniorage strictly outweigh the persistent cost of the lost deposit franchise, yielding a net aggregate welfare gain of ₹36.1 million.

The magnitude of the net welfare gain in the long-run is best understood relative to the disruption that produces it. Expressed per adopting household, the net aggregate welfare gain is approximately ₹410, which corresponds to about 0.26 percent of average annual income in our sample (₹159,000), and is equal to roughly two-thirds of the ₹627 of deposit interest a typical adopter forgoes each year by holding their average ₹20,909 non-interest-bearing CBDC balance. At the aggregate level, the long-run net welfare gain of ₹36.1 million represents approximately 2.0 percent of the ₹1.84 billion in deposits that migrate out of the banking system. Against the district's ₹67.5 billion deposit base, however, the aggregate gain is roughly 0.05 percent, which is what one would expect from a targeted policy that directly affects only a small share of bank balance sheets.

10 Conclusion

This paper studies how a live, large-scale retail central bank digital currency affects bank funding, digital payments, and household portfolios. We focus on India's retail digital rupee (e₹), a non-interest-bearing, capped CBDC implemented in a mature but cash-intensive banking system with a dominant low-cost digital rail (UPI). The question is salient because central banks around the globe are actively experimenting with retail CBDC designs, yet the academic debate has been almost entirely theoretical and delivers sign-ambiguous predictions that hinge on assumptions about CBDC design choices and household behavior. We provide the first empirical evidence on the real-world consequences of a retail CBDC along the three core margins that dominate this theoretical literature.

Our empirical strategy combines two complementary sources of quasi-experimental variation, one at the aggregate level and one at the household level. At the district level, we exploit the interaction of India's long-standing lead bank system with the Reserve Bank of India's staggered authorization of banks to offer CBDC, which generates plausibly exogenous timing of CBDC access across districts and allows us to study the effects on aggregate deposits and credit. At the individual level, we use centrally scheduled CBDC marketing campaigns at a large public-sector bank, rolled out in waves across districts, to trace how intensive promotion of CBDC affects households' portfolio choices, payment behavior, and adoption hazards over time.

Across these settings, we find that the introduction of a conservatively designed retail CBDC

generates meaningful reallocation of liquid balances and sizable shifts in digital payment flows, but only a modest and transitory contraction in bank lending. On the banking side, CBDC eligibility leads to a persistent decline in retail deposits of about 2.7%, concentrated in liquid low-yield savings accounts, while illiquid and high-yield deposits remain largely unchanged. Although total credit falls initially following the introduction of CBDC, this decline reverts to zero over time. This reversion of the credit effect, despite a persistent effect on deposits, is consistent with banks offsetting deposit outflows by adjusting their funding mix and portfolios rather than sharply cutting lending. On the payments side, CBDC primarily crowds out deposit-backed digital payments: UPI volumes drop by around 6.3% whereas the effect on ATM withdrawals is economically small and statistically insignificant. This indicates that CBDC can be a strong competitor to existing digital rails but a weak substitute for cash. On the household side, early adoption is highly skewed toward male, higher-income, highly educated, and high-deposit customers holding uninsured deposits. As a result, the first wave of CBDC users consists of households that are already well integrated into the formal financial system, so the initial phase of CBDC diffusion reallocates safety and convenience among already-included households rather than immediately closing financial inclusion gaps.

These results highlight a clear policy tradeoff. Retail CBDC reduces bank deposits and generates short-run pressure on credit supply, but it also delivers substantial convenience and safety benefits to adopting households. To quantify this tradeoff, we combine our empirical estimates with a simple sufficient-statistics framework that maps the observed decline in deposits, the temporary contraction in credit, and the revealed surplus from adoption into district-level welfare. This exercise points to moderate but economically meaningful welfare gains in the long-run when benchmarked against household income, foregone interest, and the size of the deposit outflow. It also allows us to explore the potential contributions of different channels through the lens of our banking block. Within this framework, the largest short-term cost arises from credit disintermediation. The subsequent recovery in lending is modeled as banks adjusting their funding mix. We note that our district-level estimates of total credit do not isolate credit supply from demand, nor do we directly observe wholesale financing. Thus, these specific mechanisms remain model-based interpretations, and the observed recovery could also reflect other balance-sheet adjustments or evolving loan demand.

Taken together, our findings speak directly to ongoing policy debates and to the theoretical literature on CBDC design. They suggest that even a non-interest-bearing, tightly capped, intermediated CBDC can meaningfully disintermediate deposits and reshape digital payment flows, but that its impact on lending can remain limited if banks have access to alternative funding sources and adjust their balance sheets accordingly. At the same time, the strong adoption gradient we document implies that, under conservative designs, the early benefits of CBDC accrue primarily to relatively affluent and financially sophisticated households. These patterns highlight that CBDC design choices, especially

remuneration, holding caps, and distribution, will be central in determining not only the extent of deposit disintermediation and competition in digital payments, but also who gains access to central bank money in practice and whether CBDC moves economies closer to, or further from, their stated goals for financial inclusion and stability.

References

Abad, Jorge, Galo Nuño, and Carlos Thomas. 2025. "CBDC and the Operational Framework of Monetary Policy." *Journal of Monetary Economics* 151:103762.

Adalid, Ramón, Álvaro Álvarez-Blázquez, Katrin Assenmacher, Lorenzo Burlon, Maria Dimou, Carolina López-Quiles, Natalia Martín Fuentes, Barbara Meller, Manuel A. Muñoz, Petya Radulova, Costanza Rodriguez d'Acri, Tamarah Shakir, Gabriela Šílová, Oscar Soons, and Alexia Ventula Veghazy. 2022. "Central Bank Digital Currency and Bank Intermediation." Occasional Paper Series 293, European Central Bank.

Agur, Itai, Anil Ari, and Giovanni Dell'Ariccia. 2022. "Designing Central Bank Digital Currencies." *Journal of Monetary Economics* 125:62–79.

Ahnert, Toni, Katrin Assenmacher, Peter Hoffmann, Agnese Leonello, Cyril Monnet, and Davide Porcellacchia. 2024. "The Economics of Central Bank Digital Currency." *International Journal of Central Banking* 20.

Ahnert, Toni, Peter Hoffmann, Agnese Leonello, and Davide Porcellacchia. 2025. "Central Bank Digital Currency and Financial Stability." Working Paper.

Alok, Shashwat, Pulak Ghosh, Nirupama Kulkarni, and Manju Puri. 2025. "Breaking Barriers to Financial Access: Cross-Platform Digital Payments and Credit Markets." Available at SSRN 5028744.

Andolfatto, David. 2021. "Assessing the Impact of Central Bank Digital Currency on Private Banks." *The Economic Journal* 131 (634):525–540.

Assenmacher, Katrin, Aleksander Berentsen, Claus Brand, and Nora Lamersdorf. 2021. "A Unified Framework for CBDC Design: Remuneration, Collateral Haircuts and Quantity Constraints." Working Paper Series 2578, European Central Bank.

Assenmacher, Katrin, Lea Bitter, and Annukka Ristiniemi. 2023. "CBDC and business cycle dynamics in a New Monetarist New Keynesian model." .

Auer, Raphael, Holti Banka, Nana Yaa Boakye-Adjei, Ahmed Faragallah, Jon Frost, Harish Natarajan, and Jermy Prenio. 2022a. "Central Bank Digital Currencies: A New Tool in the Financial Inclusion Toolkit?" *FSI Insights* 41.

Auer, Raphael, Giulio Cornelli, and Jon Frost. 2023. "Rise of the Central Bank Digital Currencies." *International Journal of Central Banking* 19 (4):185–214.

Auer, Raphael, Jon Frost, Leonardo Gambacorta, Cyril Monnet, Tara Rice, and Hyun Song Shin. 2022b. "Central Bank Digital Currencies: Motives, Economic Implications, and the Research Frontier." *Annual Review of Economics* 14 (1):697–721.

Barrdear, John and Michael Kumhof. 2022. "The Macroeconomics of Central Bank Digital Currencies." *Journal of Economic Dynamics and Control* 142:104148.

Berg, Tobias, Jan Keil, Felix Martini, and Manju Puri. 2024. "CBDCs, Payment Firms, and Geopolitics." Tech. rep., National Bureau of Economic Research.

Bijlsma, Michiel, Carin van der Cruijsen, Nicole Jonker, and Jelmer Reijerink. 2024. "What Triggers Consumer Adoption of Central Bank Digital Currency?" *Journal of Financial Services Research* 65:1–40.

Bindseil, Ulrich. 2020. "Tiered CBDC and the Financial System." Working Paper Series 2351, European Central Bank.

Bordo, Michael D and Andrew T Levin. 2017. "Central Bank Digital Currency and the Future of Monetary Policy." Tech. rep., National Bureau of Economic Research.

Borusyak, Kirill, Xavier Jaravel, and Jann Spiess. 2024. "Revisiting Event-Study Designs: Robust and Efficient Estimation." *The Review of Economic Studies* 91:3253–3285.

Brunnermeier, Markus K, Harold James, and Jean-Pierre Landau. 2019. "The Digitalization of Money." Tech. rep., National Bureau of Economic Research.

Brunnermeier, Markus K and Dirk Niepelt. 2019. "On the Equivalence of Private and Public Money." *Journal of Monetary Economics* 106:27–41.

Burgess, Robin and Rohini Pande. 2005. "Do rural banks matter? Evidence from the Indian social banking experiment." *American economic review* 95 (3):780–795.

Burlon, Lorenzo, Manuel A Muñoz, and Frank Smets. 2024. "The Optimal Quantity of CBDC in a Bank-Based Economy." *American Economic Journal: Macroeconomics* 16 (4):172–217.

Callaway, Brantly and Pedro H. C. Sant'Anna. 2021. "Difference-in-Differences with Multiple Time Periods." *Journal of Econometrics* 225:200–230.

Carapella, Francesca, Jin-Wook Chang, Sebastian Infante, Melissa Leistra, Arazi Lubis, and Alexandros Vardoulakis. 2024. "Financial Stability Implications of CBDC." FEDS Working Paper.

Carapella, Francesca and Jean Flemming. 2020. "Central Bank Digital Currency: A Literature Review." FEDS Notes.

Chang, Huifeng, Federico Grinberg, Lucyna Gornicka, Marcello Miccoli, and Brandon Joel Tan. 2023. "Central Bank Digital Currency and Bank Disintermediation in a Portfolio Choice Model." Working Paper.

Chapman, James, Jonathan Chiu, Seyed Mohammadreza Davoodalhosseini, Janet Hua Jiang, Francisco Rivadeneyra, and Yu Zhu. 2023. "Central Bank Digital Currencies and Banking: Literature Review and New Questions." Bank of Canada Staff Discussion Paper.

Chetty, Raj. 2009. "Sufficient Statistics for Welfare Analysis: A Bridge Between Structural and Reduced-Form Methods." *Annual Review of Economics* 1:451–488.

Chiu, Jonathan, Seyed Mohammadreza Davoodalhosseini, Janet Jiang, and Yu Zhu. 2023. "Bank Market Power and Central Bank Digital Currency: Theory and Quantitative Assessment." *Journal of Political Economy* 131:1213–1248.

Cox, D. R. 1972. "Regression Models and Life-Tables." *Journal of the Royal Statistical Society: Series B (Methodological)* 34:187–220.

Cramer, Kim Fe, Pulak Ghosh, Nirupama Kulkarni, and Nishant Vats. 2025. *Shadow Banks on the Rise: Evidence Across Market Segments*. Financial Markets Group.

Das, Mitali, Tommaso Mancini-Griffoli, Fumitaka Nakamura, Julia Otten, Gabriel Soderberg, Juan Sole, and Brandon Joel Tan. 2023. "Implications of Central Bank Digital Currencies for Monetary Policy Transmission." *IMF Fintech Notes* 2023 (010):1.

Davoodalhosseini, Seyed Mohammadreza. 2022. "Central Bank Digital Currency and Monetary Policy." *Journal of Economic Dynamics and Control* 142:104150.

de Chaisemartin, Clément, Xavier D'Haultfœuille, and Gonzalo Vazquez-Bare. 2024. "Difference-in-Difference Estimators with Continuous Treatments and No Stayers." *AEA Papers and Proceedings* 114:610–13.

Dinger, Valeriya and Ben R. Craig. 2015. "Volatile Lending and Bank Wholesale Funding." Working paper.

Drechsler, Itamar, Alexi Savov, Philipp Schnabl, and Olivier Wang. 2026. "Deposit franchise runs." *The Journal of Finance* 81 (3):1573–1617.

Dubey, Tamanna Singh and Amiyatosh Purnanandam. 2025. "Can Cashless Payments Spur Economic Growth?" Available at SSRN 4373602.

Duffie, Darrell. 2021. "Building a stronger financial system: Opportunities for a digital dollar." *Central Bank Digital Currency Considerations, Projects, Outlook* .

Duffie, Darrell and Elizabeth Economy. 2022. *Digital Currencies: the US, China, and the World at a Crossroads*. Hoover Institution Press, Stanford University.

Duffie, Darrell et al. 2019. "Digital Currencies and Fast Payment Systems: Disruption is Coming." In *Asian Monetary Forum, May, mimeo*, vol. 1.

European Central Bank. 2020. "Report on a Digital Euro." *https://www. ecb. europa. eu/pub/pdf/other/Report_on_a_digital_euro 4d7268b458. en. pdf* .

Federal Reserve Board of Governors. 2022. "Money and Payments: The U.S. Dollar in the Age of Digital Transformation."

Fernández-Villaverde, Jesús, Daniel Sanches, Lina Schilling, and Harald Uhlig. 2021. "Central Bank Digital Currency: Central Banking for All?" *Review of Economic Dynamics* 41:225–242.

García, Alejandro, Bena Lands, Xuezhi Liu, and Joshua Slive. 2020. "The Potential Effect of a Central Bank Digital Currency on Deposit Funding in Canada." Staff Analytical Notes 2020-15, Bank of Canada.

Garratt, Rodney, Jiaheng Yu, and Haoxiang Zhu. 2022. "How Central Bank Digital Currency Design Choices Impact Monetary Policy Pass-Through and Market Composition." Working Paper.

Ghosh, Pulak, Nicola Limodio, and Nishant Vats. 2025. "Household portfolio and deposit insurance: Implications for the supply of safe assets." *Available at SSRN 5488990* .

Gross, Jonas and Jonathan Schiller. 2021. "A Model for Central Bank Digital Currencies: Implications for Bank Funding and Monetary Policy." Working Paper.

Gust, Christopher, Kyungmin Kim, and Romina Ruprecht. 2023. "The Effects of CBDC on the Federal Reserve's Balance Sheet." FEDS Working Paper.

Infante, Sebastian, Kyungmin Kim, Anna Orlik, André F Silva, and Robert Tetlow. 2022. "The Macroeconomic Implications of CBDC: A Review of the Literature." FEDS Working Paper.

———. 2023. "CBDC: Issues and Prospects." *European Economy* :53–64.

———. 2024. "Retail CBDC: Implications for Banking and Financial Stability." *Annual Review of Financial Economics* 16:207–232.

Kaplan, E. L. and Paul Meier. 1958. "Nonparametric Estimation from Incomplete Observations." *Journal of the American Statistical Association* 53:457–481.

Kaul, Vivek. 2021. "Why the Reserve Bank Wants to Have Its Own Digital Currency." Mint. Available here (Accessed: March 28, 2026).

Kearns, Jeff and Ashlin Mathew. 2022. "How India's Central Bank Helped Spur a Digital Payments Boom." IMF Country Focus. Available here (Accessed: March 28, 2026).

Keister, Todd and Cyril Monnet. 2022. "Central Bank Digital Currency: Stability and Information." *Journal of Economic Dynamics and Control* 142:104501.

Keister, Todd and Daniel Sanches. 2023. "Should Central Banks Issue Digital Currency?" *The Review of Economic Studies* 90 (1):404–431.

Kim, Young Sik and Ohik Kwon. 2022. "Central Bank Digital Currency, Credit Supply, and Financial Stability." *Journal of Money, Credit and Banking* 55:297–321.

Kosanović, Nikola. 2025. "Central Bank Digital Currencies: Financial Stability Perspective." *Journal of Central Banking Theory and Practice* 14 (1):145–162.

Kumhof, Michael and Clare Noone. 2021. "Central Bank Digital Currencies — Design Principles for Financial Stability." *Economic Analysis and Policy* 71:553–572.

Kumhof, Michael, Marco Pinchetti, Phurichai Rungcharoenkitkul, and Andrej Sokol. 2023. "CBDC policies in open economies." *Available at SSRN 4388834* .

Li, Jiaqi. 2023. "Predicting the Demand for Central Bank Digital Currency: A Structural Analysis with Survey Data." *Journal of Monetary Economics* 134:73–85.

Li, Jiaqi, Andrew Usher, and Yu Zhu. 2025. "Central Bank Digital Currency and Banking Choices." Working Paper.

Maniff, Jesse Leigh. 2020. "Inclusion by Design: Crafting a Central Bank Digital Currency to Reach All Americans." Payments System Research Briefing, Federal Reserve Bank of Kansas City.

Meaning, Jack, Ben Dyson, James Barker, and Emily Clayton. 2018. "Broadening Narrow Money: Monetary Policy with a Central Bank Digital Currency." Bank of England Working Paper.

Minesso, Massimo Ferrari, Arnaud Mehl, and Livio Stracca. 2022. "Central Bank Digital Currency in an Open Economy." *Journal of Monetary Economics* 127:54–68.

Murakami, David, Ivan Shchapov, and Ganesh Viswanath-Natraj. 2024. "CBDCs, Financial Inclusion, and Optimal Monetary Policy." Working Paper.

Niepelt, Dirk. 2021. "Central bank digital currency: Considerations, projects, outlook." *VoxEU* .

———. 2024. "Money and Banking with Reserves and CBDC." *Journal of Finance* 79:2505–2552.

Paul, Pascal, Mauricio Ulate, and Jing Cynthia Wu. 2026. "A Macroeconomic Model of Central Bank Digital Currency." NBER Working Paper No. w33968.

Pennacchi, George. 2012. "Narrow Banking." *Annual Review of Financial Economics* 4 (1):141–159.

Piazzesi, Monika and Martin Schneider. 2022. "Credit Lines, Bank Deposits or CBDC? Competition & Efficiency in Modern Payment Systems." Working Paper.

PIB Thiruvananthpuram. 2023. "India's UPI: A Global Front-Runner in Digital Payment Systems." Available here (Accessed: March 28, 2026).

PwC India. 2023. "The Indian Payments Handbook – 2022–2027." PwC. Available here (Accessed: March 28, 2026).

Reserve Bank of India. 2018. "Master Circular – Lead Bank Scheme." Available here (Accessed: March 27, 2026).

———. 2022a. "Concept Note on Central Bank Digital Currency." Available here (Accessed: March 28, 2026).

———. 2022b. "Operationalisation of Central Bank Digital Currency – Retail (e₹-R) Pilot." Available here (Accessed: March 28, 2026).

———. 2024. "Reserve Bank of India Bulletin September 2024." Available here (Accessed: March 28, 2026).

———. 2026. "Digital Rupee (e ₹) – FAQs." Available here (Accessed: March 27, 2026).

Reuters. 2022. "India Cenbank to Start Pilot of Retail Digital Rupee on Dec. 1." Reuters. Available here (Accessed: March 27, 2026).

Roth, Jonathan, Pedro H. C. Sant'Anna, Alyssa Bilinski, and John Poe. 2023. "What's Trending in Difference-in-Differences? A Synthesis of the Recent Econometrics Literature." *Journal of Econometrics* 235:2218–2244.

Schilling, Linda, Jesús Fernández-Villaverde, and Harald Uhlig. 2021. "Facing the Central Bank Digital Currency Trilemma." In *Central Bank Digital Currency: Considerations, Projects, Outlook*, edited by Dirk Niepelt. Paris & London: CEPR Press, 45–50.

———. 2024. "Central Bank Digital Currency: When Price and Bank Stability Collide." *Journal of Monetary Economics* 145:103554.

Shy, Oz. 2021. "Digital Currency, Digital Payments, and the 'Last Mile' to the Unbanked." Federal Reserve Bank of Atlanta Policy Hub.

Tan, Brandon Joel. 2024. "Central Bank Digital Currency and Financial Inclusion." *Journal of Macroeconomics* 81:103620.

Wang, Xin and Xiaoxiong Hu. 2025. "Financial Development, Non-Bank E-Money, and Central Bank Digital Currency." Working Paper.

Whited, Toni M., Yufeng Wu, and Kairong Xiao. 2023. "Will Central Bank Digital Currency Disintermediate Banks?" Working Paper.

Williamson, Stephen. 2022a. "Central Bank Digital Currency: Welfare and Policy Implications." *Journal of Political Economy* 130 (11):2829–2861.

Williamson, Stephen D. 2022b. "Central Bank Digital Currency and Flight to Safety." *Journal of Economic Dynamics and Control* 142:104146.

Table 1: Staggered Eligibility of Lead Banks for CBDC

Lead Bank	Quarter Eligible for CBDC
State Bank of India	2022 Q4
Bank of Baroda	2023 Q1
Canara Bank	2023 Q1
Union Bank of India	2023 Q1
Punjab National Bank	2023 Q2
Indian Bank	2024 Q1
Bank of India	2024 Q3
UCO Bank	2025 Q2
Bank of Maharashtra	2025 Q2
Central Bank of India	Never Eligible
Indian Overseas Bank	Never Eligible
Jammu and Kashmir Bank	Never Eligible

This table lists the quarter in which each lead bank became eligible to offer the CBDC to its customers as part of the Reserve Bank of India's phased roll-out. Lead banks that were never assigned CBDC eligibility during our sample period (2018 Q1 through 2025 Q3) are listed as "Never Eligible."

Table 2: Summary Statistics

Panel A: Quarterly Bank Deposits and Credit				
	# Obs	Median	Mean	SD
Log Total Retail Deposits	20,863	24.871	24.869	1.417
Log Credit	20,863	24.398	24.374	1.525
Panel B: Individual-Level Bank Deposits and Credit				
	# Obs	Median	Mean	SD
Log Saving Balance	1,591,968	13.518	13.207	2.737
Log Fixed Deposit Balance	1,591,968	0.000	2.645	6.135
Log Recurring Deposit Balance	1,591,968	0.000	0.638	3.042
Log Total Deposit Balance	1,591,968	13.774	13.641	3.043
Panel C: Monthly Transaction Usage				
	# Obs	Median	Mean	SD
Log CBDC Volume	1,591,968	0.000	0.034	0.541
Log Number of CBDC Transactions	1,591,968	0.000	0.007	0.130
Log UPI Volume	1,591,968	11.974	10.021	4.831
Log Number of UPI Transactions	1,591,968	2.639	2.355	1.451
Log ATM Withdrawal Volume	1,591,968	0.000	3.031	4.291
Panel D: User Characteristics				
	# Obs	Median	Mean	SD
CBDC User	153,875	0.000	0.044	0.206
UPI User	153,875	0.000	0.498	0.500
Male	153,875	1.000	0.592	0.491
Age	153,875	36.710	39.394	13.791
Log Income	153,875	11.695	11.973	0.972
Sophisticated Education	153,875	0.000	0.009	0.093
Deposits (Rupees)	153,875	352,390.000	15,664,155.341	53,355,221.855

This table reports descriptive statistics for the key variables used in the analysis. Panel A reports district-level quarterly bank deposit variables, each measured as the natural logarithm of aggregate deposits across all banks within the district. Log Total Retail Deposits, and Log Credit correspond to total retail deposits, and credit, respectively. Panel B reports individual-level bank deposit variables. Log Saving Balance, Log Fixed Deposit Balance, and Log Recurring Deposit Balance are measured as the natural logarithm of the outstanding balance in rupees in each account type; Log Total Deposit Balance is the natural logarithm of the sum of these three components. Saving balance is computed as the within-month average of saving deposit balance. Panel C reports individual-level monthly transaction variables. Log Volume variables are measured as the natural logarithm of total transaction value in rupees; Log Number variables are measured as the natural logarithm of total transaction count. Log ATM Withdrawal Volume is the natural logarithm of total ATM withdrawal value in rupees. Panel D reports individual-level user characteristics. CBDC User and UPI User are binary indicators equal to one if the user has ever adopted CBDC or used UPI, respectively. Male and Sophisticated Education are binary indicators, where the latter equals one if the user holds a computer, finance, or engineering degree. Age is measured in years and Log Income is the natural logarithm of income. Deposits is the median total deposit balance over the six months prior to CBDC treatment, winsorized at the 1% level. All continuous variables in Panels A, B, and C, and Log Income in Panel D are winsorized at the 1% level.

Table 3: Determinants of CBDC Adoption—Cox Proportional Hazards Model Hazard Ratios

	CBDC Adoption								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Male	2.8970***						2.2645***	2.3009***	2.3181***
	(21.1351)						(16.8969)	(16.9701)	(17.7767)
Age Q2		2.5482***					1.3221***	1.3168***	1.2720***
		(21.1049)					(6.2303)	(6.3685)	(5.4653)
Age Q3		2.0322***					0.9567	0.9744	0.9370
		(12.4735)					(−0.7801)	(−0.4815)	(−1.1969)
Age Q4		1.1778**					0.5544***	0.5889***	0.5579***
		(2.1520)					(−7.3152)	(−6.8386)	(−7.7456)
Income Q2			1.1423*				0.9530	0.9769	0.9608
			(1.8545)				(−0.6917)	(−0.3311)	(−0.5582)
Income Q3			2.0807***				1.6741***	1.6578***	1.5787***
			(10.9646)				(8.2766)	(8.1171)	(7.1512)
Income Q4			11.0440***				6.6799***	6.2192***	5.6972***
			(38.7649)				(34.9448)	(35.4167)	(32.3191)
Pre CBDC-Eligible UPI User				1.4954***			0.8103***	0.8187***	0.8437***
				(8.6334)			(−6.6119)	(−6.4090)	(−5.0039)
Sophisticated Education					6.5134***		2.3992***	2.1937***	2.0004***
					(26.5608)		(13.1090)	(11.6986)	(8.3741)
Above Insurance Threshold						5.8280***	4.3922***	4.6822***	4.0542***
						(35.0755)	(30.9560)	(29.9270)	(29.5688)
# Obs	153,875	153,875	153,875	153,875	153,875	153,875	153,875	153,875	153,875
Pseudo R^2 Nagelkerke	0.0095	0.0063	0.0445	0.0017	0.0045	0.0261	0.0694	0.0662	0.0519
District FE								Yes	
ZIP code FE									Yes

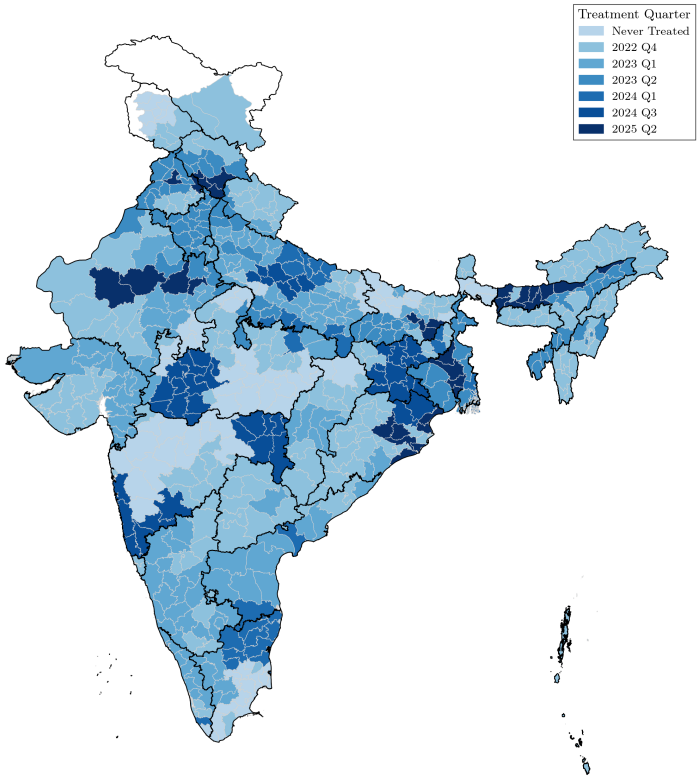
This table presents coefficients from Cox proportional hazards models estimating the determinants of CBDC adoption. The dependent variable is the time until CBDC adoption, measured in weeks since becoming targeted for CBDC marketing campaigns. Columns (1) through (6) examine the effect of each covariate separately: gender, age quartile, income quartile, prior UPI usage, sophisticated education, and deposit insurance threshold, respectively. Column (7) presents the full specification including all covariates without fixed effects, column (8) adds district fixed effects, and column (9) adds ZIP code fixed effects in place of district fixed effects. Fixed effects are implemented via stratification. The omitted categories are female, age quartile 1 (youngest), annual income quartile 1 (lowest), no prior UPI usage, no sophisticated education, and below deposit insurance threshold. Sophisticated education is defined as holding a computer, finance, or engineering degree. Deposits is the median total deposit balance over the six months prior to CBDC treatment. Above the deposit insurance threshold is defined as having a deposit balance above ₹500,000. All models use robust standard errors clustered at the district level. Reported coefficients are expressed as hazard ratios ($e^{\hat{\beta}}$) and represent the multiplicative difference in the instantaneous rate of CBDC adoption relative to the reference group at any given point in time, conditional on not yet having adopted. z-statistics are reported in parentheses. For example, the estimate of 2.8970 in column (1) implies that men adopt CBDC at approximately $e^{1.0637} = 2.8970$ times the hazard rate of women, meaning that in any given week, men are adopting at 2.8970x the rate of women, conditional on not having adopted yet. Table D.1 presents the estimated coefficients without any transformations. The sample includes all individuals from the start of the CBDC roll-out through the end of March 2024. The Breslow method is used to adjust for ties (individuals adopting at the same time). ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Table 4: Composition of the District-Level Welfare Change

	Short-Run Impact ($\theta_{SR} = 0.476$)	Long-Run Impact ($\theta_{SR} = 0$)
(i) Household surplus	+54.6	+54.6
	(+74.7%)	(+74.7%)
(ii) Lost deposit franchise	−36.9	−36.9
	(−50.5%)	(−50.5%)
(iii) Wholesale deadweight loss	−3.8	0.0
	(−5.2%)	(0.0%)
(iv) Lending curvature loss	−3.4	0.0
	(−4.7%)	(0.0%)
(v) Credit-contraction externality	−43.9	0.0
	(−60.1%)	(0.0%)
(vi) Net seigniorage	+18.4	+18.4
	(+25.3%)	(+25.3%)
Net welfare $\Delta\mathcal{W}$	−15.0	+36.1
	(−20.5%)	(+49.5%)

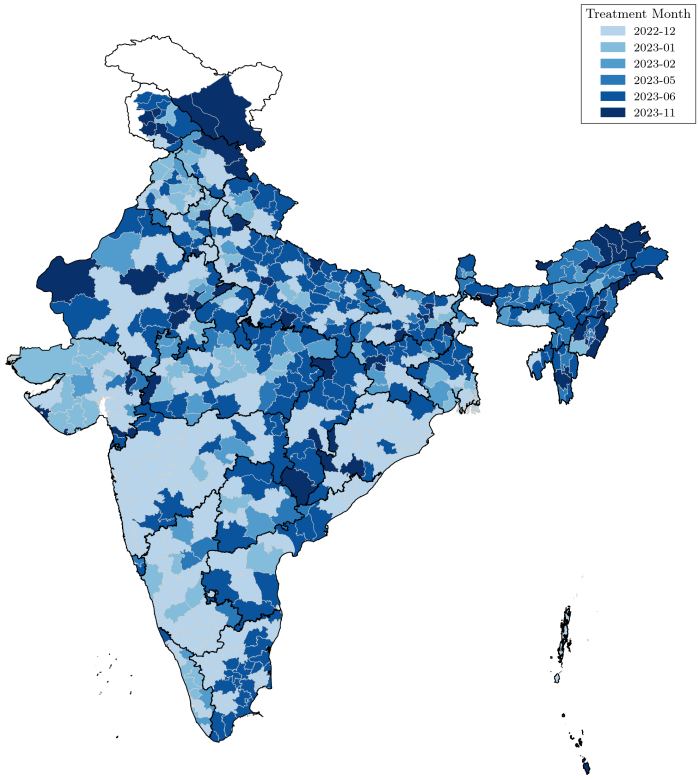
This table presents the calculations of each term in the equation for net welfare gain (13). For the calculations, we assume the following values: deposit outflow $|dD| = 1.84$ billion, pass-through $\theta_{SR} = 0.476$, adoption rate $a = 0.044$, adopters $Na = 88{,}000$, adopter surplus $\overline{\Delta v} = 620$ rupees, franchise spread $\lambda = 2$ percentage points, wholesale slope $\psi_{SR}|dD| \approx 1.5$ percentage points, credit wedge $\eta = 0.05$, net seigniorage $s^{\text{CB}} = 1\%$, savings rate $r^S = 3\%$, and lending curvature penalty $c''(\ell)_{SR}|dD| \approx 1.65$ percentage points. In the long-term calculation, we assume that $\psi_{LR} \to 0$ and $\theta_{LR} \to 0$. The welfare estimates are presented in millions of rupees. The percentage of gross gain is presented in parentheses. Gross gain is defined as the sum of the terms (i) and (vi): 54.6 + 18.4 = 73 million rupees.

Figure 1: Geographic Roll-Out of CBDC Across Lead Bank Districts



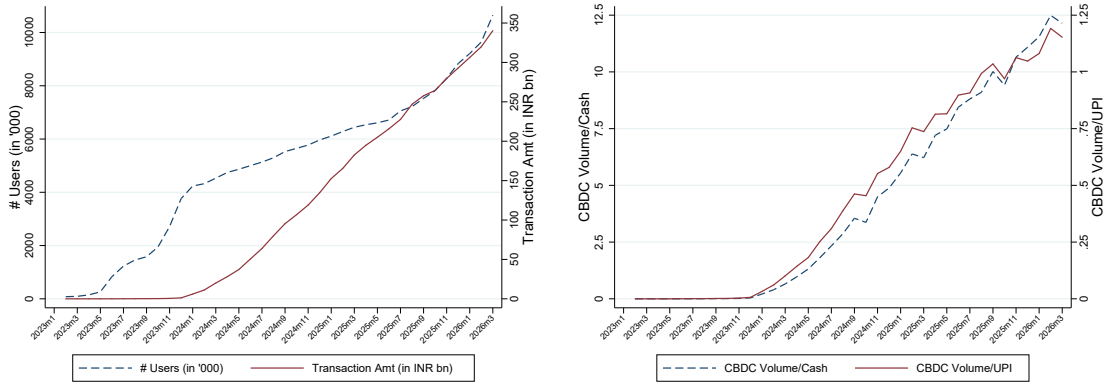
This figure presents a district-level map of India colored by the quarter in which the district's assigned lead bank first became eligible to offer CBDC. Districts shaded in lighter blue received CBDC eligibility earlier, while darker shades indicate later eligibility cohorts. Districts whose lead bank was never assigned CBDC eligibility during our sample period are shaded in the lightest gray. Each district is assigned to exactly one lead bank, and the treatment timing of the district is determined by the quarter in which that lead bank became eligible for CBDC, as listed in Table 1.

Figure 2: Geographic Roll-Out of CBDC Marketing Campaigns Across Districts in Sample Bank



This figure presents a district-level map of India colored by the month in which the district became targeted with intensive CBDC marketing campaigns. Districts shaded in lighter blue received marketing campaigns earlier, while darker shades indicate later treatment cohorts.

Figure 3: Aggregate Evolution of Retail CBDC Usage in India

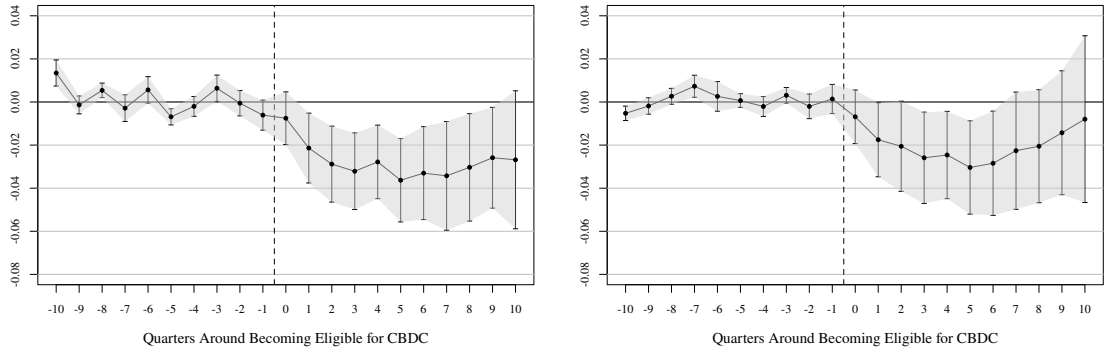


(a) CBDC Users and Transaction Volumes

(b) CBDC Volumes Relative to UPI and Cash

This figure plots the evolution of aggregate retail CBDC usage in India between February 2023 and March 2026. Panel (a) reports the total number of registered CBDC users (in thousands, left axis) and total monthly CBDC transaction volume (in ₹ billions, right axis). Panel (b) reports the ratio of CBDC transaction volume to total UPI transaction volume and to total cash transaction volume, expressed in percent. CBDC users and transaction volumes are constructed from bank-reported administrative data on CBDC wallets and transactions. UPI and cash volumes are taken from publicly available Reserve Bank of India statistics and are matched by month. The data on the number of CBDC users and transaction volume are provided by the Fintech department of the Reserve Bank of India.

Figure 4: Effect of CBDC Eligibility on Bank Deposits and Credit



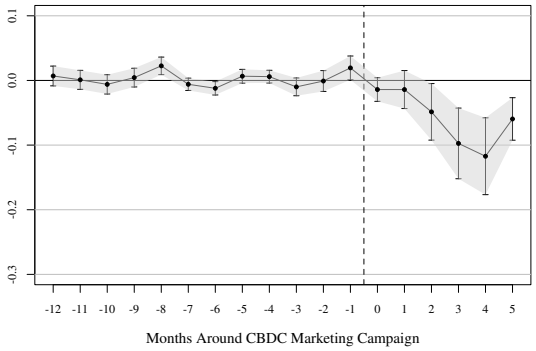
(a) Log Total Retail Deposit

(b) Log Credit

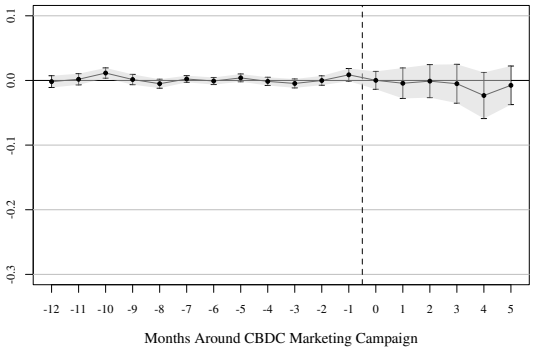
	ATT Estimates		
	Simple	Group-Specific	Event Study
Log Total Retail Deposit	−0.0273***	−0.0237***	−0.0277***
	(0.0084)	(0.0071)	(0.0089)
Log Credit	−0.0201*	−0.018**	−0.020*
	(0.0104)	(0.0089)	(0.0104)

This figure plots the event-study dynamic coefficients estimated using the methodology outlined in Callaway and Sant'Anna (2021). The outcome variable is the natural logarithm of total retail deposits in panel (a), and credit in panel (b), each measured at the district × quarter level. Treatment timing is defined as the quarter in which the lead bank of each district first received RBI authorization to offer CBDC. The x-axis denotes quarters relative to the district's first quarter of CBDC eligibility. The control group consists of districts whose lead bank had not yet received CBDC eligibility within our sample period. All outcome variables are winsorized at the 1% level. The shaded bands represent 95% confidence intervals. The table below the figures reports the estimates of the aggregate average treatment effect on the treated (ATT) by taking a weighted average of all group-time specific, partially aggregated group-specific and partially aggregated event-time specific average treatment effects following Callaway and Sant'Anna (2021). Standard errors clustered at the district level shown in parentheses. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

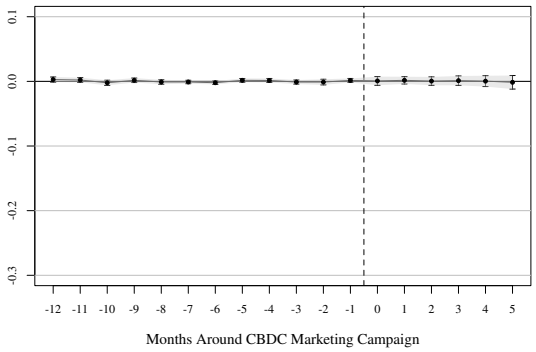
Figure 5: Monthly Effect of CBDC Marketing Campaign on Deposit Balances



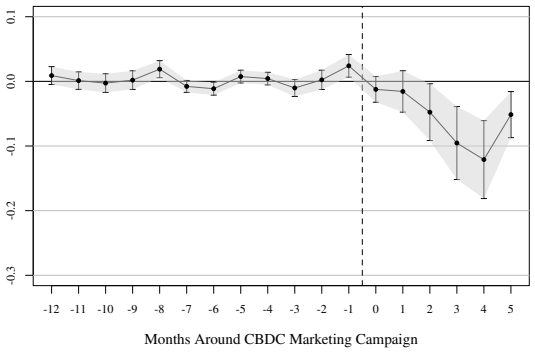
(a) Log Saving Balance



(b) Log Fixed Deposit Balance



(c) Log Recurring Deposit Balance

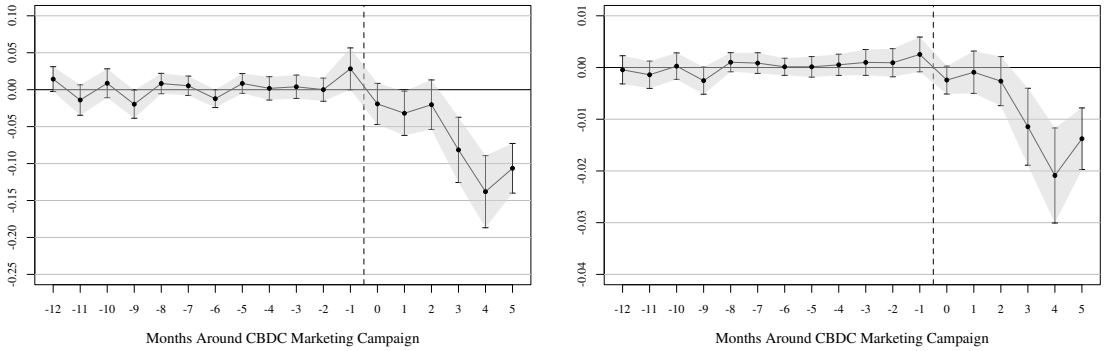


(d) Log Total Deposit Balance

	ATT Estimates		
	Simple	Group-Specific	Event Study
Log Saving Balance	−0.0585***	−0.0634***	−0.0585***
	(0.0171)	(0.0181)	(0.0174)
Log Fixed Deposit Balance	−0.0068	−0.0086	−0.0069
	(0.0125)	(0.0133)	(0.0120)
Log Recurring Deposit Balance	0.0005	0.0006	0.0005
	(0.0031)	(0.0034)	(0.0033)
Log Total Deposit Balance	−0.0575***	−0.0631***	−0.0572***
	(0.0185)	(0.0195)	(0.0175)

This figure plots the event-study dynamic coefficients estimated using the methodology outlined in Callaway and Sant'Anna (2021). The outcome variable is the natural logarithm of a user's month average savings balance in rupees in panel (a), the natural logarithm of a user's month end fixed deposit balance in panel (b), the natural logarithm of a user's month end recurring deposit balance in panel (c), and the natural logarithm of a user's total deposit balance, defined as the sum of savings, fixed deposit, and recurring deposit, in panel (d), each measured at the user × month level. Treatment timing is defined as the month in which the user's district first became treated for CBDC marketing campaigns. The x-axis denotes months relative to the district's first month of CBDC marketing campaigns. The control group consists of users in districts that had not yet received CBDC marketing campaigns. All outcome variables are winsorized at the 1% level. The shaded bands represent 95% confidence intervals. The table below the figures reports the estimates of the aggregate average treatment effect on the treated (ATT) by taking a weighted average of all group-time specific, partially aggregated group-specific and partially aggregated event-time specific average treatment effects following Callaway and Sant'Anna (2021). Standard errors clustered at the district level are shown in parentheses. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Figure 6: Monthly Effect of CBDC Marketing Campaign on UPI Usage



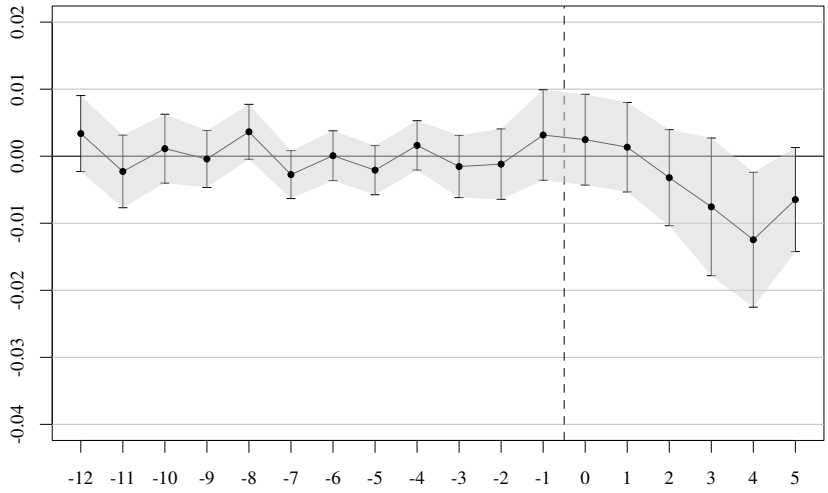
(a) Log UPI Volume

(b) Log Number of UPI Transactions

	ATT Estimates		
	Simple	Group-Specific	Event Study
Log UPI Volume	−0.0648***	−0.0666***	−0.0662***
	(0.0142)	(0.0160)	(0.0142)
Log Number of UPI Transactions	−0.0085***	−0.0089***	−0.0087***
	(0.0021)	(0.0023)	(0.0021)

This figure plots the event-study dynamic coefficients estimated using the methodology outlined in Callaway and Sant'Anna (2021). The outcome variable is the natural logarithm of total UPI transaction volume in rupees in panel (a), and the natural logarithm of the total number of UPI transactions in panel (b), each measured at the user × month level. Treatment timing is defined as the month in which the user's district first became treated for CBDC marketing campaigns. The x-axis denotes months relative to the district's first month of CBDC marketing campaigns. The control group consists of users in districts that had not yet received CBDC marketing campaigns. All outcome variables are winsorized at the 1% level. The shaded bands represent 95% confidence intervals. The table below the figures reports the estimates of the aggregate average treatment effect on the treated (ATT) by taking a weighted average of all group-time specific, partially aggregated group-specific and partially aggregated event-time specific average treatment effects following Callaway and Sant'Anna (2021). Standard errors clustered at the district level are shown in parentheses. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

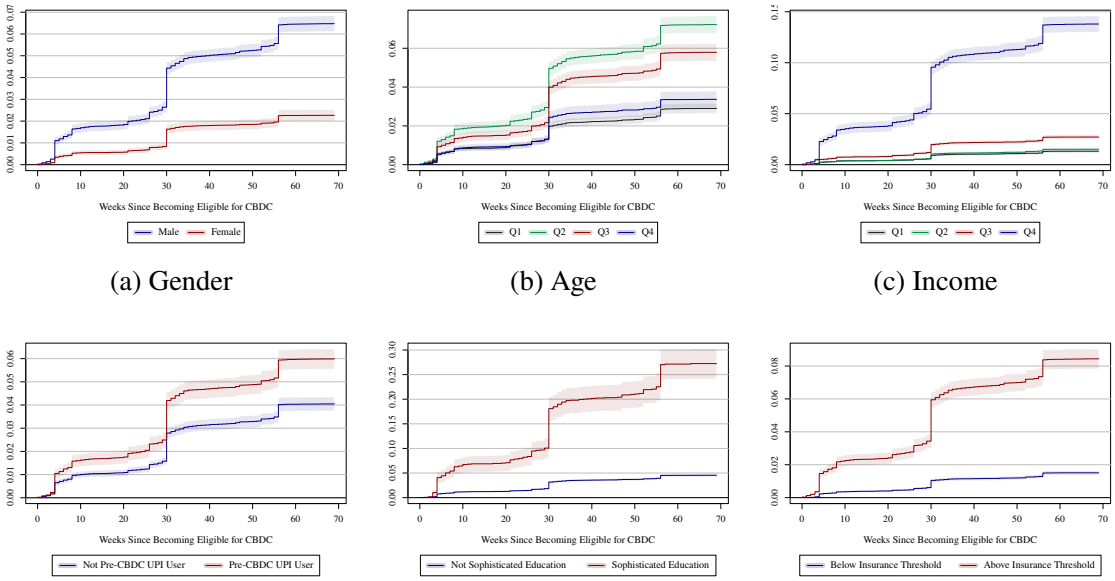
Figure 7: Monthly Effect of CBDC Marketing Campaign on ATM Withdrawals



	ATT Estimates		
	Simple	Group-Specific	Event Study
Log ATM Withdrawal Volume	−0.0042	−0.0044	−0.0043
	(0.0033)	(0.0037)	(0.0032)

This figure plots the event-study dynamic coefficients estimated using the methodology outlined in Callaway and Sant'Anna (2021). The outcome variable is the natural logarithm of total ATM withdrawal volume in rupees measured at the user × month level. Treatment timing is defined as the month in which the user's district first became treated for CBDC marketing campaigns. The x-axis denotes months relative to the district's first month of CBDC marketing campaigns. The control group consists of users in districts that had not yet received CBDC marketing campaigns. All outcome variables are winsorized at the 1% level. The shaded bands represent 95% confidence intervals. The table below the figures reports the estimates of the aggregate average treatment effect on the treated (ATT) by taking a weighted average of all group-time specific, partially aggregated group-specific and partially aggregated event-time specific average treatment effects following Callaway and Sant'Anna (2021). Standard errors clustered at the district level are shown in parentheses. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Figure 8: Kaplan-Meier Estimates of Cumulative CBDC Adoption



(a) Gender (b) Age (c) Income

(d) UPI Experience (e) Education (f) Deposit Insurance Threshold

This figure displays Kaplan-Meier estimates of cumulative CBDC adoption probability across demographic subgroups over the 70-week observation period following the targeted CBDC marketing campaigns. The y-axis shows the cumulative probability of adoption and the x-axis measures the number of weeks since becoming targeted for marketing campaigns. Panel (a) stratifies by gender. Panel (b) stratifies by age quartile, where Q1 represents the youngest age and Q4 represents the oldest age. Panel (c) stratifies by income quartile, where Q1 represents the lowest income and Q4 represents the highest income. Panel (d) stratifies by pre-CBDC UPI usage, distinguishing between users who used UPI at least once before being targeted for CBDC marketing campaigns and those who did not. Panel (e) stratifies by education, distinguishing between users with sophisticated education (computer, finance, or engineering degree) and those without. Panel (f) stratifies by the deposit insurance threshold, where above the deposit insurance threshold is defined as having a median total deposit balance over the six months prior to treatment above ₹500,000. Shaded regions represent 95% confidence intervals using robust standard errors clustered at the district level.

Internet Appendix for:

"Central Bank Digital Currency (CBDC), Bank Deposits, and Digital Payments: First Evidence from a Live Retail CBDC"

Contents

A Institutional Background **A3**

A.1 Unified Payments Interface (UPI) A3

A.2 Central Bank Digital Currency (CBDC) A5

B Data **A8**

C Robustness **A9**

C.1 Effect of CBDC on Aggregate Bank Deposits and Credit A9

C.2 Effect of CBDC on Individual-Level Bank Deposits A13

C.3 Effect of CBDC on Different Types of Digital Transactions A15

C.4 Effect of CBDC on Alternative Transaction Methods A17

D Adoption of CBDC **A26**

D.1 Empirical Strategy A26

D.2 Robustness: Determinants of CBDC Adoption A28

D.3 Characteristics of UPI Users A30

E Conceptual Framework **A31**

E.1 Deposits: Proof of Proposition 4.1 . . . . . A31

E.2 Digital payments: Proofs of Lemma 4.2 and Proposition 4.2 . . . . . A32

E.3 Adoption: Proofs of Lemma 4.3 and Proposition 4.3 . . . . . A32

E.4 Extension: Heterogeneity in Risk and Trust A33

E.5 Extension: Interest-Bearing or Fee-Charging CBDC A35

E.6 Extension: CBDC Holding Caps A35

E.7 Extension: Endogenous Term Deposits and Illiquid Portfolio Choice A37

F Imperfect Treatment, Partial CBDC Exposure, and Bounds under Callaway and Sant'Anna ATT **A42**

G The Costs & Benefits of CBDC **A48**

G.1 Proofs A48

G.2 Term-by-Term Evaluation A49
G.3 Sensitivity Analysis to Calibrated Parameters A50

Appendix A Institutional Background

A.1 Unified Payments Interface (UPI)

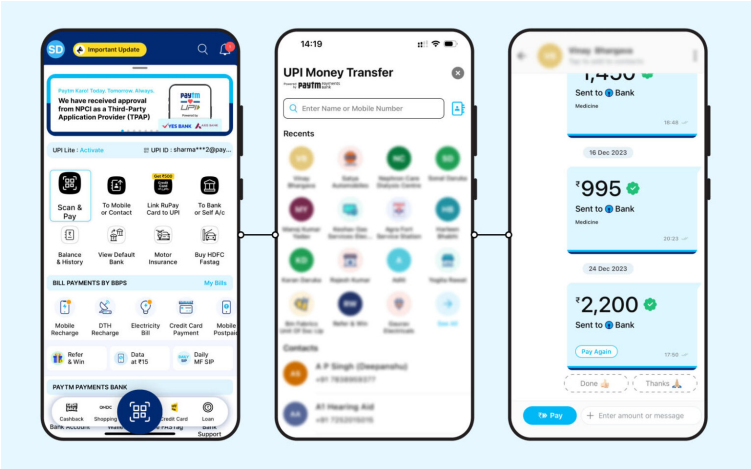
Launched in April 2016, UPI is a single-window mobile payment system developed by the National Payments Corporation of India (NPCI), an initiative of the RBI and the Indian Banks' Association (IBA) that operates as an umbrella organization for retail payments and settlement systems. The new payment system allows both customers and merchants to transfer money between bank accounts safely and at no cost. It eliminates the need to enter bank details or other sensitive information when initiating a transaction. Instead, users simply scan a QR code to transact. UPI transformed digital payments in India by allowing users to link bank accounts to mobile applications and initiate transactions safely. It has since become the dominant retail digital payment system. Appendix Figure A.1 illustrates the UPI transaction flow within a typical payment application, showing the steps involved in initiating a person-to-person transfer: searching for a recipient, selecting a contact, and confirming the payment. Appendix Figure A.2 illustrates a typical merchant point of sale, where QR codes from multiple UPI application providers are displayed side by side.

In fact, UPI is starting to fundamentally change the way Indians transact as adoption is burgeoning (Kearns and Mathew, 2022). In March 2026, it processed over 18 billion transactions with an estimated value of over 24 trillion Rupees, up from just 2 million transactions in December 2016.[16] According to PwC's Indian payments handbook, 75% of digital retail payments in the 2022-23 fiscal year used UPI, and the figure is forecasted to grow to 90% by 2027 (PwC India, 2023). In addition to high local adoption, UPI has gone global with eight countries—Bhutan, France, Mauritius, Nepal, Qatar, Singapore, Sri Lanka, and the UAE—accepting it as a form of payment, with even more countries expected to become available in the future.[17]

UPI's growth has started the process of replacing cash as a method of payment. As of 2023, cash now accounts for less than 60% of total transaction volume, down from 90% in 2017. Furthermore, the RBI states that digital payments now make up more than 40% of all payments in India. High adoption and usage of this new payment method are clear, and it is even forecasted to continue growing in the next several years. The transition towards a digitized and cashless payments system is well underway with the introduction of UPI, but the RBI takes it a step further with the introduction of a CBDC.

[16] See more monthly statistics on UPI usage at the NPCI's website.
[17] For updates on the list of countries, visit the NPCI's website here.

Figure A.1: UPI Transaction Flow



This figure shows the interface of a UPI payment application, illustrating the steps involved in initiating a person-to-person transfer and the transaction history screen. Source: Paytm.

Figure A.2: UPI QR Codes at Point of Sale



This figure shows UPI QR codes displayed at a merchant point of sale, illustrating the widespread acceptance infrastructure for UPI payments across multiple third-party application providers including Google Pay, Paytm, and PhonePe. Source: Nasir Kachroo / NurPhoto / Getty Images.

A.2 Central Bank Digital Currency (CBDC)

Figure A.3: CBDC Marketing Campaigns



(a) Karnataka Bank



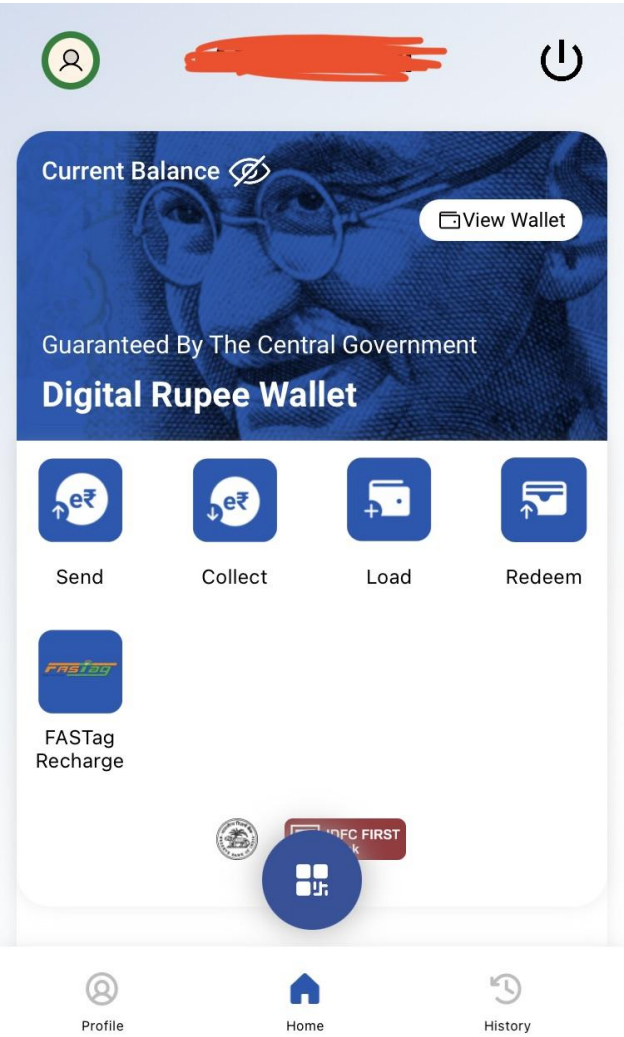
(b) Bank of India



(c) State Bank of India

This figure presents examples of CBDC marketing materials distributed by participating banks, illustrating the type of promotional campaigns whose staggered roll-out across districts constitutes the treatment variable in our individual-level analysis.

Figure A.4: CBDC Wallet Interface



This figure shows the home screen of a CBDC wallet application, displaying the four core transaction functions—Send, Collect, Load, and Redeem—that correspond to the transaction types observed in our CBDC transaction data. Source: screenshot of digital rupee wallet application.

Figure A.5: CBDC QR Code at Point of Sale



This figure shows a Digital Rupee acceptance display at a Reliance Retail merchant location, illustrating the point-of-sale infrastructure for CBDC payments powered by Yes Bank. Source: The Hindu.

Appendix B Data

Table B.1: Summary Statistics of Weekly Transaction Usage

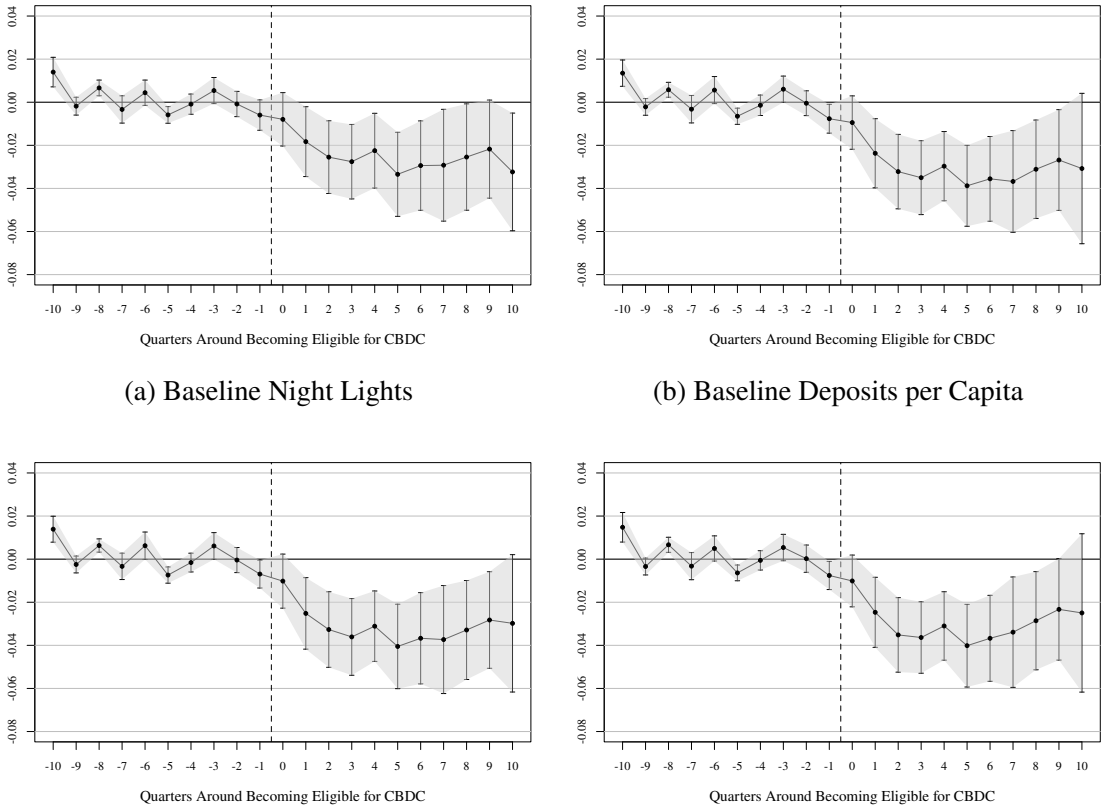
	# Obs	Median	Mean	SD
Log UPI Volume	6,964,860	9.419	7.118	5.158
Log Number of UPI Transactions	6,964,860	1.099	1.244	1.079

This table reports the summary statistics for transaction-level variables at the weekly level. Log UPI Volume is measured as the natural logarithm of the total transaction value in rupees over the relevant period. Log Number of UPI Transactions is measured as the natural logarithm of the total count of transactions.

Appendix C Robustness

C.1 Effect of CBDC on Aggregate Bank Deposits and Credit

Figure C.1: Effect of CBDC Eligibility on Bank Deposits (Conditional on Baseline District Characteristics)



(a) Baseline Night Lights

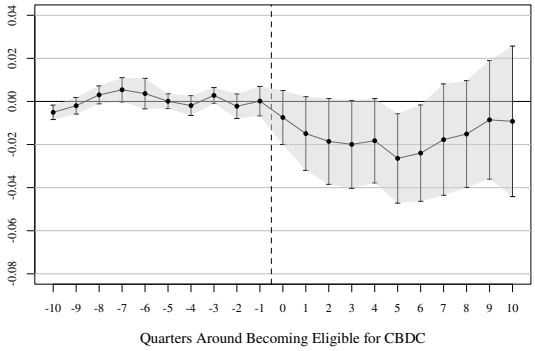
(b) Baseline Deposits per Capita

(c) Baseline Credit per Capita

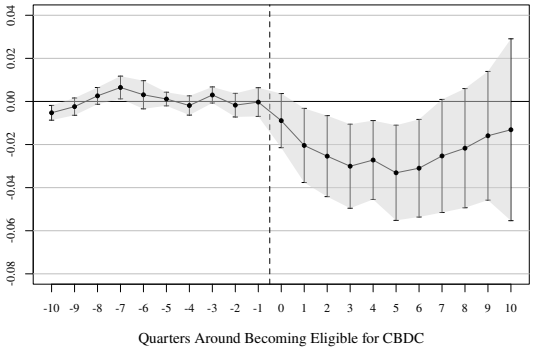
(d) All Covariates

This figure replicates panel (a) of Figure 4 allowing for conditional parallel trends, estimated using the methodology outlined in Callaway and Sant'Anna (2021) with district-level covariates measured in the pre-treatment period. The outcome variable is the natural logarithm of total retail deposits, measured at the district × quarter level. Panel (a) conditions on baseline night light intensity, panel (b) on baseline deposits per capita, panel (c) on baseline credit per capita, and panel (d) on all three jointly. Covariates are fixed at their pre-treatment values and do not vary over time, so that identification requires parallel trends to hold conditional on baseline district characteristics rather than unconditionally. Treatment timing is defined as the quarter in which the lead bank of each district first received RBI authorization to offer CBDC. The x-axis denotes quarters relative to the district's first quarter of CBDC eligibility. The control group consists of districts whose lead bank had not yet received CBDC eligibility within our sample period. All outcome variables are winsorized at the 1% level. Standard errors are clustered at the district level. The shaded bands represent 95% confidence intervals.

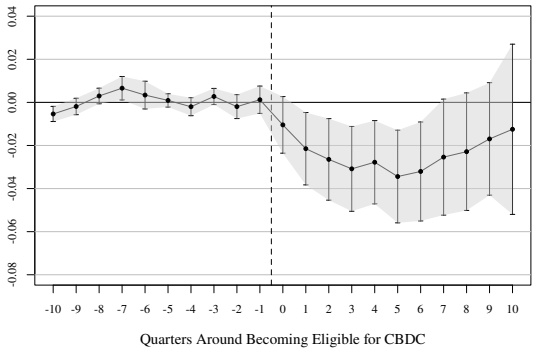
Figure C.2: Effect of CBDC Eligibility on Bank Credit (Conditional on Baseline District Characteristics)



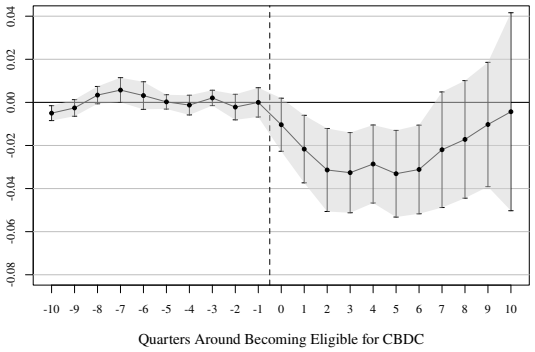
(a) Baseline Night Lights



(b) Baseline Deposits per Capita



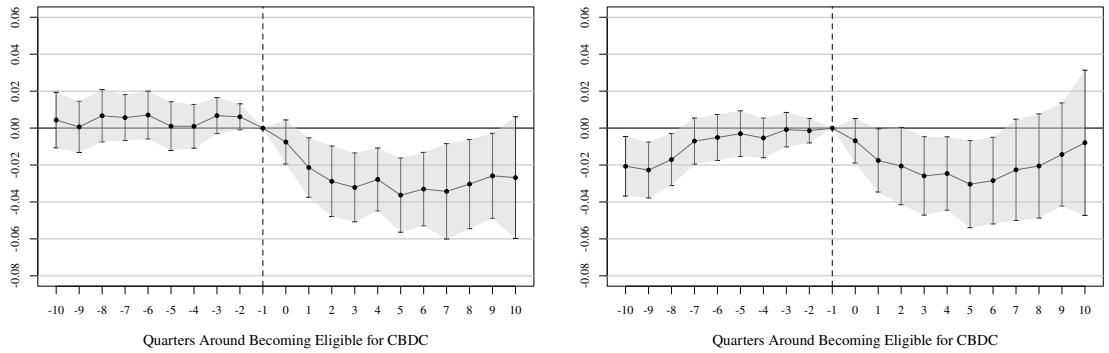
(c) Baseline Credit per Capita



(d) All Covariates

This figure replicates panel (b) of Figure 4 allowing for conditional parallel trends, estimated using the methodology outlined in Callaway and Sant'Anna (2021) with district-level covariates measured in the pre-treatment period. The outcome variable is the natural logarithm of credit, measured at the district × quarter level. Panel (a) conditions on baseline night light intensity, panel (b) on baseline deposits per capita, panel (c) on baseline credit per capita, and panel (d) on all three jointly. Covariates are fixed at their pre-treatment values and do not vary over time, so that identification requires parallel trends to hold conditional on baseline district characteristics rather than unconditionally. Treatment timing is defined as the quarter in which the lead bank of each district first received RBI authorization to offer CBDC. The x-axis denotes quarters relative to the district's first quarter of CBDC eligibility. The control group consists of districts whose lead bank had not yet received CBDC eligibility within our sample period. All outcome variables are winsorized at the 1% level. Standard errors are clustered at the district level. The shaded bands represent 95% confidence intervals.

Figure C.3: Effect of CBDC Eligibility on Bank Deposits and Credit (Callaway Sant'Anna Long Difference)

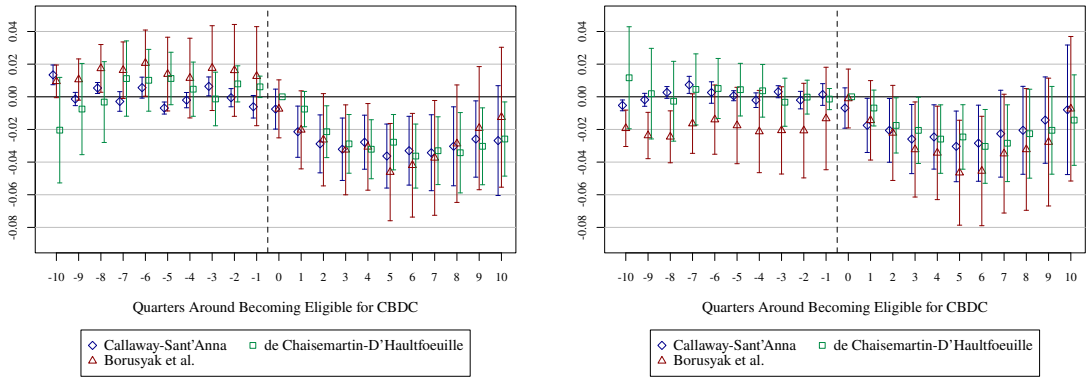


(a) Log Total Retail Deposit

(b) Log Credit

This figure replicates Figure 4 with pre-treatment effects calculated as long differences, making the pre-treatment coefficients comparable to conventional two-way fixed effects event study plots. The outcome variable is the natural logarithm of total retail deposits in panel (a), and credit in panel (b), each measured at the district × quarter level. Treatment timing is defined as the quarter in which the lead bank of each district first received RBI authorization to offer CBDC. The x-axis denotes quarters relative to the district's first quarter of CBDC eligibility. The control group consists of districts whose lead bank had not yet received CBDC eligibility within our sample period. All outcome variables are winsorized at the 1% level. Standard errors are clustered at the district level. The shaded bands represent 95% confidence intervals.

Figure C.4: Effect of CBDC Eligibility on Bank Deposits and Credit (Robustness Using Alternative DiD Estimators)



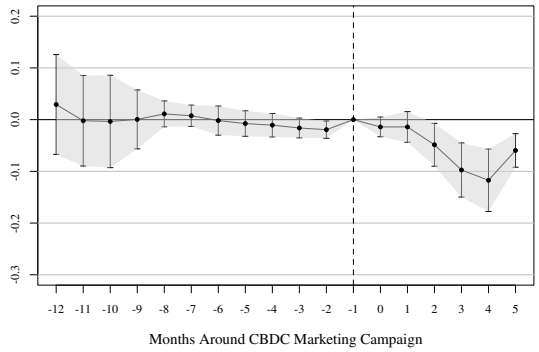
(a) Log Total Retail Deposit

(b) Log Credit

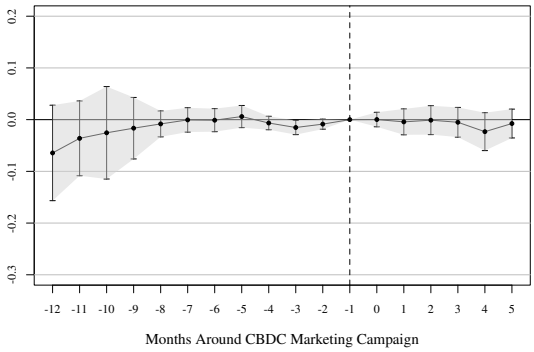
This figure plots the event-study dynamic coefficients for the effect of CBDC eligibility on bank deposits and credit using three alternative estimators: (1) Callaway and Sant'Anna (2021) estimator (in blue with hollow diamond markers), (2) Borusyak, Jaravel, and Spiess (2024) estimator (in red with hollow triangle markers), and (3) de Chaisemartin, D'Haultfœuille, and Vazquez-Bare (2024) estimator (in green with hollow square markers). The outcome variable is the natural logarithm of total retail deposits in panel (a), and credit in panel (b), each measured at the district × quarter level. Treatment timing is defined as the quarter in which the lead bank of each district first received RBI authorization to offer CBDC. The x-axis denotes quarters relative to the district's first quarter of CBDC eligibility. The control group consists of districts whose lead bank had not yet received CBDC eligibility within our sample period. All outcome variables are winsorized at the 1% level. Standard errors are clustered at the district level. The bars represent 95% confidence intervals.

C.2 Effect of CBDC on Individual-Level Bank Deposits

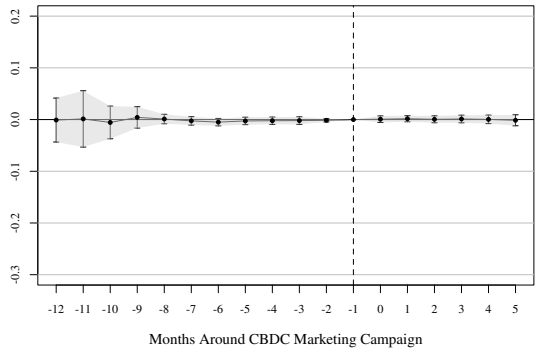
Figure C.5: Monthly Effect of CBDC Marketing Campaign on Deposit Balances (Callaway Sant'Anna Long Difference)



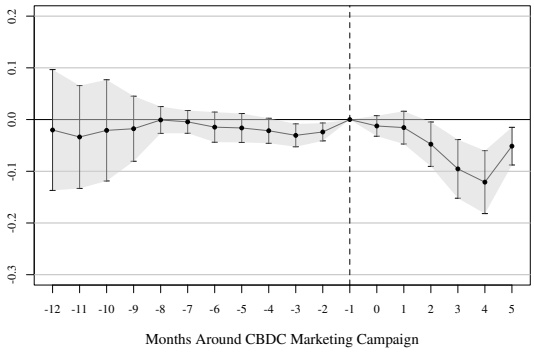
(a) Log Saving Balance



(b) Log Fixed Deposit Balance



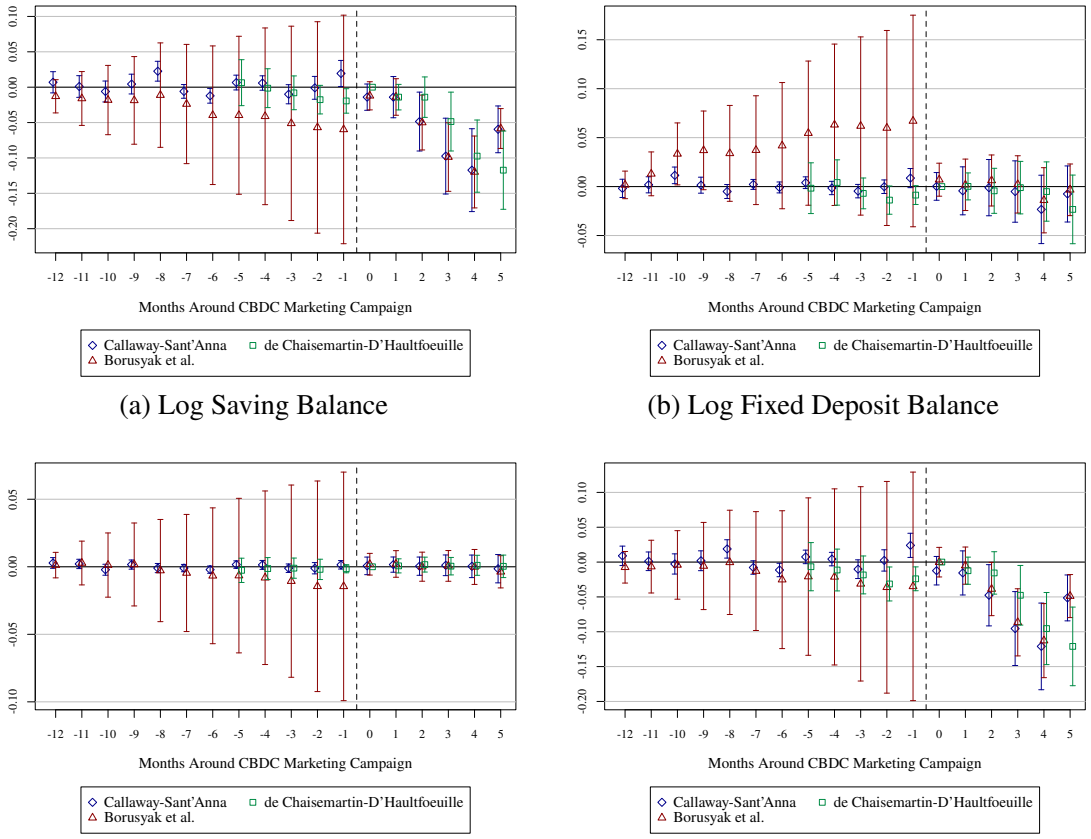
(c) Log Recurring Deposit Balance



(d) Log Total Deposit Balance

This figure replicates Figure 5 with pre-treatment effects calculated as long differences, making the pre-treatment coefficients comparable to conventional two-way fixed effects event study plots. The outcome variable is the natural logarithm of a user's month average savings balance in rupees in panel (a), the natural logarithm of a user's month end fixed deposit balance in panel (b), the natural logarithm of a user's month end recurring deposit balance in panel (c), and the natural logarithm of a user's total deposit balance, defined as the sum of savings, fixed deposit, and recurring deposit, in panel (d), each measured at the user × month level. Treatment timing is defined as the month in which the user's district first became treated for CBDC marketing campaigns. The x-axis denotes months relative to the district's first month of CBDC marketing campaigns. The control group consists of users in districts that had not yet received CBDC marketing campaigns. All outcome variables are winsorized at the 1% level. The shaded bands represent 95% confidence intervals. Standard errors are clustered at the district level.

Figure C.6: Monthly Effect of CBDC Marketing Campaign on Deposit Balances (Robustness Using Alternative DiD Estimators)



(a) Log Saving Balance

(b) Log Fixed Deposit Balance

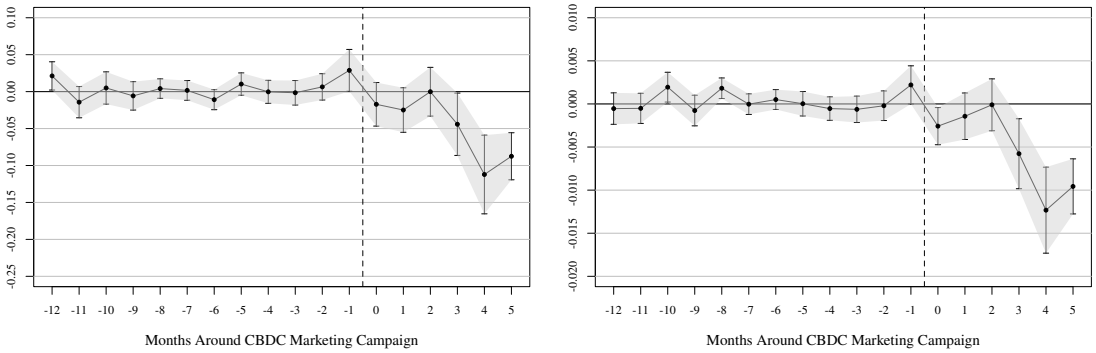
(c) Log Recurring Deposit Balance

(d) Log Total Deposit Balance

This figure plots the event-study dynamic coefficients for the effect of CBDC eligibility on deposit balances using three alternative estimators: (1) Callaway and Sant'Anna (2021) estimator (in blue with hollow diamond markers), (2) Borusyak, Jaravel, and Spiess (2024) estimator (in red with hollow triangle markers), and (3) de Chaisemartin, D'Haultfœuille, and Vazquez-Bare (2024) estimator (in green with hollow square markers). The outcome variable is the natural logarithm of a user's month average savings balance in rupees in panel (a), the natural logarithm of a user's month end fixed deposit balance in panel (b), the natural logarithm of a user's month end recurring deposit balance in panel (c), and the natural logarithm of a user's total deposit balance, defined as the sum of savings, fixed deposit, and recurring deposit, in panel (d), each measured at the user × month level. Treatment timing is defined as the month in which the user's district first became treated for CBDC marketing campaigns. The x-axis denotes months relative to the district's first month of CBDC marketing campaigns. The control group consists of users in districts that had not yet received CBDC marketing campaigns. All outcome variables are winsorized at the 1% level. The shaded bands represent 95% confidence intervals. Standard errors are clustered at the district level.

C.3 Effect of CBDC on Different Types of Digital Transactions

Figure C.7: Monthly Effect of CBDC Marketing Campaign on P2P UPI Usage



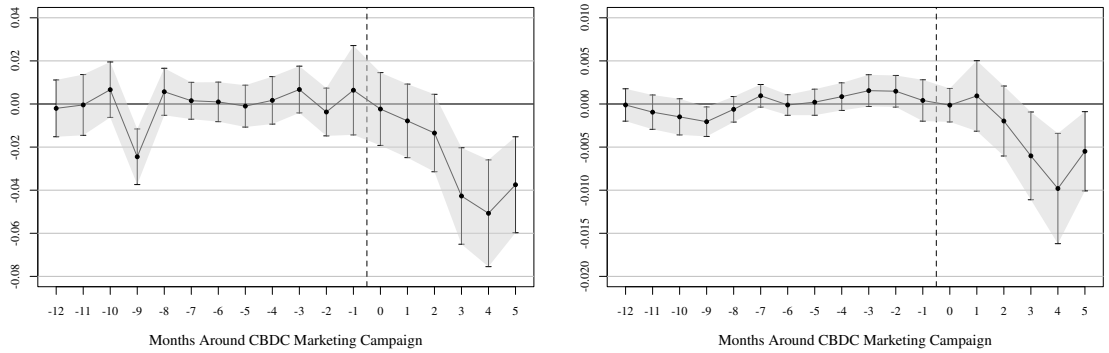
(a) Log P2P UPI Volume

(b) Log Number of P2P UPI Transactions

	ATT Estimates		
	Simple	Group-Specific	Event Study
Log P2P UPI Volume	−0.0462***	−0.0477***	−0.0477***
	(0.0145)	(0.0161)	(0.0143)
Log Number of P2P UPI Transactions	−0.0051***	−0.0054***	−0.0053***
	(0.0012)	(0.0014)	(0.0013)

This figure plots the event-study dynamic coefficients estimated using the methodology outlined in Callaway and Sant'Anna (2021). The outcome variable is the natural logarithm of total person-to-person (P2P) UPI transaction volume in rupees in panel (a), and the natural logarithm of the total number of P2P UPI transactions in panel (b), each measured at the user × month level. Treatment timing is defined as the month in which the user's district first became treated for CBDC marketing campaigns. The x-axis denotes months relative to the district's first month of CBDC marketing campaigns. The control group consists of users in districts that had not yet received CBDC marketing campaigns. All outcome variables are winsorized at the 1% level. The shaded bands represent 95% confidence intervals. The table below the figures reports the estimates of the aggregate average treatment effect on the treated (ATT) by taking a weighted average of all group-time specific, partially aggregated group-specific and partially aggregated event-time specific average treatment effects following Callaway and Sant'Anna (2021). Standard errors clustered at the district level are shown in parentheses. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Figure C.8: Monthly Effect of CBDC Marketing Campaign on P2M UPI Usage



(a) Log P2M UPI Volume

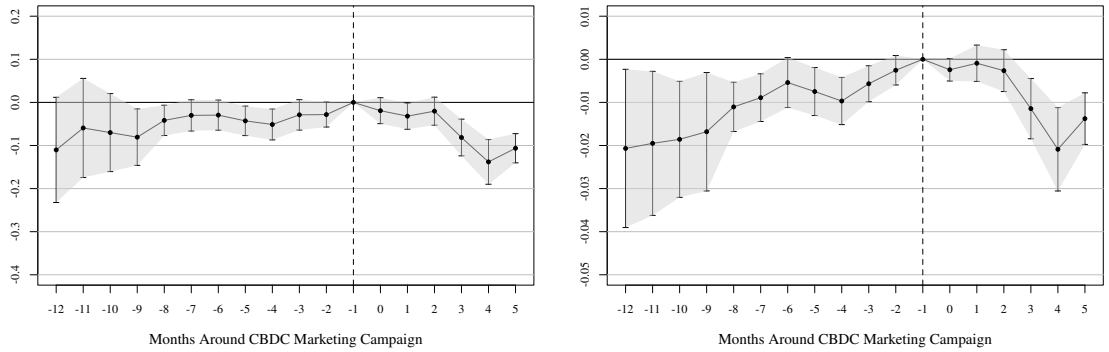
(b) Log Number of P2M UPI Transactions

	ATT Estimates		
	Simple	Group-Specific	Event Study
Log P2M UPI Volume	-0.0253^{***}	-0.026^{***}	-0.0258^{***}
	(0.0074)	(0.0080)	(0.0078)
Log Number of P2M UPI Transactions	-0.0037^{**}	-0.0038^{**}	-0.0037^{**}
	(0.0017)	(0.0019)	(0.0017)

This figure plots the event-study dynamic coefficients estimated using the methodology outlined in Callaway and Sant'Anna (2021). The outcome variable is the natural logarithm of total person-to-merchant (P2M) UPI transaction volume in rupees in panel (a), and the natural logarithm of the total number of P2M UPI transactions in panel (b), each measured at the user × month level. Treatment timing is defined as the month in which the user's district first became treated for CBDC marketing campaigns. The x-axis denotes months relative to the district's first month of CBDC marketing campaigns. The control group consists of users in districts that had not yet received CBDC marketing campaigns. All outcome variables are winsorized at the 1% level. The shaded bands represent 95% confidence intervals. The table below the figures reports the estimates of the aggregate average treatment effect on the treated (ATT) by taking a weighted average of all group-time specific, partially aggregated group-specific and partially aggregated event-time specific average treatment effects following Callaway and Sant'Anna (2021). Standard errors clustered at the district level are shown in parentheses. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

C.4 Effect of CBDC on Alternative Transaction Methods

Figure C.9: Monthly Effect of CBDC Marketing Campaign on UPI Usage (Callaway Sant'Anna Long Difference)

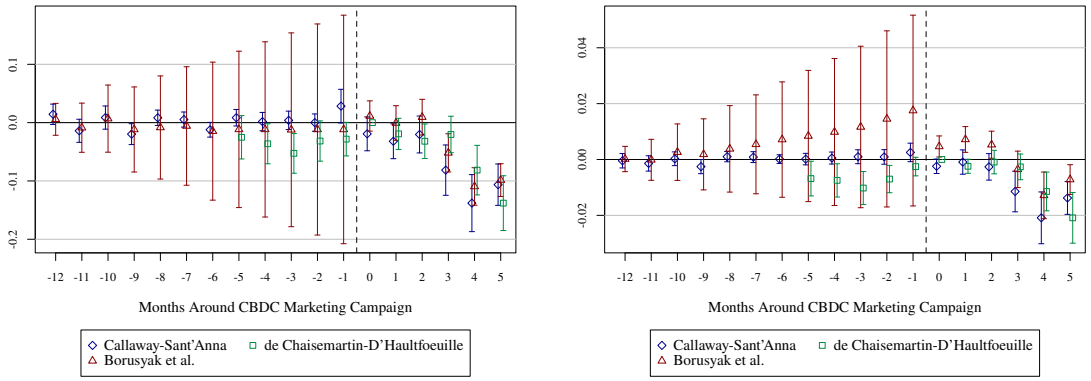


(a) Log UPI Volume

(b) Log Number of UPI Transactions

This figure replicates Figure 6 with pre-treatment effects calculated as long differences, making the pre-treatment coefficients comparable to conventional two-way fixed effects event study plots. The outcome variable is the natural logarithm of total UPI transaction volume in rupees in panel (a), and the natural logarithm of the total number of UPI transactions in panel (b), each measured at the user × month level. Treatment timing is defined as the month in which the user's district first became treated for CBDC marketing campaigns. The x-axis denotes months relative to the district's first month of CBDC marketing campaigns. The control group consists of users in districts that had not yet received CBDC marketing campaigns. All outcome variables are winsorized at the 1% level. The shaded bands represent 95% confidence intervals. Standard errors are clustered at the district level.

Figure C.10: Monthly Effect of CBDC Marketing Campaign on UPI Usage (Robustness Using Alternative DiD Estimators)

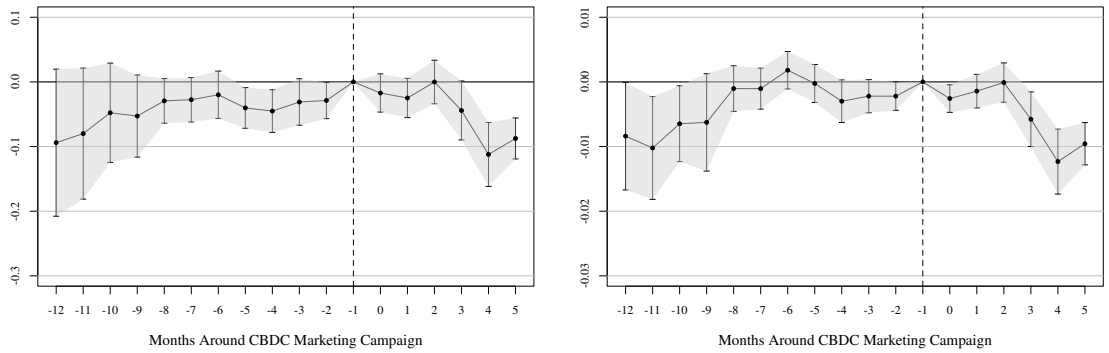


(a) Log UPI Volume

(b) Log Number of UPI Transactions

This figure plots the event-study dynamic coefficients for the effect of CBDC eligibility on UPI usage using three alternative estimators: (1) Callaway and Sant'Anna (2021) estimator (in blue with hollow diamond markers), (2) Borusyak, Jaravel, and Spiess (2024) estimator (in red with hollow triangle markers), and (3) de Chaisemartin, D'Haultfœuille, and Vazquez-Bare (2024) estimator (in green with hollow square markers). The outcome variable is the natural logarithm of total UPI transaction volume in rupees in panel (a), and the natural logarithm of the total number of UPI transactions in panel (b), each measured at the user × month level. Treatment timing is defined as the month in which the user's district first became treated for CBDC marketing campaigns. The x-axis denotes months relative to the district's first month of CBDC marketing campaigns. The control group consists of users in districts that had not yet received CBDC marketing campaigns. All outcome variables are winsorized at the 1% level. The bars represent 95% confidence intervals. Standard errors are clustered at the district level.

Figure C.11: Monthly Effect of CBDC Marketing Campaign on UPI P2P Usage (Callaway Sant'Anna Long Difference)

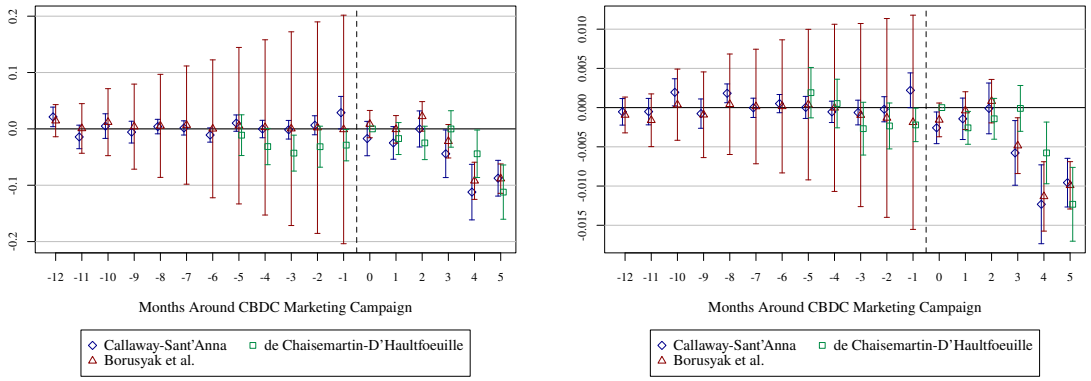


(a) Log P2P UPI Volume

(b) Log Number of P2P UPI Transactions

This figure replicates Figure C.7 with pre-treatment effects calculated as long differences, making the pre-treatment coefficients comparable to conventional two-way fixed effects event study plots. The outcome variable is the natural logarithm of total person-to-person (P2P) UPI transaction volume in rupees in panel (a), and the natural logarithm of the total number of P2P UPI transactions in panel (b), each measured at the user × month level. Treatment timing is defined as the month in which the user's district first became treated for CBDC marketing campaigns. The x-axis denotes months relative to the district's first month of CBDC marketing campaigns. The control group consists of users in districts that had not yet received CBDC marketing campaigns. All outcome variables are winsorized at the 1% level. The shaded bands represent 95% confidence intervals. Standard errors are clustered at the district level.

Figure C.12: Monthly Effect of CBDC Marketing Campaign on UPI P2P Usage (Robustness Using Alternative DiD Estimators)

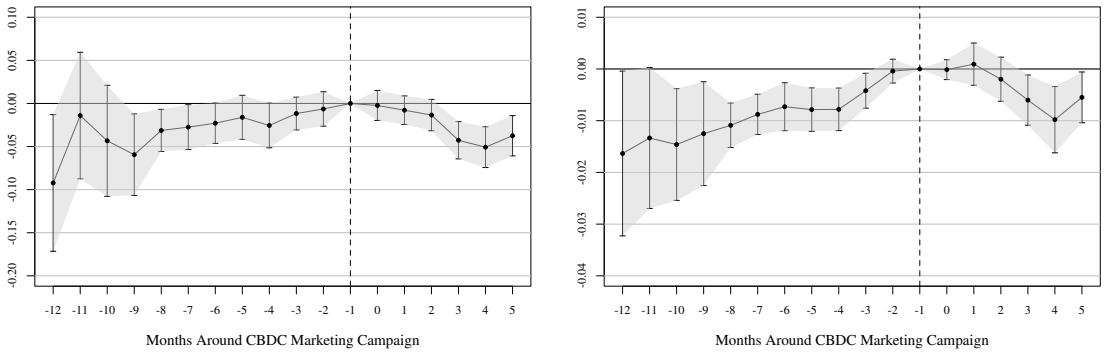


(a) Log P2P UPI Volume

(b) Log Number of P2P UPI Transactions

This figure plots the event-study dynamic coefficients for the effect of CBDC eligibility on person-to-person (P2P) UPI usage using three alternative estimators: (1) Callaway and Sant'Anna (2021) estimator (in blue with hollow diamond markers), (2) Borusyak, Jaravel, and Spiess (2024) estimator (in red with hollow triangle markers), and (3) de Chaisemartin, D'Haultfœuille, and Vazquez-Bare (2024) estimator (in green with hollow square markers). The outcome variable is the natural logarithm of total P2P UPI transaction volume in rupees in panel (a), and the natural logarithm of the total number of P2P UPI transactions in panel (b), each measured at the user × month level. Treatment timing is defined as the month in which the user's district first became treated for CBDC marketing campaigns. The x-axis denotes months relative to the district's first month of CBDC marketing campaigns. The control group consists of users in districts that had not yet received CBDC marketing campaigns. All outcome variables are winsorized at the 1% level. The bars represent 95% confidence intervals. Standard errors are clustered at the district level.

Figure C.13: Monthly Effect of CBDC Marketing Campaign on UPI P2M Usage (Callaway Sant'Anna Long Difference)

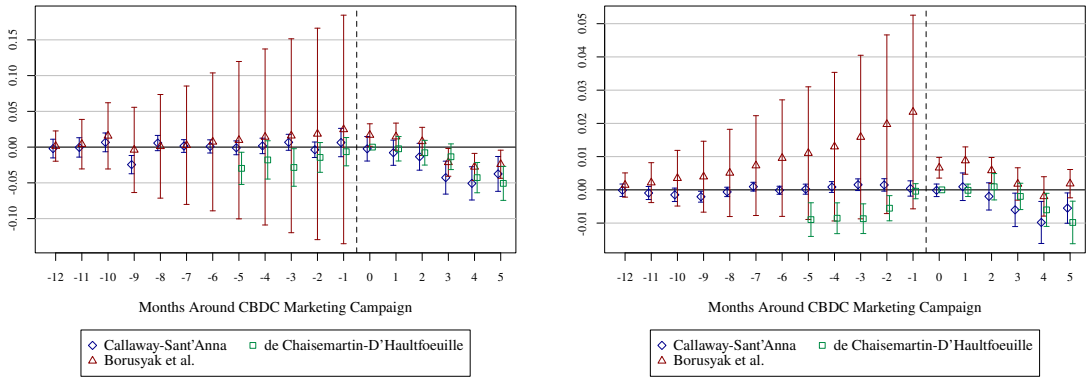


(a) Log P2M UPI Volume

(b) Log Number of P2M UPI Transactions

This figure replicates Figure C.8 with pre-treatment effects calculated as long differences, making the pre-treatment coefficients comparable to conventional two-way fixed effects event study plots. The outcome variable is the natural logarithm of total person-to-merchant (P2M) UPI transaction volume in rupees in panel (a), and the natural logarithm of the total number of P2M UPI transactions in panel (b), each measured at the user × month level. Treatment timing is defined as the month in which the user's district first became treated for CBDC marketing campaigns. The x-axis denotes months relative to the district's first month of CBDC marketing campaigns. The control group consists of users in districts that had not yet received CBDC marketing campaigns. All outcome variables are winsorized at the 1% level. The shaded bands represent 95% confidence intervals. Standard errors are clustered at the district level.

Figure C.14: Monthly Effect of CBDC Marketing Campaign on UPI P2M Usage (Robustness Using Alternative DiD Estimators)

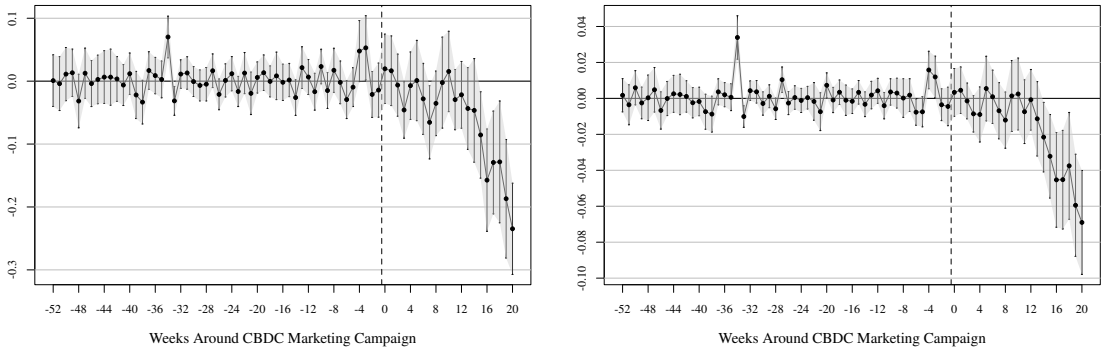


(a) Log P2M UPI Volume

(b) Log Number of P2M UPI Transactions

This figure plots the event-study dynamic coefficients for the effect of CBDC eligibility on person-to-merchant (P2M) UPI usage using three alternative estimators: (1) Callaway and Sant'Anna (2021) estimator (in blue with hollow diamond markers), (2) Borusyak, Jaravel, and Spiess (2024) estimator (in red with hollow triangle markers), and (3) de Chaisemartin, D'Haultfœuille, and Vazquez-Bare (2024) estimator (in green with hollow square markers). The outcome variable is the natural logarithm of total P2M UPI transaction volume in rupees in panel (a), and the natural logarithm of the total number of P2M UPI transactions in panel (b), each measured at the user × month level. Treatment timing is defined as the month in which the user's district first became treated for CBDC marketing campaigns. The x-axis denotes months relative to the district's first month of CBDC marketing campaigns. The control group consists of users in districts that had not yet received CBDC marketing campaigns. All outcome variables are winsorized at the 1% level. The bars represent 95% confidence intervals. Standard errors are clustered at the district level.

Figure C.15: Weekly Effect of CBDC Marketing Campaign on UPI Usage



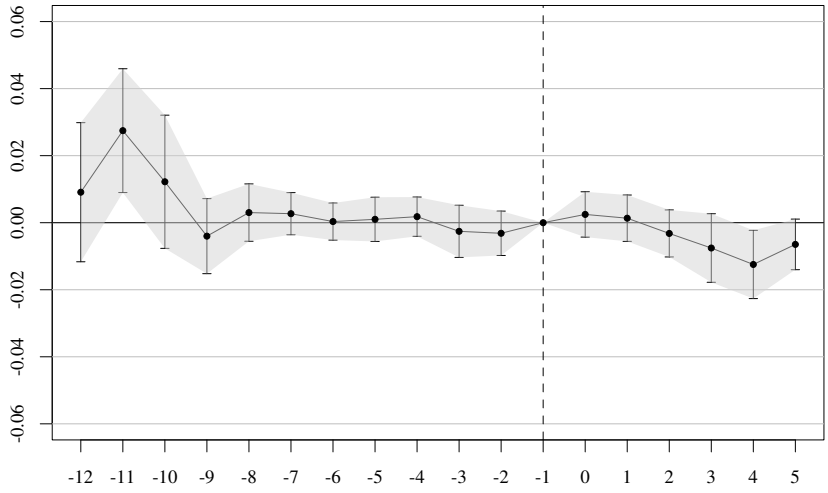
(a) Log UPI Volume

(b) Log Number of UPI Transactions

	ATT Estimates		
	Simple	Group-Specific	Event Study
Log UPI Volume	−0.0572**	−0.0572**	−0.0572**
	(0.0260)	(0.0267)	(0.0274)
Log Number of UPI Transactions	−0.0167**	−0.0167**	−0.0167**
	(0.0075)	(0.0081)	(0.0078)

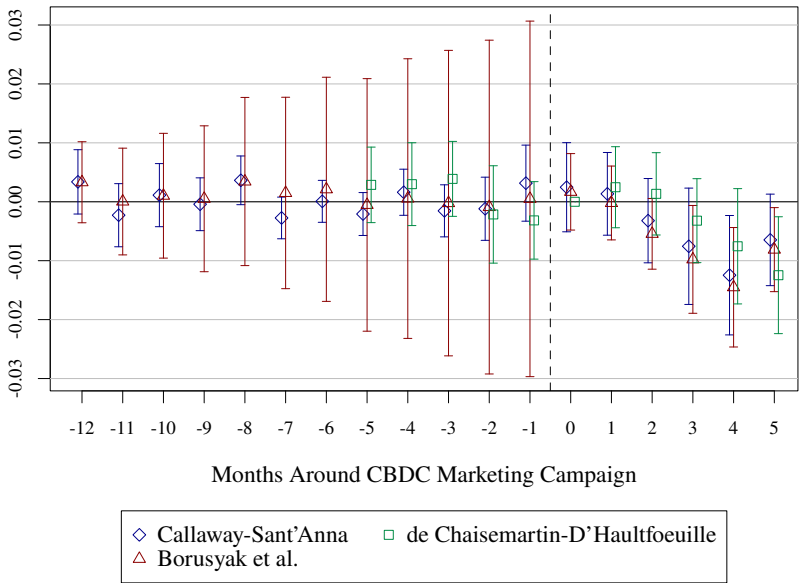
This figure plots the event-study dynamic coefficients estimated using the methodology outlined in Callaway and Sant'Anna (2021). The outcome variable is the natural logarithm of total UPI transaction volume in rupees in panel (a), and the natural logarithm of the total number of UPI transactions in panel (b), each measured at the user × week level. Treatment timing is defined as the week in which the user's district first became treated for CBDC marketing campaigns. The x-axis denotes weeks relative to the district's first week of CBDC marketing campaigns. The control group consists of users in districts that had not yet received CBDC marketing campaigns. All outcome variables are winsorized at the 1% level. The shaded bands represent 95% confidence intervals. The table below the figures reports the estimates of the aggregate average treatment effect on the treated (ATT) by taking a weighted average of all group-time specific, partially aggregated group-specific and partially aggregated event-time specific average treatment effects following Callaway and Sant'Anna (2021). Standard errors clustered at the district level are shown in parentheses. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Figure C.16: Monthly Effect of CBDC Marketing Campaign on ATM Withdrawals (Callaway Sant'Anna Long Difference)



This figure replicates Figure 7 with pre-treatment effects calculated as long differences, making the pre-treatment coefficients comparable to conventional two-way fixed effects event study plots. The outcome variable is the natural logarithm of total ATM withdrawal volume in rupees measured at the user × month level. Treatment timing is defined as the month in which the user's district first became treated for CBDC marketing campaigns. The x-axis denotes months relative to the district's first month of CBDC marketing campaigns. The control group consists of users in districts that had not yet received CBDC marketing campaigns. All outcome variables are winsorized at the 1% level. The shaded bands represent 95% confidence intervals. Standard errors are clustered at the district level.

Figure C.17: Monthly Effect of CBDC Marketing Campaign on ATM Withdrawals (Robustness Using Alternative DiD Estimators)



This figure plots the event-study dynamic coefficients for the effect of CBDC eligibility on ATM withdrawals using three alternative estimators: (1) Callaway and Sant'Anna (2021) estimator (in blue with hollow diamond markers), (2) Borusyak, Jaravel, and Spiess (2024) estimator (in red with hollow triangle markers), and (3) de Chaisemartin, D'Haultfœuille, and Vazquez-Bare (2024) estimator (in green with hollow square markers). The outcome variable is the natural logarithm of total ATM withdrawal volume in rupees measured at the user × month level. Treatment timing is defined as the month in which the user's district first became treated for CBDC marketing campaigns. The *x*-axis denotes months relative to the district's first month of CBDC marketing campaigns. The control group consists of users in districts that had not yet received CBDC marketing campaigns. All outcome variables are winsorized at the 1% level. The bars represent 95% confidence intervals. Standard errors are clustered at the district level.

Appendix D Adoption of CBDC

D.1 Empirical Strategy

We estimate the effect that user characteristics have on the probability of that individual being a CBDC adopter. Since we only observe CBDC transactions until March 2024, for the individuals who did not use the CBDC during our sample period, we do not observe the exact time when they adopted the CBDC, if they did so. Thus, our CBDC adoption data is right-censored, and using a standard linear probability model or logistic regression model to estimate the effect of user demographics on their likelihood to adopt the CBDC would not be valid since these models cannot distinguish between individuals who never adopted and those who had not yet adopted by the end of the sample period. Instead, we employ two survival analysis models to estimate the effect: the Cox proportional hazards model (Cox, 1972) to determine the effect of user characteristics on CBDC adoption and the Kaplan-Meier estimator (Kaplan and Meier, 1958) to estimate the likelihood to adopt CBDC for different subsamples.

In survival analysis, the outcome variable is the time to an event, referred to as *survival time*, and the goal is to estimate the probability that an individual "survives" without experiencing the event beyond a certain time t, referred to as the survival function $S(t) = \mathbb{P}[T > t]$, where T is a random variable representing an individual's true survival time. In our analysis, the outcome variable is the number of weeks to CBDC adoption, and we measure the probability that an individual has not adopted the CBDC by some time t. For right-censored individuals, the survival time is the number of weeks from when they were targeted by CBDC marketing campaigns to the end of March 2024. The Kaplan-Meier estimator is a non-parametric method that estimates this probability for each time t, which allows us to compare adoption rates across subgroups. The Cox proportional hazards model (Cox PH model) focuses, instead, on the hazard rate $h(t)$, which is defined as the instantaneous risk of an event at time t given that the subject has survived until time t. Formally, the survival function S and the hazard rate h are related as follows

$$S(t) = \exp\left(-\int_0^t h(u)du\right). \tag{D.1}$$

Intuitively, the hazard rate measures how quickly individuals are adopting CBDC at any given point in time. The Cox PH model estimates the effect of covariates on the hazard rate while leaving the baseline hazard unspecified. Formally, the model specifies the hazard rate for individual i at time t as

$$h(t \mid \mathbf{x}_i) = h_0(t)\exp(\mathbf{x}_i^\top \boldsymbol{\beta}), \tag{D.2}$$

where $h_0(t)$ is the unspecified baseline hazard at time t, $\mathbf{x}_i$ is a vector of individual-level covariates,

and $\boldsymbol{\beta}$ is a vector of coefficients that will be estimated. To interpret the estimated coefficients, consider two individuals i and j with covariate vectors $\mathbf{x}_i$ and $\mathbf{x}_j$. The ratio of their hazard rates at any time t is

$$\frac{h(t \mid \mathbf{x}_i)}{h(t \mid \mathbf{x}_j)} = \frac{h_0(t)\exp(\mathbf{x}_i^\top \boldsymbol{\beta})}{h_0(t)\exp(\mathbf{x}_j^\top \boldsymbol{\beta})} = \exp\left((\mathbf{x}_i - \mathbf{x}_j)^\top \boldsymbol{\beta}\right). \tag{D.3}$$

The resulting quantity is the *hazard ratio*. For a single covariate X_k, differing by one unit between the two individuals, the hazard ratio simplifies to $\exp(\beta_k)$, which represents the multiplicative effect of a one-unit increase in covariate X_k on the hazard rate, holding all other covariates fixed. The coefficient β_k is the log-hazard ratio of covariate k. The key assumption underlying this derivation is *proportional hazards*: the ratio of hazard rates between any two individuals is constant over time, which is what allows the hazard ratio $\exp(\beta_k)$ to serve as a time-invariant summary of the effect of the covariate on the likelihood of adopting the CBDC, conditional on not having adopted it yet. In other words, one can interpret the estimated hazard ratio $\exp(\hat{\beta}_k)$ of a single covariate k as the multiplicative effect of a one-unit change in the covariate k on the adoption rate at any given week, conditional on not having adopted the CBDC yet.

We include the following covariates in the Cox model: gender, age quartile, income quartile, an indicator for prior UPI usage, an indicator for sophisticated education, and an indicator for an individual having balances above the deposit insurance threshold. Prior UPI usage is an indicator variable that equals one for the individuals who used UPI at least once prior to becoming treated for the CBDC marketing campaigns during our sample period. The sophisticated education indicator equals one for individuals with a degree in computer science, finance, or engineering, as these individuals are most likely to understand the technical nature of CBDC and may therefore adopt earlier than others. Deposits are the median total deposit balance over the six months prior to CBDC treatment. An individual holds balances above the deposit insurance threshold if their deposits are above ₹500,000. For each of these covariates, we present Kaplan-Meier survival curves estimated separately for each subgroup to descriptively compare the timing of CBDC adoption across groups.

As a complement to the survival analysis for CBDC adoption, we also estimate a linear probability model of UPI usage in Appendix D.3 using the same covariates, except for the prior UPI usage indicator.

D.2 Robustness: Determinants of CBDC Adoption

Table D.1: Determinants of CBDC Adoption: Cox Proportional Hazards Model

	CBDC Adoption								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Male	1.0637***						0.8174***	0.8333***	0.8408***
	(0.0503)						(0.0484)	(0.0491)	(0.0473)
Age Q2		0.9354***					0.2792***	0.2752***	0.2406***
		(0.0443)					(0.0448)	(0.0432)	(0.0440)
Age Q3		0.7091***					−0.0443	−0.0260	−0.0650
		(0.0569)					(0.0568)	(0.0539)	(0.0543)
Age Q4		0.1637**					−0.5898***	−0.5295***	−0.5836***
		(0.0761)					(0.0806)	(0.0774)	(0.0753)
Income Q2			0.1331*				−0.0481	−0.0233	−0.0399
			(0.0718)				(0.0695)	(0.0705)	(0.0716)
Income Q3			0.7327***				0.5153***	0.5055***	0.4566***
			(0.0668)				(0.0623)	(0.0623)	(0.0639)
Income Q4			2.4019***				1.8991***	1.8276***	1.7400***
			(0.0620)				(0.0543)	(0.0516)	(0.0538)
Pre CBDC-Eligible UPI User				0.4024***			−0.2104***	−0.2000***	−0.1700***
				(0.0466)			(0.0318)	(0.0312)	(0.0340)
Sophisticated Education					1.8739***		0.8751***	0.7856***	0.6933***
					(0.0705)		(0.0668)	(0.0672)	(0.0828)
Above Insurance Threshold						1.7627***	1.4798***	1.5438***	1.3998***
						(0.0503)	(0.0478)	(0.0516)	(0.0473)
# Obs	153,875	153,875	153,875	153,875	153,875	153,875	153,875	153,875	153,875
Pseudo R^2 Nagelkerke	0.0095	0.0063	0.0445	0.0017	0.0045	0.0261	0.0694	0.0662	0.0519
District FE								Yes	
ZIP code FE									Yes

This table presents coefficients from Cox proportional hazards models estimating the determinants of CBDC adoption. The dependent variable is the time until CBDC adoption, measured in weeks since becoming targeted for CBDC marketing campaigns. Columns (1) through (6) examine the effect of each covariate separately: gender, age quartile, income quartile, prior UPI usage, sophisticated education, and deposit insurance threshold, respectively. Column (7) presents the full specification including all covariates without fixed effects, column (8) adds district fixed effects, and column (9) adds ZIP code fixed effects in place of district fixed effects. Fixed effects are implemented via stratification. The omitted categories are female, age quartile 1 (youngest), annual income quartile 1 (lowest), no prior UPI usage, no sophisticated education, and below deposit insurance threshold. Sophisticated education is defined as holding a computer, finance, or engineering degree. Deposits is the median total deposit balance over the six months prior to CBDC treatment. Above the deposit insurance threshold is defined as having a deposit balance above ₹500,000. All models use robust standard errors clustered at the district level. Positive coefficients indicate higher hazard rates and faster adoption relative to the reference group. For example, the estimate of 1.0637 in column (1) implies that men adopt CBDC at approximately $e^{1.0637} = 2.8970$ times the hazard rate of women, meaning that in any given week, men are adopting at 2.8970x the rate of women, conditional on not having adopted yet. The sample includes all individuals from the start of the CBDC roll-out through the end of March 2024. The Breslow method is used to adjust for ties (individuals adopting at the same time). Standard errors in parentheses are estimated by clustering at the district level. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Table D.2: Determinants of CBDC Adoption: Linear Probability Model

	CBDC Adoption								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Constant	0.0213***	0.0264***	0.0122***	0.0377***	0.0425***	0.0138***	-0.0167***		
	(0.0013)	(0.0012)	(0.0007)	(0.0014)	(0.0013)	(0.0006)	(0.0012)		
Male	0.0389***						0.0279***	0.0296***	0.0306***
	(0.0014)						(0.0013)	(0.0013)	(0.0014)
Age Q2		0.0411***					0.0168***	0.0184***	0.0170***
		(0.0019)					(0.0016)	(0.0016)	(0.0017)
Age Q3		0.0281***					0.0004	0.0024	0.0009
		(0.0023)					(0.0020)	(0.0020)	(0.0021)
Age Q4		0.0056**					-0.0219***	-0.0180***	-0.0203***
		(0.0023)					(0.0022)	(0.0023)	(0.0024)
Income Q2			0.0020**				-0.0047***	-0.0034***	-0.0037***
			(0.0009)				(0.0010)	(0.0011)	(0.0013)
Income Q3			0.0133***				0.0050***	0.0044***	0.0031**
			(0.0012)				(0.0010)	(0.0011)	(0.0012)
Income Q4			0.1182***				0.0964***	0.0940***	0.0909***
			(0.0036)				(0.0029)	(0.0027)	(0.0027)
Pre CBDC-Eligible UPI User				0.0188***			-0.0099***	-0.0088***	-0.0090***
				(0.0022)			(0.0016)	(0.0017)	(0.0018)
Sophisticated Education					0.2125***		0.1565***	0.1534***	0.1476***
					(0.0144)		(0.0138)	(0.0137)	(0.0139)
Above Insurance Threshold						0.0668***	0.0538***	0.0573***	0.0534***
						(0.0028)	(0.0022)	(0.0023)	(0.0024)
# Obs	153,875	153,875	153,875	153,875	153,875	153,875	153,875	153,875	153,875
R^2	0.01	0.01	0.06	0.00	0.01	0.03	0.08	0.10	0.17
District FE								Yes	
ZIP code FE									Yes

This table presents coefficients from a linear probability model estimating the determinants of CBDC adoption. The dependent variable is an indicator variable that is equal to one if the individual adopted CBDC during our sample period. Columns (1) through (6) examine the effect of each covariate separately: gender, age quartile, income quartile, prior UPI usage, sophisticated education, and deposit insurance threshold, respectively. Column (7) presents the full specification including all covariates without fixed effects, column (8) adds district fixed effects, and column (9) adds ZIP code fixed effects in place of district fixed effects. The omitted categories are female, age quartile 1 (youngest), annual income quartile 1 (lowest), no prior UPI usage, no sophisticated education, and below deposit insurance threshold. Sophisticated education is defined as holding a computer, finance, or engineering degree. Deposits is the median total deposit balance over the six months prior to CBDC treatment. Above the deposit insurance threshold is defined as having a deposit balance above ₹500,000. The sample includes all individuals from the start of the CBDC roll-out through the end of March 2024. Standard errors in parentheses are estimated by clustering at the district level. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

D.3 Characteristics of UPI Users

Table D.3: Probability of Being a UPI User by User Characteristics and Education

	UPI User							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Constant	0.4144***	0.6379***	0.4732***	0.4971***	0.4558***	0.5016***		
	(0.0073)	(0.0055)	(0.0058)	(0.0049)	(0.0047)	(0.0078)		
Male	0.1411***					0.1432***	0.1460***	0.1428***
	(0.0057)					(0.0052)	(0.0047)	(0.0049)
Age Q2		-0.0576***				-0.0972***	-0.1070***	-0.1111***
		(0.0047)				(0.0045)	(0.0042)	(0.0044)
Age Q3		-0.1684***				-0.2155***	-0.2228***	-0.2246***
		(0.0049)				(0.0046)	(0.0042)	(0.0043)
Age Q4		-0.3511***				-0.4053***	-0.4062***	-0.4048***
		(0.0063)				(0.0065)	(0.0055)	(0.0057)
Income Q2			0.0167***			0.0195***	0.0109***	0.0092**
			(0.0042)			(0.0037)	(0.0039)	(0.0043)
Income Q3			0.0412***			0.0338***	0.0276***	0.0241***
			(0.0041)			(0.0035)	(0.0034)	(0.0036)
Income Q4			0.0429***			0.0147***	0.0111***	0.0063
			(0.0052)			(0.0038)	(0.0040)	(0.0044)
Sophisticated Education				0.0986***		-0.0039	-0.0150	-0.0164
				(0.0144)		(0.0131)	(0.0125)	(0.0132)
Above Insurance Threshold					0.0922***	0.1504***	0.1468***	0.1372***
					(0.0053)	(0.0043)	(0.0043)	(0.0041)
# Obs	153,875	153,875	153,875	153,875	153,875	153,875	153,875	153,875
R^2	0.02	0.07	0.00	0.00	0.01	0.12	0.14	0.21
District FE							Yes	
ZIP code FE								Yes

This table presents coefficients from a linear probability model estimating the determinants of being a UPI user based on self-reported gender, age, income, education, and deposit balance. Columns (1) through (5) examine the effect of each covariate separately: gender, age quartile, income quartile, sophisticated education, and deposit insurance threshold, respectively. Column (6) presents the full specification including all covariates without fixed effects, column (7) adds district fixed effects, and column (8) adds ZIP code fixed effects in place of district fixed effects. The omitted categories are female, age quartile 1 (youngest), annual income quartile 1 (lowest), no sophisticated education, and below deposit insurance threshold. Sophisticated education is defined as holding a computer, finance, or engineering degree. Deposits is the median total deposit balance over the six months prior to CBDC treatment. Above the deposit insurance threshold is defined as having a deposit balance above ₹500,000. Standard errors in parentheses are estimated by clustering at the district level. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Appendix E Conceptual Framework

This section lays out the conceptual framework that underpins our empirical analysis and supplies the formal arguments for the results stated in Section 4. We begin by presenting the formal proofs of the main propositions from that section. We then extend the framework to incorporate heterogeneity in risk and trust and to study alternative CBDC designs, including interest-bearing CBDC, holding caps, and endogenous term-deposit choices. For ease of reference, we keep the notation and timing from the main text and restate only the elements that are needed to follow the arguments below.

For completeness, we briefly restate the timing and risk structure. At the beginning of the period, households observe their wealth W_i and the CBDC design, choose portfolios (D_i^S, D_i^T, H_i, C_i), and, if applicable, whether to adopt CBDC. At the end of the period, bank default is realized and deposits pay according to the insurance scheme in equation (2), while CBDC and cash always pay par. Deposit insurance applies jointly to total deposits $D_i^{\text{tot}} = D_i^S + D_i^T$ up to the limit $\bar{D}$, with expected loss rate δ per rupee of uninsured deposits.

E.1 Deposits: Proof of Proposition 4.1

Proof of Lemma 4.1. The household chooses (D^S, D^T, H) to maximize $M(D^S, D^T, H, 0)$ subject to $D^S + D^T + H = W$. Given $r^T > r^S$ and that term deposits are illiquid only within the period, any excess wealth beyond the level of liquid balances needed for transactions is strictly better invested in D^T than in D^S or H. Therefore, for any optimal choice of total liquid balances $L_0 = D_0^S + H_0$, all remaining wealth $W - L_0$ is placed in term deposits, proving (i). If L_0 is uniquely optimal, then as W rises, L_0 stays fixed and $D_0^T = W - L_0$ increases, proving (ii). □

Proof of Proposition 4.1. Let $(\tilde{D}^S, \tilde{D}^T, \tilde{H}, \tilde{C})$ be an optimal portfolio with $\tilde{D}^T < D_0^T$, and write $L_0 = D_0^S + H_0$ and $\Delta = D_0^T - \tilde{D}^T > 0$.

We first show $\tilde{D}^S \geq \Delta$. Since $\tilde{D}^T < D_0^T$, term deposits can be increased, so optimality requires that shifting a rupee from the CBDC wallet into term deposits not raise utility. That shift gains at least $\min\{r^T, (1+r^T)(1-\delta)-1\}$ and costs $\phi_i S'(\tilde{C})$, so $\phi_i S'(\tilde{C}) \geq \min\{r^T, (1+r^T)(1-\delta)-1\} \geq \phi_i S'(L_0)$ by Assumption 2. Concavity of S then gives $\tilde{C} \leq L_0$, and the same argument applied to cash, which is dominated in return by savings deposits and enters utility only through the fixed cash-intensive transaction need, gives $\tilde{H} \leq H_0$. Hence $\tilde{H} + \tilde{C} \leq L_0$, which by the budget constraint is equivalent to $\tilde{D}^S \geq \Delta$.

Now construct the alternative portfolio

$$D_1^T = D_0^T, \quad D_1^S = \tilde{D}^S - \Delta, \quad H_1 = \tilde{H}, \quad C_1 = \tilde{C},$$

which is feasible by the preceding paragraph and preserves total wealth. By the budget constraint,

$D_1^S + H_1 + C_1 = \tilde{D}^S + \tilde{H} + \tilde{C} - \Delta = L_0$, so the alternative portfolio holds total liquidity fixed at the baseline level while restoring term deposits from $\tilde{D}^T$ to D_0^T. Because the reallocation is financed entirely out of savings deposits, it leaves total deposits D^{tot} unchanged and therefore does not alter the household's uninsured exposure; no condition on δ is required. The only change in expected monetary payoffs is the return differential on the reallocated balances, which is strictly positive since $r^T > r^S$. Moreover, because $H_1 = \tilde{H}$ and $C_1 = \tilde{C}$, the payments block of utility is unchanged: the convenience terms $\mu_i \gamma^{\text{CBDC}}$ and $\mu_i \gamma^{\text{UPI}}$, the segmentation benefit $\phi_i S(C_i)$, and the adoption cost k_i all take the same values as under the candidate portfolio, so both $U_i^{\text{no CBDC}}$ and U_i^{CBDC} are invariant to the construction. By Assumption 1, holding L_0 fixed leaves non-payment utility unchanged as well. Hence any portfolio with $\tilde{D}^T < D_0^T$ is strictly dominated and there exists an optimal CBDC portfolio with $D_1^T = D_0^T$ and CBDC funded from the liquid margin. □

E.2 Digital payments: Proofs of Lemma 4.2 and Proposition 4.2

Proof of Lemma 4.2. For a CBDC adopter with $C_i > 0$, each unit of digital payment can be routed through UPI or CBDC. The per-unit net utilities are

$$u^{\text{UPI}} = \gamma^{\text{UPI}}, \quad u^{\text{CBDC}}(C_i, \phi_i, k_i) = \gamma^{\text{CBDC}} + \frac{\phi_i S(C_i)}{\mu_i} - \frac{k_i}{\mu_i},$$

where we spread the fixed cost k_i over μ_i transactions. It is clear then that u^{CBDC} is weakly increasing in C_i and in ϕ_i, and weakly decreasing in k_i. For any given (C_i, ϕ_i, k_i), the household allocates all inframarginal digital payments to the channel with higher per-unit net utility, so, under the assumption $\gamma^{\text{CBDC}} \approx \gamma^{\text{UPI}}$, the optimal share x_i^{CBDC} is weakly increasing in $u^{\text{CBDC}} - u^{\text{UPI}}$ and inherits its monotonicity properties. For sufficiently large C_i and ϕ_i, $u^{\text{CBDC}} > u^{\text{UPI}}$ and $x_i^{\text{CBDC}} \to 1$. □

Proof of Proposition 4.2. For CBDC adopters, as C_i rises from zero, Lemma 4.2 implies that x_i^{CBDC} increases and x_i^{UPI} decreases, reducing the share of digital payments routed via UPI. Aggregating across adopters, UPI volume and transaction counts fall even if total digital volume is unchanged. Assumption 3 implies that a subset of transactions remains in cash regardless of CBDC adoption, so aggregate cash usage responds only weakly as CBDC spreads. Hence aggregate UPI usage declines sharply relative to cash. □

E.3 Adoption: Proofs of Lemma 4.3 and Proposition 4.3

Proof of Lemma 4.3. Fix (ϕ_i, μ_i, k_i) and suppose $D_{0i}^{\text{tot}} > \bar{D}$. Consider a small increase $dC > 0$ in CBDC funded by a reduction in deposits. As long as $D_i^{\text{tot}} > \bar{D}$, this marginal change reduces uninsured deposits by dC and saves expected loss $\delta(1 + r^S)dC$, so the net interest cost of holding CBDC instead of deposits on this margin is $r^S(1 - \delta)$. Once deposits fall back to $\bar{D}$, the cost becomes r^S. Aggregating

over the relevant range yields an effective cost $\tilde{r}_i(W_i) \in [0, r^S]$ that is weakly decreasing in D_{0i}^{tot}. The household's problem in C is then

$$\max_{0 \le C \le W_i} \{\phi_i S(C) - \tilde{r}_i(W_i) C\},$$

which is strictly concave because S is concave. An optimal $C_i^*(W_i)$ exists and satisfies the usual first-order condition. Comparing $V^1(W_i)$ and $V^0(W_i)$ yields the adoption inequality (7). □

Proof of Proposition 4.3. Household j has $D_{0j}^{\text{tot}} \le \bar{D}$, so its effective cost of CBDC relative to deposits is r^S. Household i has $D_{0i}^{\text{tot}} > \bar{D}$, so by the argument in the proof of Lemma 4.3, its effective cost $\tilde{r}_i(W_i)$ is strictly below r^S. Given the same (ϕ, μ, k), the left-hand side of (7) is the same function of C for both households, while the right-hand side is strictly smaller for i than for j for any given C. If j finds it optimal to adopt, so that (7) holds at C_j^*, then the same inequality holds strictly for i at C_i^*, proving part (i). For part (ii), choose $(\phi, \mu, k, r^S, \delta)$ so that the inequality holds at cost $\tilde{r}_i(W_i)$ but fails at cost r^S, which is possible whenever $\delta > 0$ and k is in an intermediate range. □

E.4 Extension: Heterogeneity in Risk and Trust

In the main analysis, heterogeneity in CBDC adoption arises from differences in wealth W_i, digital usage μ_i, the valuation of safety and segmentation ϕ_i, and adoption costs k_i. Proposition 4.3 focuses on the role of wealth and total deposits. In this subsection we highlight more explicitly how cross-sectional variation in ϕ_i and k_i generates additional adoption heterogeneity for a given level of deposits, which we interpret as capturing differences in risk perceptions and trust in the banking system and in the central bank.

Fix a household i with wealth W_i and baseline portfolio $(D_{0i}^S, D_{0i}^T, H_{0i})$ in the no-CBDC benchmark. Let $D_{0i}^{\text{tot}} = D_{0i}^S + D_{0i}^T$ denote total deposits, and suppose $D_{0i}^{\text{tot}} > \bar{D}$ so that some deposits are uninsured. Conditional on wealth and total deposits, Lemma 4.3 implies that there exists an optimal CBDC balance $C_i^*(W_i)$ that solves

$$\max_{0 \le C \le W_i} \{\phi_i S(C) - \tilde{r}_i(W_i) C\},$$

where $\tilde{r}_i(W_i) \in [0, r^S]$ is the effective net cost of holding CBDC relative to deposits, capturing both interest foregone and the avoided expected loss on uninsured deposits. Adoption is optimal if and only if

$$\phi_i S(C_i^*(W_i)) + \mu_i(\gamma^{\text{CBDC}} - \gamma^{\text{UPI}}) \ge \tilde{r}_i(W_i) C_i^*(W_i) + k_i.$$

Holding $(W_i, D_{0i}^{\text{tot}}, \mu_i)$ fixed, this condition defines a threshold relationship between ϕ_i and k_i.

Proposition E.1 (Risk / trust heterogeneity and adoption)**.** *Fix* $(W_i, D_{0i}^{tot}, \mu_i)$ *with* $D_{0i}^{tot} > \bar{D}$ *and let* $\tilde{r}_i(W_i)$ *be given. There exists a decreasing function* $\Phi_i(k)$ *such that, for any* $k_i \geq 0$*,*

$$\textit{household } i \textit{ adopts CBDC} \quad \Longleftrightarrow \quad \phi_i \geq \Phi_i(k_i).$$

Moreover, for any fixed $\phi_i > 0$*, there exists a corresponding threshold* $K_i(\phi)$ *such that household* i *adopts CBDC if and only if* $k_i \leq K_i(\phi_i)$*.*

Proof. Given $(W_i, D_{0i}^{\text{tot}}, \mu_i)$ and $\tilde{r}_i(W_i)$, the optimal CBDC balance $C_i^*(W_i)$ is strictly increasing in ϕ_i and does not depend on k_i, because k_i enters utility additively and does not affect the marginal trade-off between $S(C)$ and $\tilde{r}_i(W_i)C$. The left-hand side of the adoption inequality,

$$\phi_i S(C_i^*(W_i)) + \mu_i(\gamma^{\text{CBDC}} - \gamma^{\text{UPI}}),$$

is therefore strictly increasing in ϕ_i and independent of k_i, while the right-hand side,

$$\tilde{r}_i(W_i)C_i^*(W_i) + k_i,$$

is strictly increasing in k_i and independent of ϕ_i for a given $C_i^*(W_i)$. By continuity, there exists for each k_i a unique cutoff $\Phi_i(k_i)$ that equates the two sides, with adoption for $\phi_i \geq \Phi_i(k_i)$. The monotonicity properties imply that $\Phi_i(\cdot)$ is decreasing. The statement about $K_i(\phi)$ follows by symmetry, defining for each ϕ_i the maximal k_i for which the inequality holds. □

Proposition E.1 shows that, conditional on wealth and total deposits, adoption is monotone in the household's valuation of safety and segmentation ϕ_i and in the inverse of the adoption cost k_i. In our empirical setting, higher ϕ_i can be interpreted as reflecting greater concern about bank fragility, fraud risk, or cyber risk, and higher trust in the central bank. Lower k_i captures greater financial and digital sophistication, such as having a computer science, finance, or engineering degree, prior experience with digital payments, or greater comfort with mobile apps.

This extension provides a simple theoretical counterpart to the empirical finding in Section 8 that CBDC adoption is higher among households with sophisticated educational backgrounds and among those who are more engaged with digital finance, even after conditioning on income and deposit balances. It also highlights that variation in perceived risk and trust can amplify the wealth-based adoption gradient in Proposition 4.3: among households with similar total deposits, those with higher ϕ_i and lower k_i will adopt CBDC earlier and at higher rates.

E.5 Extension: Interest-Bearing or Fee-Charging CBDC

We briefly consider how the model's implications change if the CBDC design is altered. Suppose CBDC pays gross return $1 + r^C$ and may charge a per-transaction fee $f \geq 0$. The adoption condition (7) then becomes

$$\phi_i S(C_i^*(W_i)) + \mu_i\big(\gamma^{\text{CBDC}} - f - \gamma^{\text{UPI}}\big) \;\geq\; \big(\tilde{r}_i(W_i) - r^C\big)C_i^*(W_i) + k_i,$$

where $\tilde{r}_i(W_i)$ is defined as in Lemma 4.3.

If r^C increases toward r^S, the effective cost of holding CBDC falls, expanding the set of households for which CBDC adoption is optimal and increasing optimal wallet balances $C_i^*(W_i)$. For r^C sufficiently high, CBDC can become a closer substitute for term deposits as well as for savings, potentially weakening Proposition 4.1. Conversely, a positive fee f reduces the net convenience of CBDC relative to UPI and dampens substitution on the intensive margin, holding adoption fixed. In this sense, the non-interest-bearing, fee-free design in our setting represents a lower-bound case for CBDC take-up and its impact on deposits and digital payments.

E.6 Extension: CBDC Holding Caps

In the baseline model, households can hold any nonnegative CBDC balance C_i up to their total wealth W_i. In many policy discussions, however, central banks consider imposing caps on individual CBDC holdings. In fact, in our setting the Reserve Bank of India has placed a cap of ₹100,000 on the amount an individual can hold in a CBDC wallet. CBDC holding caps are often motivated by financial stability concerns: they aim to limit the extent of deposit disintermediation in normal times and to reduce the risk that CBDC becomes a dominant run asset in periods of stress, while still allowing households to use CBDC as a convenient retail payment instrument.

In this subsection we study a simple version of such a design, where the household faces an exogenous cap $\bar{C} > 0$ on the CBDC wallet:

$$0 \leq C_i \leq \min\{W_i, \bar{C}\}.$$

We retain the baseline assumption that CBDC is non-interest-bearing and fee-free ($r^C = 0$, no per-transaction fee), and we keep the rest of the environment unchanged.

The household's portfolio problem is now

$$\max_{D_i^S, D_i^T, H_i, C_i} M_i(D_i^S, D_i^T, H_i, C_i) + \max\{U_i^{\text{no CBDC}}, U_i^{\text{CBDC}}\}$$

subject to the budget constraint (1), the return and insurance structure in (2), and the additional

constraint $C_i \leq \bar{C}$. As before, we focus on cases where the household chooses to adopt CBDC and set $C_i > 0$.

Given an adopting household i, the optimal CBDC balance without a cap, $C_i^*(W_i)$, solves

$$\max_{0 \leq C \leq W_i} \{\phi_i S(C) - \tilde{r}_i(W_i) C\},$$

with adoption condition given by (7). Under the cap, the effective optimal balance becomes

$$\hat{C}_i(W_i) = \min\{C_i^*(W_i), \bar{C}\}.$$

If $C_i^*(W_i) \leq \bar{C}$, the cap is slack and the household's behavior coincides with the baseline solution. If $C_i^*(W_i) > \bar{C}$, the household is constrained at $\hat{C}_i(W_i) = \bar{C}$ and cannot fully implement its unconstrained preferred CBDC position.

Because CBDC continues to be funded from savings deposits as in Proposition 4.1, the cap limits the maximum reallocation from savings into CBDC. For a constrained household with $C_i^*(W_i) > \bar{C}$, the reduction in savings deposits is given by $\bar{C}$ rather than $C_i^*(W_i)$, and the corresponding decline in uninsured exposure is smaller. In this sense, $\bar{C}$ directly bounds the extent of deposit disintermediation and the shift in digital payments away from UPI.

The adoption margin is affected in two ways. First, the effective gain from adoption on the left-hand side of (7) is computed at $\hat{C}_i(W_i)$ rather than $C_i^*(W_i)$:

$$\phi_i S(\hat{C}_i(W_i)) + \mu_i(\gamma^{\text{CBDC}} - \gamma^{\text{UPI}}),$$

which is weakly smaller when $C_i^*(W_i) > \bar{C}$. Second, the interest-cost term on the right-hand side,

$$\tilde{r}_i(W_i)\, \hat{C}_i(W_i) + k_i,$$

is also lower when the cap binds. The net effect is that for households with modest desired CBDC balances (small $C_i^*(W_i)$), the cap is irrelevant, while for households with very strong incentives to hold CBDC (large $C_i^*(W_i)$) the cap truncates their adoption benefit and reduces, but does not eliminate, their incentive to adopt.

Overall, the introduction of a CBDC holding cap has three main implications relative to the baseline model. First, it limits the maximum reallocation out of savings deposits into CBDC for households that would otherwise choose large CBDC balances, thereby capping the extent of deposit disintermediation and the reduction in uninsured exposure for the most exposed households. Second, it attenuates the intensive-margin substitution from UPI to CBDC among these households, because their wallet balances cannot scale proportionally with their desired level of CBDC usage. Third,

because the cap binds only for households with relatively high wealth or high ϕ_i, the qualitative pattern of adoption remains unchanged: CBDC take-up is still concentrated among wealthier and more sophisticated depositors. However, the quantitative impact of their adoption on deposits and digital payments is compressed by the policy limit $\bar{C}$.

E.7 Extension: Endogenous Term Deposits and Illiquid Portfolio Choice

This extension relaxes the baseline assumption that term deposits are fixed over the CBDC adoption horizon. We allow households to choose an illiquid term-deposit position in addition to liquid savings, cash, and CBDC, and we study when the introduction of CBDC induces households to unwind term deposits. The extension is designed to clarify the conditions under which the baseline treatment of term deposits as predetermined is a valid approximation.

Environment and timing. Time is discrete and indexed by $t \in \{0, 1\}$. At date 0, before the introduction of CBDC, the household allocates initial financial wealth W_0 across three assets: a liquid savings deposit D_0^S, an illiquid term deposit D_0^T, and cash holdings H_0. These choices satisfy

$$D_0^S + D_0^T + H_0 \leq W_0.$$

Savings deposits are fully liquid and earn gross return $1 + r^S$. Term deposits earn gross return $1 + r^T$, where $r^T > r^S$, but cannot be freely accessed before maturity. Cash earns zero nominal return. At date 1, after CBDC becomes available, the household decides whether to adopt CBDC and how to allocate liquid resources across savings deposits, cash, and CBDC. Term deposits mature only at the end of date 1. Thus, over the horizon relevant for CBDC adoption and payments, term deposits are illiquid unless withdrawn early at a cost.

To allow for partial reallocation out of term deposits, suppose that at date 1 the household may liquidate a fraction $\omega \in [0, 1]$ of its outstanding term deposit position. Early liquidation incurs a proportional penalty $\kappa \in (0, 1)$, so each unit of term deposits unwound delivers only $(1 - \kappa)$ units of date-1 liquid resources. The remaining share $(1 - \omega)$ is held to maturity and yields gross payoff $(1 + r^T)(1 - \omega)D_0^T$ at the end of the period.

Date-1 liquid resources. Conditional on the portfolio chosen at date 0, the household enters date 1 with pre-CBDC liquid wealth

$$\widetilde{L}_1 = (1 + r^S)D_0^S + H_0.$$

If the household liquidates a fraction ω of its term deposits early, total liquid resources available for date-1 portfolio and payments decisions are

$$L_1(\omega) = (1 + r^S)D_0^S + H_0 + (1 - \kappa)\omega D_0^T.$$

These liquid resources can be allocated across savings deposits D_1^S, cash H_1, and CBDC C_1:

$$D_1^S + H_1 + C_1 \leq L_1(\omega).$$

The remaining term deposit position, $(1 - \omega)D_0^T$, is held to maturity and pays $(1 + r^T)(1 - \omega)D_0^T$ at the end of date 1.

Payments and CBDC adoption. As in the baseline model, households make a continuum of payment transactions during date 1. A transaction can be executed using cash, existing digital payment rails linked to deposit accounts, or CBDC if the household adopts. Let $a \in \{0, 1\}$ denote CBDC adoption. Adoption requires payment of a fixed cost $k \geq 0$, which may capture onboarding, learning, and technology frictions. Conditional on adoption, the household may use CBDC for digital transactions and obtains payment convenience benefits summarized by the function

$$\Psi(a, C_1; \theta),$$

where θ indexes household-specific payment needs or valuation of digital convenience. We assume that $\Psi(0, C_1; \theta) = 0$ and that for adopters the marginal convenience benefit of CBDC is weakly increasing in the amount of CBDC held over the relevant range.

Preferences. Household utility is additively separable between consumption/non-payment utility and payment convenience. Let terminal non-payment resources be denoted by

$$X_1(\omega) = D_1^S + H_1 + C_1 + (1 + r^T)(1 - \omega)D_0^T.$$

The household's date-1 problem is

$$V(D_0^S, D_0^T, H_0; \theta) = \max_{a \in \{0,1\},\, \omega \in [0,1],\, D_1^S, H_1, C_1} \{u(X_1(\omega)) + \Psi(a, C_1; \theta) - ak\}$$

subject to

$$D_1^S + H_1 + C_1 \leq (1 + r^S)D_0^S + H_0 + (1 - \kappa)\omega D_0^T,$$

and

$$D_1^S, H_1, C_1 \geq 0.$$

Consistent with Assumption 1, the non-payment part of utility depends on total available resources rather than on the exact composition of liquid balances across savings, cash, and CBDC. Composition matters only through payment convenience and through differences in returns and liquidity across assets.

Interpretation of the liquidation choice. The choice of ω captures whether the household finds it worthwhile to sacrifice the superior return on term deposits in order to create additional date-1 liquidity. Liquidating one additional unit of term deposits increases current liquid resources by only $(1 - \kappa)$ units but reduces terminal wealth by the forgone maturity payoff $(1 + r^T)$. Thus, early liquidation is attractive only if the value of the extra liquidity and payment convenience exceeds the return advantage of remaining invested in the term deposit.

Let λ_1 denote the shadow value of one additional unit of date-1 liquid resources in the household's optimization problem. For an interior choice of ω, the first-order condition is

$$(1 - \kappa)\lambda_1 = (1 + r^T)u'(X_1(\omega)) .$$

The left-hand side is the marginal value of converting one unit of term deposits into date-1 liquidity, net of the early-withdrawal penalty. The right-hand side is the marginal utility value of preserving that unit until maturity and receiving the higher gross return $1 + r^T$. Hence, the household liquidates term deposits only if the shadow value of immediate liquidity is sufficiently high.

A sufficient condition for no early liquidation. The baseline model corresponds to the region in which households do not unwind term deposits over the horizon relevant for CBDC adoption and usage. A sufficient condition for $\omega^\star = 0$ is that

$$(1 - \kappa)\lambda_1 < (1 + r^T)u'(X_1(0)) .$$

This condition states that the marginal value of additional liquid resources must be smaller than the marginal value of preserving the term deposit until maturity. It is more likely to hold when: (i) the term-deposit rate r^T is high relative to the return on CBDC and savings deposits; (ii) the early-withdrawal penalty κ is sizable; and (iii) CBDC primarily provides transaction convenience rather than a high pecuniary return.

These conditions closely match our empirical setting. The digital rupee is non-interest-bearing, so it does not compete with term deposits as a yield-bearing asset. Instead, it competes with liquid

savings balances and existing digital payment instruments. In this environment, the marginal value of CBDC arises primarily through liquidity and payment convenience, which is unlikely to dominate the combined return advantage and illiquidity cost associated with term deposits. As a result, the optimal response to CBDC operates mainly through liquid savings deposits rather than through term deposits.

Date-0 portfolio choice. At date 0, the household anticipates the date-1 problem and chooses its initial portfolio to maximize

$$U_0 = \max_{D_0^S, D_0^T, H_0} \left\{ u_0(W_0 - D_0^S - D_0^T - H_0) + \beta \, \mathbb{E}\left[V(D_0^S, D_0^T, H_0; \theta) \right] \right\},$$

subject to

$$D_0^S + D_0^T + H_0 \leq W_0, \qquad D_0^S, D_0^T, H_0 \geq 0.$$

Relative to liquid savings, term deposits provide a higher return but sacrifice flexibility. Liquid savings provide lower returns but preserve the option value of future portfolio reallocation when new payment technologies such as CBDC become available. The household therefore chooses D_0^T by balancing a standard yield-liquidity trade-off.

This formulation makes clear that the introduction of CBDC can, in principle, affect ex ante demand for term deposits even if households do not unwind existing term deposits ex post. If households anticipate that future CBDC access will raise the value of liquid balances, they may choose lower D_0^T and higher D_0^S at date 0. Our empirical analysis, however, focuses on short- to medium-run responses after the introduction of CBDC, a horizon over which the main observable adjustment occurs in liquid savings rather than in term deposits.

Proposition E.2. *Suppose CBDC is non-interest-bearing or offers a return strictly below the term-deposit rate r^T, and suppose the early-withdrawal penalty κ is sufficiently large. Then there exists a region of the parameter space such that the optimal early liquidation choice satisfies $\omega^\star = 0$. In that region, the household does not unwind term deposits in response to CBDC, and all portfolio adjustment occurs through liquid savings deposits, cash, and CBDC.*

Proof sketch. Because CBDC provides payment convenience but no competitive yield in our empirical setting, the gain from reallocating one unit of term deposits into CBDC is bounded by the marginal value of extra liquidity in date-1 transactions. When the term-deposit return advantage r^T and/or the early-withdrawal penalty κ are sufficiently large, the value of preserving term deposits until maturity strictly exceeds the benefit of converting them into liquid balances. The Kuhn–Tucker condition for the liquidation choice then implies a corner solution at $\omega^\star = 0$. Hence, CBDC affects the composition of liquid assets without inducing substitution out of term deposits.

Implications. This extension delivers two implications. First, it rationalizes the baseline assumption that term deposits are effectively fixed over the CBDC adoption horizon: when CBDC is low-yield and term deposits are sufficiently illiquid, households do not find it optimal to unwind term deposits to fund CBDC. Second, it clarifies how the model would change under alternative CBDC designs. If CBDC were to offer an interest rate close to or above r^T, or if term deposits could be liquidated at negligible cost, then the no-liquidation condition would fail. In that case, CBDC would compete not only with liquid savings deposits but also with illiquid term deposits, and the full intertemporal portfolio problem would need to be solved with endogenous adjustment along both margins.

Appendix F Imperfect Treatment, Partial CBDC Exposure, and Bounds under Callaway and Sant'Anna ATT

Our district-level treatment indicator is the timing of *lead-bank* authorization to offer the retail CBDC. This choice is attractive because it maps naturally into the lead bank system and provides plausibly exogenous staggered variation in access to CBDC. However, control districts are not literally CBDC free. Households in these districts may still access CBDC through non-lead banks that operate locally or digitally. In other words, the binary treatment indicator in our difference-in-differences specifications is an imperfect proxy for the underlying intensity of CBDC penetration. In this section we show formally, within the Callaway and Sant'Anna (2021) framework, that our estimand is an intention-to-treat (ITT) effect of lead-bank authorization under partial CBDC exposure. We then derive simple bounds that describe how much this attenuates the estimated treatment effect relative to the effect of moving from zero CBDC penetration to the observed treatment intensity.

Setup and Callaway and Sant'Anna ATT. Let d index districts and let $t = 1, \ldots, T$ index time periods. For each district d, let $G_d \in \{1, \ldots, T, \infty\}$ denote the period in which its lead bank first becomes authorized to offer CBDC. We set $G_d = \infty$ for never-treated districts. Treatment status is given by the binary indicator

$$D_{d,t} = \mathbf{1}\{t \geq G_d,\, G_d < \infty\}.$$

The observed outcome is $Y_{d,t}$. In the analysis below, $Y_{d,t}$ denotes either the natural logarithm of total retail deposits or the natural logarithm of total credit extended by all scheduled commercial banks in district d at time t. We denote the potential outcomes without and with lead-bank authorization by $Y_{d,t}(0)$ and $Y_{d,t}(1)$.

For a group g that first becomes treated at time g, the group and time specific average treatment effect on the treated in Callaway and Sant'Anna (2021) is

$$ATT(g,t) = \mathbb{E}\big[Y_{d,t}(1) - Y_{d,t}(0) \mid G_d = g\big], \qquad t \geq g.$$

The overall average treatment effect on the treated that the Callaway and Sant'Anna estimator reports is a weighted average of these group and time specific effects, where the weights $\omega_{g,t}$ are non negative and sum to one.

$$ATT^{CS} = \sum_{g,t} \omega_{g,t}\, ATT(g,t),$$

In our setting, the treatment indicator $D_{d,t}$, which records lead-bank authorization, is not the

primitive measure of CBDC exposure. Instead, the outcome depends on an underlying continuous measure of CBDC penetration that aggregates CBDC usage from both lead and non-lead banks.

CBDC penetration and a linear structural relationship. Let $C_{d,t}$ denote the true intensity of CBDC penetration in district d at time t. This variable aggregates CBDC activity from all scheduled commercial banks in the district. We assume a linear structural relationship between the outcome and CBDC penetration:

$$Y_{d,t}(c) = \alpha_d + \lambda_t + \beta c + \varepsilon_{d,t}, \tag{F.1}$$

where $Y_{d,t}(c)$ is the potential outcome that would be obtained if CBDC penetration were equal to c in district d at time t. The terms α_d and λ_t are district and time effects and $\varepsilon_{d,t}$ is an error term. The parameter β is the causal effect of a one unit change in CBDC penetration. Under equation (F.1), a change from CBDC penetration c_0 to c_1 changes the outcome by

$$Y_{d,t}(c_1) - Y_{d,t}(c_0) = \beta(c_1 - c_0). \tag{F.2}$$

Lead-bank authorization changes $C_{d,t}$, but control districts are not CBDC free. Even when $D_{d,t} = 0$, CBDC may be available through non-lead banks. For a given treatment group g and a time period t, define

$$C^T_{g,t} = \mathbb{E}[C_{d,t} \mid G_d = g], \qquad C^C_{g,t} = \mathbb{E}[C_{d,t} \mid d \in C(g,t)], \tag{F.3}$$

where $C(g,t)$ is the comparison group that the Callaway and Sant'Anna estimator uses for (g,t), such as never-treated or not-yet-treated districts. The quantity $C^T_{g,t}$ is the average CBDC penetration among treatment group g districts after lead-bank authorization. The quantity $C^C_{g,t}$ is the average CBDC penetration in the comparison group in the same period.

Under equations (F.1) and (F.2), we can express the group and time specific ATT in terms of CBDC penetration. Let $C^{(1)}_{d,t}$ and $C^{(0)}_{d,t}$ denote the CBDC penetration levels under lead-bank authorization and under the counterfactual without lead-bank authorization. Then

$$\begin{aligned} ATT(g,t) &= \mathbb{E}\big[Y_{d,t}(1) - Y_{d,t}(0) \mid G_d = g\big] \\ &= \mathbb{E}\big[Y_{d,t}(C^{(1)}_{d,t}) - Y_{d,t}(C^{(0)}_{d,t}) \mid G_d = g\big] \\ &= \mathbb{E}\big[\beta\big(C^{(1)}_{d,t} - C^{(0)}_{d,t}\big) \mid G_d = g\big] \\ &= \beta\,\mathbb{E}\big[C^{(1)}_{d,t} - C^{(0)}_{d,t} \mid G_d = g\big]. \end{aligned} \tag{F.4}$$

The usual identifying assumptions are the no anticipation assumption

$$\mathbb{E}[C^{(0)}_{d,g-1} \mid G_d = g] = C^T_{g,g-1} \tag{F.5}$$

and the parallel trends assumption

$$\mathbb{E}[C^{(0)}_{d,t} - C^{(0)}_{d,g-1} \mid G_d = g] = \mathbb{E}[C_{d,t} - C_{d,g-1} \mid d \in C(g,t)] = C^C_{g,t} - C^C_{g,g-1}. \tag{F.6}$$

We make one further assumption, namely that average CBDC penetration in the base period is equal across treated and comparison districts

$$C^T_{g,g-1} = C^C_{g,g-1}. \tag{F.7}$$

Two features of the Indian institutional environment support this assumption. First, lead-bank assignments were fixed under the Lead Bank Scheme of 1969 on historical and administrative grounds, and the RBI's CBDC authorization sequence reflected banks' internal technical readiness rather than local conditions. A district's lead-bank identity is therefore unlikely to carry material information about the presence or CBDC activity of individuals banking with other banks operating in the district. Thus, CBDC penetration is unlikely to systematically differ across treated and comparison districts prior to lead-bank authorization. Second, because lead banks intermediate the majority of retail deposits and payment relationships in their districts, CBDC penetration due to non-lead banks is likely to be small. So any residual difference between the two groups in the pre-treatment period is correspondingly small.

Under assumptions (F.5), (F.6), and (F.7), we obtain

$$\mathbb{E}\big[C^{(1)}_{d,t} \mid G_d = g\big] = C^T_{g,t}, \qquad \mathbb{E}\big[C^{(0)}_{d,t} \mid G_d = g\big] = C^C_{g,t}, \tag{F.8}$$

which implies that equation (F.4) becomes

$$ATT(g,t) = \beta\big(C^T_{g,t} - C^C_{g,t}\big). \tag{F.9}$$

Equation (F.9) shows that each group and time specific ATT is an intention-to-treat effect. It measures the effect on the outcome of increasing CBDC penetration from the average level in the comparison group, $C^C_{g,t}$, to the average level in the treated group, $C^T_{g,t}$, at time t.

From ATTs to a zero-to-treated effect and bounds. For any pair (g,t), a useful benchmark is the effect of moving from zero CBDC penetration to the treated level $C^T_{g,t}$. We denote this by

$$\Delta^{0\to T}_{g,t} = \beta C^T_{g,t}. \tag{F.10}$$

We combine equations (F.9) and (F.10) by defining the relative CBDC penetration in the comparison group as

$$\rho_{g,t} = \frac{C^C_{g,t}}{C^T_{g,t}}. \tag{F.11}$$

With this definition we have

$$ATT(g,t) = \beta C^T_{g,t}(1-\rho_{g,t}) = \Delta^{0\to T}_{g,t}(1-\rho_{g,t}), \tag{F.12}$$

and the scaling factor that relates the zero-to-treated effect to the group and time specific ATT is

$$\frac{\Delta^{0\to T}_{g,t}}{ATT(g,t)} = \frac{1}{1-\rho_{g,t}}. \tag{F.13}$$

If $0 \le \rho_{g,t} < 1$, then the comparison group has strictly lower CBDC penetration than the treated group but is not CBDC free. In this case each group and time specific ATT is attenuated toward zero relative to the corresponding zero-to-treated effect. The parameter $\rho_{g,t}$ determines the size of this attenuation.

The overall ATT is a weighted average of group and time specific ATTs,

$$ATT^{CS} = \sum_{g,t} \omega_{g,t}\, ATT(g,t), \qquad \sum_{g,t} \omega_{g,t} = 1,\ \omega_{g,t} \ge 0. \tag{F.14}$$

We define the corresponding weighted zero-to-treated effect by

$$\Delta^{0\to T}_{CS} = \sum_{g,t} \omega_{g,t}\, \Delta^{0\to T}_{g,t} = \sum_{g,t} \omega_{g,t}\, \beta C^T_{g,t}. \tag{F.15}$$

Using equation (F.12), we can write

$$\Delta^{0\to T}_{CS} = \sum_{g,t} \omega_{g,t}\, ATT(g,t) \cdot \frac{1}{1-\rho_{g,t}}. \tag{F.16}$$

Suppose we can bound the relative CBDC penetration in the comparison group in all group and time

cells by a common upper bound

$$\rho_{g,t} \leq \bar{\rho} < 1 \quad \text{for all } (g,t). \tag{F.17}$$

Then $(1-\rho_{g,t})^{-1} \leq (1-\bar{\rho})^{-1}$ for all (g,t), and equation (F.16) implies

$$\left|\Delta_{CS}^{0\to T}\right| \leq \left|\frac{1}{1-\bar{\rho}} \sum_{g,t} \omega_{g,t} ATT(g,t)\right| = \frac{1}{1-\bar{\rho}} \left|ATT^{CS}\right|. \tag{F.18}$$

Under the linear structure in equation (F.1) and the uniform bound in equation (F.17), the average effect of moving from zero CBDC penetration to the treated group level of CBDC penetration, $\Delta_{CS}^{0\to T}$, is at most a factor $(1-\bar{\rho})^{-1}$ times as large in magnitude than the Callaway and Sant'Anna overall ATT, ATT^{CS}. When control districts have partial CBDC exposure that reaches at most a fraction $\bar{\rho}$ of treated districts' penetration in the relevant two by two comparisons, the estimator ATT^{CS} understates the corresponding zero-to-treated effect by at most the factor $(1-\bar{\rho})^{-1}$.

Specialization to deposits and credit using bank footprints. In the empirical work, we estimate separate Callaway and Sant'Anna (2021) ATTs for the natural logarithm of total retail deposits and for the natural logarithm of total credit at the district level. We denote these by ATT^{CS}_{Dep} and ATT^{CS}_{Cred}. The bound in equation (F.18) maps each of these estimates into a zero-to-treated effect once an outcome-specific value of $\bar{\rho}$ is supplied. We now describe a possible method for constructing $\bar{\rho}$ from bank geographic footprints together with bank-level CBDC authorization.

First, for each district d and bank b, compute the pre-CBDC share of deposits or credit that bank b intermediates in district d. We denote these shares by $s^{\text{Dep}}_{b,d}$ and $s^{\text{Cred}}_{b,d}$. Second, combine these shares with bank-level CBDC authorization to form district-level exposure indices $\tilde{C}^{\text{Dep}}_{d,t}$ and $\tilde{C}^{\text{Cred}}_{d,t}$

$$\tilde{C}^{\text{Dep}}_{d,t} = \sum_b s^{\text{Dep}}_{b,d} \cdot 1\{\text{bank } b \text{ authorized for CBDC by time } t\},$$
$$\tilde{C}^{\text{Cred}}_{d,t} = \sum_b s^{\text{Cred}}_{b,d} \cdot 1\{\text{bank } b \text{ authorized for CBDC by time } t\}.$$

Third, calculate the ratio of these indices $\frac{C^C_{d,t}}{C^T_{d,t}}$ for each relevant group-time cell, and take the upper envelope across all cells to recover the values of $\tilde{\rho}^{\text{Dep}}$ and $\tilde{\rho}^{\text{Cred}}$. Applying equation (F.18) with $\bar{\rho} = \tilde{\rho}^{\text{Dep}}$ for deposits and $\bar{\rho} = \tilde{\rho}^{\text{Cred}}$ for credit then yields

$$\left|\Delta^{0\to T}_{CS,\text{Dep}}\right| \leq \frac{1}{1-\tilde{\rho}^{\text{Dep}}} \left|ATT^{CS}_{\text{Dep}}\right|, \qquad \left|\Delta^{0\to T}_{CS,\text{Cred}}\right| \leq \frac{1}{1-\tilde{\rho}^{\text{Cred}}} \left|ATT^{CS}_{\text{Cred}}\right|, \tag{F.19}$$

where $\Delta^{0\to T}_{CS,\text{Dep}}$ and $\Delta^{0\to T}_{CS,\text{Cred}}$ denote the average effects of moving from zero CBDC penetration to

the treated-group penetration level for deposits and credit, respectively. This construction requires district-by-bank footprint shares matched to bank-level CBDC usage. We therefore present it as a possible procedure that maps our reported estimates into zero-to-treated effects, which a reader or regulator with access to such data can apply directly.

Because lead banks intermediate the majority of retail deposits and payment relationships in their districts, CBDC penetration due to non-lead banks in comparison districts is likely to be modest, so $\bar{\rho}$ is likely to be bounded below one. At the same time, comparison districts are not literally CBDC-free, so $\bar{\rho}$ is likely to be positive. We therefore interpret the reported Callaway and Sant'Anna ATTs for lead-bank authorization as intention-to-treat effects under partial CBDC exposure. The bounds above show that the corresponding zero-to-treated effects on deposits and credit must be larger in magnitude, and that the maximum size of this gap is given by the factors $(1 - \tilde{\rho}^{\mathrm{Dep}})^{-1}$ and $(1 - \tilde{\rho}^{\mathrm{Cred}})^{-1}$.

Appendix G The Costs & Benefits of CBDC

G.1 Proofs

In this subsection, we provide proofs for Proposition 9.1 and Corollary 9.1.

Proof of Proposition 9.1. We defined the bank's total wholesale funding cost as

$$R(B) = \rho_0 B + \frac{1}{2}\psi B^2, \tag{G.1}$$

so $R'(B) = \rho_0 + \psi B$ and $R''(B) = \psi \geq 0$. Using the within-period funding constraint, we can rewrite the bank's objective as

$$\Pi^{\text{bank}} = r^\ell \ell - (r^S - \lambda)D - R(\ell - D - e) - c(\ell). \tag{G.2}$$

The first order condition for ℓ gives us

$$r^\ell - R'(B) - c'(\ell) = 0 \iff r^\ell - \rho_0 - \psi B - c'(\ell) = 0. \tag{G.3}$$

The marginal loan earns r^ℓ, is funded at the margin by wholesale funds costing $R'(B)$, and carries origination cost $c'(\ell)$. The second-order condition is

$$\frac{\partial^2 \Pi^{\text{bank}}}{\partial \ell^2} = -R''(B) - c''(\ell) = -(\psi + c''(\ell)) < 0, \tag{G.4}$$

so the first order condition characterizes an interior maximum.

We now totally differentiate the first order condition with respect to D, holding r^ℓ, ρ_0, ψ fixed. Since $B = \ell - D - e$, we have $dB = d\ell - dD$. Thus, totally differentiating the first order condition implies

$$-\psi dB - c''(\ell)d\ell = 0 \iff -\psi(d\ell - dD) - c''(\ell)d\ell = 0. \tag{G.5}$$

Collecting terms in $d\ell$,

$$(\psi + c''(\ell))d\ell = \psi dD \implies \theta \equiv \frac{d\ell}{dD} = \frac{\psi}{\psi + c''(\ell)}. \tag{G.6}$$

Since $\psi \geq 0$, and $c'' > 0$, the numerator is non-negative and strictly smaller than the denominator, so

$\theta \in [0, 1)$. The wholesale response follows from $dB = d\ell - dD$:

$$\frac{dB}{dD} = \frac{d\ell}{dD} - 1 = \theta - 1 = -(1 - \theta). \tag{G.7}$$

□

Proof of Corollary 9.1. Differentiating the deposit-to-loan pass-through θ with respect to ψ gives

$$\frac{\partial \theta}{\partial \psi} = \frac{c''(\ell)}{(\psi + c''(\ell))^2} > 0. \tag{G.8}$$

Suppose that $\psi_{SR} > \psi_{LR}$, and let $\theta_h = \frac{\psi_h}{\psi_h + c''(\ell)}$, where $h \in \{SR, LR\}$. Then, it is clear that $\theta_{SR} > \theta_{LR}$. □

G.2 Term-by-Term Evaluation

In this subsection, we derive the estimates presented in Table 4. All quantities are annual flows for a single median district in 2022 Q3. For the calculations, we assume the following values: deposit outflow $|dD| = 1.843$ billion, pass-through $\theta_{SR} = 0.476 \to \theta_{LR} = 0$, adoption rate $a = 0.044$, adopters $Na = 88{,}000$, adopter surplus $\overline{\Delta v} = 620$ rupees, franchise spread $\lambda = 2\%$, wholesale slope $\psi_{SR}|dD| = 1.5\% \to \psi_{LR}|dD| = 0$, credit wedge $\eta = 0.05$, net seigniorage $s^{\mathrm{CB}} = 1\%$, savings rate $r^S = 3\%$, and lending curvature penalty $c''(\ell)_{SR}|dD| \approx 1.65$ percentage points. Substituting these into (13) term by term:

$$\text{(i) household surplus} \quad Na\overline{\Delta v} = 88{,}000 \times 620 = 54.6 \text{ m}, \tag{G.9}$$

$$\text{(ii) lost franchise} \quad \lambda|dD| = 0.02 \times 1.843 \text{ bn} = 36.9 \text{ m}, \tag{G.10}$$

$$\text{(iii) wholesale deadweight} \quad \tfrac{1}{2}\big(\psi_{SR}|dD|\big)(1-\theta)^2|dD| = \tfrac{1}{2} \times 0.015 \times (0.524)^2 \times 3.8 \text{ bn} = 1.0 \text{ m}, \tag{G.11}$$

$$\text{(iv) lending curvature loss} \quad \tfrac{1}{2}\big(c''(\ell)_{SR}|dD|\big)\theta^2|dD| = \tfrac{1}{2} \times 0.0165 \times (0.476)^2 \times 1.843 \text{ bn} = 3.4 \text{ m}, \tag{G.12}$$

$$\text{(v) credit externality} \quad \eta\theta_{SR}|dD| = 0.05 \times 0.476 \times 1.843 \text{ bn} = 43.9 \text{ m}, \tag{G.13}$$

$$\text{(vi) net seigniorage} \quad s^{\mathrm{CB}}|dD| = 0.01 \times 1.843 \text{ bn} = 18.4 \text{ m}. \tag{G.14}$$

Summing with the signs of (13) gives the short-run impact figure,

$$\Delta \mathcal{W}_{\text{impact}} = 54.6 - 36.9 - 3.8 - 3.4 - 43.9 + 18.4 = -15 \text{ m}, \tag{G.15}$$

which is the first column of Table 4. In the long-run credit reverts, so $\theta_{LR} = 0$ and $\psi_{LR} = 0$ and terms (iii), (iv), and (v) vanish entirely. Terms (i), (ii), and (vi) are unaffected, since none depend on θ and ψ. This gives

$$\Delta \mathcal{W}_{\text{LR}} = 54.6 - 36.9 - 0 - 0 - 0 + 18.4 = 36.1 \text{ m}. \tag{G.16}$$

Gross gains are the sum of the two benefit terms, $Na\overline{\Delta v} + s^{\text{CB}}|dD| = 73$ m, against which the shares in the parentheses of Table 4 are computed.

G.3 Sensitivity Analysis to Calibrated Parameters

The term-by-term evaluation in Section G.2 rests on quantities we estimate and on a set of prices and wedges we assume. Because these inputs are assumptions rather than measurements, we now conduct a robustness check on our welfare result by varying the inputs. In the following tables, we vary the average adopter surplus $\overline{\Delta v}$, the deposit franchise spread λ, the wholesale slope $\psi|dD|$, and the credit wedge η.

Table G.1: Sensitivity of Short-Term Net District Welfare to $\overline{\Delta v}$, λ, and $\psi_{SR}|dD|$ with $\eta = 0.025$

	Deposit franchise spread λ				
Adopter surplus $\overline{\Delta v}$	1.50%	1.75%	2.00%	2.25%	2.50%
₹550	17	13	8	3	−1
₹600	22	17	12	8	3
₹650	26	21	17	12	8
₹700	30	26	21	17	12
₹750	35	30	26	21	16

(a) Wholesale Slope $\psi_{SR}|dD| = 0$

	Deposit franchise spread λ				
Adopter surplus $\overline{\Delta v}$	1.50%	1.75%	2.00%	2.25%	2.50%
₹550	14	9	4	−0	−5
₹600	18	13	9	4	−0
₹650	22	18	13	9	4
₹700	27	22	18	13	8
₹750	31	27	22	17	13

(b) Wholesale Slope $\psi_{SR}|dD| = 0.0075$

	Deposit franchise spread λ				
Adopter surplus $\overline{\Delta v}$	1.50%	1.75%	2.00%	2.25%	2.50%
₹550	10	5	1	−4	−8
₹600	14	10	5	1	−4
₹650	19	14	10	5	0
₹700	23	19	14	9	5
₹750	28	23	18	14	9

(c) Wholesale Slope $\psi_{SR}|dD| = 0.015$

	Deposit franchise spread λ				
Adopter surplus $\overline{\Delta v}$	1.50%	1.75%	2.00%	2.25%	2.50%
₹550	6	2	−3	−7	−12
₹600	11	6	2	−3	−8
₹650	15	11	6	1	−3
₹700	20	15	10	6	1
₹750	24	19	15	10	6

(d) Wholesale Slope $\psi_{SR}|dD| = 0.0225$

	Deposit franchise spread λ				
Adopter surplus $\overline{\Delta v}$	1.50%	1.75%	2.00%	2.25%	2.50%
₹550	3	−2	−6	−11	−16
₹600	7	3	−2	−7	−11
₹650	12	7	2	−2	−7
₹700	16	11	7	2	−2
₹750	20	16	11	7	2

(e) Wholesale Slope $\psi_{SR}|dD| = 0.03$

This table presents the evaluation of $\Delta\mathcal{W}$ in the short-run ($\theta_{SR} = 0.476$) under different values of the adopter surplus $\overline{\Delta v}$, the deposit franchise spread λ, and the wholesale slope $\psi_{SR}|dD|$ with $\eta = 0.025$, while holding other parameters fixed and recalculating $c''(\ell)_{SR}|dD|$ according to Proposition 9.1. All welfare estimates are presented in millions of rupees.

Table G.2: Sensitivity of Short-Term Net District Welfare to $\overline{\Delta v}$, λ, and $\psi_{SR}|dD|$ with $\eta = 0.05$

	Deposit franchise spread λ				
Adopter surplus $\overline{\Delta v}$	1.50%	1.75%	2.00%	2.25%	2.50%
₹550	−5	−9	−14	−19	−23
₹600	−0	−5	−9	−14	−19
₹650	4	−0	−5	−10	−14
₹700	9	4	−1	−5	−10
₹750	13	8	4	−1	−6

(a) Wholesale Slope $\psi_{SR}|dD| = 0$

	Deposit franchise spread λ				
Adopter surplus $\overline{\Delta v}$	1.50%	1.75%	2.00%	2.25%	2.50%
₹550	−8	−13	−18	−22	−27
₹600	−4	−9	−13	−18	−22
₹650	1	−4	−9	−13	−18
₹700	5	0	−4	−9	−14
₹750	9	5	0	−5	−9

(b) Wholesale Slope $\psi_{SR}|dD| = 0.0075$

	Deposit franchise spread λ				
Adopter surplus $\overline{\Delta v}$	1.50%	1.75%	2.00%	2.25%	2.50%
₹550	−12	−17	−21	−26	−30
₹600	−8	−12	−17	−21	−26
₹650	−3	−8	−12	−17	−22
₹700	1	−3	−8	−13	−17
₹750	6	1	−4	−8	−13

(c) Wholesale Slope $\psi_{SR}|dD| = 0.015$

	Deposit franchise spread λ				
Adopter surplus $\overline{\Delta v}$	1.50%	1.75%	2.00%	2.25%	2.50%
₹550	−16	−20	−25	−29	−34
₹600	−11	−16	−20	−25	−30
₹650	−7	−11	−16	−21	−25
₹700	−2	−7	−12	−16	−21
₹750	2	−3	−7	−12	−16

(d) Wholesale Slope $\psi_{SR}|dD| = 0.0225$

	Deposit franchise spread λ				
Adopter surplus $\overline{\Delta v}$	1.50%	1.75%	2.00%	2.25%	2.50%
₹550	−19	−24	−28	−33	−38
₹600	−15	−19	−24	−29	−33
₹650	−10	−15	−20	−24	−29
₹700	−6	−11	−15	−20	−24
₹750	−2	−6	−11	−15	−20

(e) Wholesale Slope $\psi_{SR}|dD| = 0.03$

This table presents the evaluation of $\Delta\mathcal{W}$ in the short-run ($\theta_{SR} = 0.476$) under different values of the adopter surplus $\overline{\Delta v}$, the deposit franchise spread λ, and the wholesale slope $\psi_{SR}|dD|$ with $\eta = 0.05$, while holding other parameters fixed and recalculating $c''(\ell)_{SR}|dD|$ according to Proposition 9.1. All welfare estimates are presented in millions of rupees.

Table G.3: Sensitivity of Short-Term Net District Welfare to $\overline{\Delta v}$, λ, and $\psi_{SR}|dD|$ with $\eta = 0.1$

	Deposit franchise spread λ				
Adopter surplus $\overline{\Delta v}$	1.50%	1.75%	2.00%	2.25%	2.50%
₹550	−49	−53	−58	−62	−67
₹600	−44	−49	−53	−58	−63
₹650	−40	−44	−49	−54	−58
₹700	−35	−40	−45	−49	−54
₹750	−31	−36	−40	−45	−49

(a) Wholesale Slope $\psi_{SR}|dD| = 0$

	Deposit franchise spread λ				
Adopter surplus $\overline{\Delta v}$	1.50%	1.75%	2.00%	2.25%	2.50%
₹550	−52	−57	−61	−66	−71
₹600	−48	−52	−57	−62	−66
₹650	−43	−48	−53	−57	−62
₹700	−39	−44	−48	−53	−57
₹750	−35	−39	−44	−48	−53

(b) Wholesale Slope $\psi_{SR}|dD| = 0.0075$

	Deposit franchise spread λ				
Adopter surplus $\overline{\Delta v}$	1.50%	1.75%	2.00%	2.25%	2.50%
₹550	−56	−60	−65	−70	−74
₹600	−51	−56	−61	−65	−70
₹650	−47	−52	−56	−61	−65
₹700	−43	−47	−52	−56	−61
₹750	−38	−43	−47	−52	−57

(c) Wholesale Slope $\psi_{SR}|dD| = 0.015$

	Deposit franchise spread λ				
Adopter surplus $\overline{\Delta v}$	1.50%	1.75%	2.00%	2.25%	2.50%
₹550	−59	−64	−69	−73	−78
₹600	−55	−60	−64	−69	−73
₹650	−51	−55	−60	−64	−69
₹700	−46	−51	−55	−60	−65
₹750	−42	−46	−51	−56	−60

(d) Wholesale Slope $\psi_{SR}|dD| = 0.0225$

	Deposit franchise spread λ				
Adopter surplus $\overline{\Delta v}$	1.50%	1.75%	2.00%	2.25%	2.50%
₹550	−63	−68	−72	−77	−81
₹600	−59	−63	−68	−72	−77
₹650	−54	−59	−63	−68	−73
₹700	−50	−54	−59	−64	−68
₹750	−45	−50	−55	−59	−64

(e) Wholesale Slope $\psi_{SR}|dD| = 0.03$

This table presents the evaluation of $\Delta\mathcal{W}$ in the short-run ($\theta_{SR} = 0.476$) under different values of the adopter surplus $\overline{\Delta v}$, the deposit franchise spread λ, and the wholesale slope $\psi_{SR}|dD|$ with $\eta = 0.1$, while holding other parameters fixed and recalculating $c''(\ell)_{SR}|dD|$ according to Proposition 9.1. All welfare estimates are presented in millions of rupees.

Table G.4: Sensitivity of Short-Term Net District Welfare to $\overline{\Delta v}$, λ, and $\psi_{SR}|dD|$ with $\eta = 0.15$

	Deposit franchise spread λ				
Adopter surplus $\overline{\Delta v}$	1.50%	1.75%	2.00%	2.25%	2.50%
₹550	−92	−97	−102	−106	−111
₹600	−88	−93	−97	−102	−106
₹650	−84	−88	−93	−97	−102
₹700	−79	−84	−88	−93	−98
₹750	−75	−79	−84	−89	−93

(a) Wholesale Slope $\psi_{SR}|dD| = 0$

	Deposit franchise spread λ				
Adopter surplus $\overline{\Delta v}$	1.50%	1.75%	2.00%	2.25%	2.50%
₹550	−96	−101	−105	−110	−114
₹600	−92	−96	−101	−105	−110
₹650	−87	−92	−96	−101	−106
₹700	−83	−87	−92	−97	−101
₹750	−78	−83	−88	−92	−97

(b) Wholesale Slope $\psi_{SR}|dD| = 0.0075$

	Deposit franchise spread λ				
Adopter surplus $\overline{\Delta v}$	1.50%	1.75%	2.00%	2.25%	2.50%
₹550	−100	−104	−109	−113	−118
₹600	−95	−100	−104	−109	−114
₹650	−91	−95	−100	−105	−109
₹700	−86	−91	−96	−100	−105
₹750	−82	−87	−91	−96	−100

(c) Wholesale Slope $\psi_{SR}|dD| = 0.015$

	Deposit franchise spread λ				
Adopter surplus $\overline{\Delta v}$	1.50%	1.75%	2.00%	2.25%	2.50%
₹550	−103	−108	−112	−117	−122
₹600	−99	−103	−108	−113	−117
₹650	−94	−99	−104	−108	−113
₹700	−90	−95	−99	−104	−108
₹750	−86	−90	−95	−99	−104

(d) Wholesale Slope $\psi_{SR}|dD| = 0.0225$

	Deposit franchise spread λ				
Adopter surplus $\overline{\Delta v}$	1.50%	1.75%	2.00%	2.25%	2.50%
₹550	−107	−111	−116	−121	−125
₹600	−102	−107	−112	−116	−121
₹650	−98	−103	−107	−112	−117
₹700	−94	−98	−103	−108	−112
₹750	−89	−94	−99	−103	−108

(e) Wholesale Slope $\psi_{SR}|dD| = 0.03$

This table presents the evaluation of $\Delta\mathcal{W}$ in the short-run ($\theta_{SR} = 0.476$) under different values of the adopter surplus $\overline{\Delta v}$, the deposit franchise spread λ, and the wholesale slope $\psi_{SR}|dD|$ with $\eta = 0.15$, while holding other parameters fixed and recalculating $c''(\ell)_{SR}|dD|$ according to Proposition 9.1. All welfare estimates are presented in millions of rupees.

Table G.5: Sensitivity of Long-Term Net District Welfare to $\overline{\Delta v}$, and λ

	Deposit franchise spread λ				
Adopter surplus $\overline{\Delta v}$	1.50%	1.75%	2.00%	2.25%	2.50%
₹550	39	35	30	25	21
₹600	44	39	34	30	25
₹650	48	43	39	34	30
₹700	52	48	43	39	34
₹750	57	52	48	43	38

This table presents the evaluation of $\Delta\mathcal{W}$ in the long-run ($\theta_{LR} = 0$ and $\psi_{LR} = 0$) under different values of the adopter surplus $\overline{\Delta v}$, and the deposit franchise spread λ, while holding other parameters fixed. All welfare estimates are presented in millions of rupees.

The sensitivity analysis reveals that the long-term welfare gain remains positive over the entire-grid, while the short-term loss hinges on $\eta \geq 0.05$. As the marginal social value of bank credit

decreases $\eta \to 0$, the effect of the credit-contraction in the short-term decreases, leading to a positive welfare effect for the introduction of a CBDC. Ultimately, while the precise money-metric figure depends on the calibration, its convenience-dominated composition and long-term positive gain do not. These figures suggest that our conclusion is robust to different assumptions on the four parameters $\left(\overline{\Delta v}, \lambda, \psi, \eta\right)$.