Who wants privacy anyway?

Comparative Evidence on CBDC Design Preference

Vagisha Srivastava

Georgia Institute of Technology

Submitted for The Workshop on the Economics of Information Security (WEIS) 2026

February 1, 2026

Abstract

Central bank digital currencies (CBDCs) require design decisions that embed security, privacy, and governance trade-offs into monetary infrastructure. Yet little empirical evidence exists on how users evaluate these trade-offs or whether preferences vary systematically across institutional contexts. This paper aims to address this gap through conjoint experiments in the United States and India, where respondents chose between hypothetical CBDC designs varying in security, privacy, government control, cost, speed, and intermediation structure. It tests whether theoretically motivated individual-level characteristics such as trust in the central bank, trust in digital payments, fraud victimization, privacy confidence, and cryptocurrency usage affect design preferences. The central finding is the consensus effect: within each payment ecosystem, users converge on shared preference structures regardless of their personal trust, fraud experience, privacy attitudes, or cryptocurrency usage. Of [462] within-country moderation tests, only [6.7%] reach even uncorrected significance, a rate barely above the chance false-positive baseline, while 13 of 21 cross-country comparisons show statistically significant divergence at Bonferroni-corrected thresholds. The sole within-country exception is privacy confidence in the United States, where respondents skeptical of Federal Reserve privacy protection show stronger preferences for architecturally private designs. These findings suggest that institutional context, rather than individual psychology, shapes CBDC design preferences, and that design optimization cannot substitute for institutional credibility.

Keywords: CBDC; Conjoint experiment; Privacy preferences; Security-privacy trade-offs; Institutional trust; Payment ecosystems

1 Introduction

Central bank digital currencies represent one of the most consequential information security design challenges facing monetary authorities worldwide. Unlike conventional cybersecurity problems where technical countermeasures can be evaluated against well-defined threat models, CBDC design requires navigating fundamental trade-offs between security, privacy, and state control that implicate both technical architecture and institutional relationships. The 134 central banks actively researching or piloting CBDCs must make design decisions that will shape financial surveillance capabilities, data protection regimes, and the boundary between public and private payment infrastructure for decades. Yet, very little empirical evidence exists on how citizens evaluate these trade-offs, and whether such preferences can be addressed through design features or are more fundamentally shaped by institutional context.

This paper addresses this gap through conjoint experiments in two large, institutionally distinct economies: the United States and India. Respondents are presented with hypothetical CBDC designs varying across seven dimensions —security, privacy, government control, transaction cost, speed, cross-border usability, and intermediation; and their preferred configurations are observed. Crucially, this paper tests whether individual-level characteristics that theory predicts should matter, like trust in central banks, fraud victimization, privacy confidence, and cryptocurrency experience, actually moderate these preferences within countries. This study focuses on retail CBDC, digital currency intended for everyday consumer transactions, rather than wholesale CBDC used for interbank settlement, where design preferences involve different stakeholders and different governance considerations.

The central finding challenges conventional assumptions: individual characteristics largely fail to differentiate CBDC design preferences within countries, even as substantial cross-country differences persist. Of 36 within-country moderation tests, only one reaches statistical significance. This paper terms this pattern the *consensus effect*: within each payment ecosystem, users converge on similar preference structures regardless of their personal trust, experience, or attitudes. This suggests that CBDC acceptance depends less on matching designs to diverse individual concerns and more on whether the institutional context permits CBDC to inherit legitimacy from existing infrastructure.

1.1 Why CBDC Design Matters for Information Security Research?

For the economics of information security community, CBDC presents a compelling case study in the political economy of security architecture. Four features make it distinctive -

First, CBDC collapses the traditional principal–agent structure of payment security. In existing electronic payment systems, security is provided by intermediaries—banks, card networks, payment processors—who bear liability for fraud and data breaches, creating incentive alignment for protective investment. CBDC potentially disintermediates these actors, concentrating security provision and surveillance capability in the central bank. This raises novel questions about who bears security costs, who captures security benefits, and how accountability structures change when the state becomes the direct provider of retail payment infrastructure.

Second, CBDC makes explicit trade-offs that remain implicit in existing systems. Current digital payments involve substantial privacy compromises that users accept through inattention, lock-in, or lack of alternatives. Card networks track purchases; mobile payment apps harvest location data; banks re-

port transactions to tax authorities. CBDC design processes, by contrast, require explicit choices about transaction anonymity thresholds, data retention periods, conditions for government access, and the technical architecture of identity verification. This explicitness creates both policy opportunity and political vulnerability where design choices become visible and contestable in ways that emergent surveillance architectures are not.

Third, CBDC is governance-laden infrastructure where technical and institutional design are inseparable. A "privacy-preserving" CBDC architecture using zero-knowledge proofs offers fundamentally different actual protection depending on whether the implementing institution operates under GDPR constraints in the EU, Section 702 surveillance authorities in the United States, or India's data localization requirements. Technical features cannot be evaluated independently of the institutional context that determines how those features will be used, circumvented, or extended over time. This inseparability of technical and institutional design is a core theme in security economics that CBDC crystallizes with unusual clarity.

Fourth, CBDC enables natural experiments in cross-national security preference divergence. The United States and India represent starkly different payment ecosystem contexts. India's Unified Payments Interface (UPI) has achieved near-universal adoption with government-embedded oversight, KYC requirements, and normalized data-sharing – creating a population accustomed to state involvement in digital finance. The United States maintains fragmented payment rails, polarized attitudes toward federal surveillance, and no equivalent public digital infrastructure. Comparing CBDC preferences across these contexts highlights how baseline ecosystem experience shapes the meaning and acceptability of specific security and privacy configurations.

1.2 The Research Problem: Do Individual Concerns Drive Design Preferences?

Policy debates about CBDC adoption frequently assume that privacy concerns, distrust of government, or resistance to surveillance will constitute primary barriers that thoughtful design can address. If privacy-concerned users can be offered anonymous small-value transactions, if government-skeptical users can be assured of legal protections against account freezing, if security-anxious users can be promised fraud insurance — then perhaps CBDC can achieve broad adoption through feature configurations that accommodate heterogeneous concerns.

This assumption, that CBDC design can accommodate heterogeneous individual concerns through targeted feature configurations, is grounded in technology acceptance research that models adoption as an individual-level decision shaped by personal perceptions of usefulness, risk, and trust. If this assumption holds, users with low trust in central banks should demand more privacy protection, fraud victims should prioritize security features, and cryptocurrency enthusiasts should prefer decentralized architectures. Central banks could then optimize adoption by matching design features to the distributional characteristics of their populations.

Yet several well-established theoretical traditions predict a different pattern. Economic sociology has long argued that economic action is embedded in institutional structures rather than driven by atomized individual calculation (Granovetter, 1985). The moral economy tradition holds that shared normative expectations about legitimate economic practice, rather than individual cost-benefit reasoning, govern how communities evaluate economic arrangements (Thompson, 1971). Domestication theory

in science and technology studies argues that technology meaning is constructed through situated social practices within specific institutional contexts rather than evaluated against abstract utility functions (Silverstone Hirsch, 1992; Hyysalo, 2010). Applied to CBDC, these perspectives predict that users within a shared payment ecosystem will converge on similar preference structures shaped by collective institutional experience, regardless of individual variation in trust, privacy attitudes, or technology experience. This tension between individual-level and institutional-level accounts of preference formation motivates the research design. Before testing whether specific individual characteristics moderate specific design preferences, a prior question must be addressed: RQ0: Are CBDC design preferences primarily shaped by individual-level characteristics, or by the institutional and infrastructural context in which users are embedded? RQ0 is not constructed as a hypothesis but as the framing question within which the specific moderation hypotheses and cross-country comparisons are situated. If individual-level theories are correct, RQ1 through RQ4 should yield significant moderation effects within countries. If institutional-context theories are correct, cross-country divergence (RQ5) should dominate while within-country moderation should be weak or absent. The specific research questions that structure this contest are:

- **RQ1:** Do individuals with higher trust in central banks tolerate more government control in CBDC design?
- **RQ2:** Do trust in digital payments and fraud victimization moderate security feature preferences?
- **RQ3:** Does privacy confidence predict demand for privacy-preserving CBDC features?
- **RQ4:** Do cryptocurrency users demand more autonomy and privacy in CBDC design?
- **RQ5:** Do CBDC design preferences differ systematically between the United States and India?

1.3 Summary of Key Findings

The results yield a striking pattern: dramatic cross-country divergence coexists with near-universal within-country agreement.

Finding 1: Transaction cost and speed dominate preferences in both countries. Respondents in both the US and India prioritize free transactions (USA: 61.5%, India: 56.7% marginal mean) and instantaneous settlement (USA: 57.7%, India: 55.1%) over other features. Cost exhibits the largest preference gradient in both samples—a 25.2 percentage point range in the USA and 13.3 percentage points in India. This confirms that functional attributes, not governance concerns, are the primary drivers of stated CBDC preferences.

Finding 2: Americans weight privacy and autonomy features substantially more heavily in their choices, while Indian respondents show relatively flat preference gradients across governance configurations. Importantly, majorities in both countries prefer full privacy and minimal government control when these features are presented; the cross-country difference is one of intensity rather than direction. Cross-country differences are largest on governance dimensions. Americans prefer full privacy 7.8 percentage points more than Indians (59.8% vs. 52.0%, $p < .001$) and prefer designs where government cannot block transactions by 6.6 percentage points (58.2% vs. 51.6%, $p < .001$). Indians show flat preference gradients across privacy levels (3.2pp range) and government control levels (2.4pp range), suggesting

these features do not meaningfully differentiate their choices.

Finding 3: CBDC governance preferences are collectively held within countries. Respondents embedded in the same institutional context converge on similar evaluative standards regardless of their individual trust levels, fraud experience, privacy confidence, or cryptocurrency usage. This consensus is documented across 36 within-country moderation tests, where only one reaches statistical significance (2.8%), a rate consistent with chance rather than meaningful moderation. Trust in central banks, fraud victimization, and cryptocurrency experience all produce identical preference structures within countries, confirming that the forces shaping CBDC governance preferences operate at the institutional rather than individual level.

Finding 4: The single exception is institutional privacy distrust in the USA. Low-confidence Americans—those who doubt the Federal Reserve will protect their financial privacy—prefer full privacy 4.9 percentage points more than high-confidence Americans (62.7% vs. 57.7%, $p = .015$). This finding is intuitive: users who distrust the central bank's commitment to privacy protection demand privacy-by-design features that do not depend on institutional trustworthiness. Architecturally private systems substitute for institutional credibility.

Finding 5: Ecosystem context dominates individual psychology. The pattern of cross-country divergence combined with within-country consensus suggests that payment infrastructure maturity and institutional familiarity shape preferences more powerfully than individual-level characteristics. India's UPI success has normalized government involvement in digital payments; the USA's fragmented system and surveillance debates produce demand for privacy and autonomy. These ecosystem-level factors appear to calibrate what "acceptable" CBDC design means, regardless of individual attitudes.

1.4 Contributions

This paper makes four contributions to the economics of information security and CBDC literatures.

Methodological contribution. This paper provides the first multi-attribute conjoint experiment comparing CBDC design preferences across the United States and India. Conjoint methods force real trade-offs among bundled features, overcoming the limitations of Likert-scale surveys that treat security, privacy, and trust as separate latent constructs.

Empirical contribution. This paper documents the consensus effect: within-country homogeneity of preferences despite substantial cross-country divergence. The 97.2% null rate across moderation tests is itself a substantive empirical finding.

Theoretical contribution. This paper adjudicates between individual-difference and contextual theories of technology preferences. The results support contextual accounts: the meaning of CBDC features is constructed relative to existing payment infrastructure, not evaluated against abstract standards.

Policy contribution. This paper identifies institutional privacy distrust as a barrier that cannot be addressed through trust-building alone. Users who doubt that the central bank will protect their financial privacy demand architecturally private systems – privacy-by-design rather than privacy-by-promise. For these users, stated commitments to data protection are insufficient; only structural guarantees, such as strong encryption, data minimization, and limited data retention, are likely to be credible and effective.

1.5 Paper Structure

The remainder of this paper proceeds as follows. Section 2 reviews the CBDC adoption literature and develops the theoretical framework. Section 3 describes the conjoint experimental design, sampling strategy, and analytical approach. Section 4 presents results, beginning with overall attribute preferences, then cross-country comparisons (RQ5), and finally within-country heterogeneity tests (RQ1–RQ4). Section 5 discusses theoretical implications, policy recommendations, and limitations. Section 6 concludes.

2 Literature Review and Theoretical Framework

This section reviews three bodies of literature relevant to CBDC design preferences: (1) empirical studies of CBDC adoption and acceptance; (2) research on privacy, surveillance, and governance trade-offs in digital payments; and (3) comparative work on payment ecosystems and institutional trust. Four theoretical perspectives are then synthesizes that generate competing predictions, formalize the research questions, and specify hypotheses.

2.1 CBDC Adoption: From Technology Acceptance to Governance Design

2.1.1 Standard Adoption Determinants

The empirical literature on CBDC adoption has grown rapidly, drawing heavily on technology acceptance frameworks. Studies across China, India, the Netherlands, and other contexts consistently find that perceived usefulness, ease of use, and trust predict adoption intentions (Søilen & Benhayoun, 2021; Ogunmola & Das, 2024; Liu et al., 2024; Bijlsma et al., 2021). In India specifically, TAM- and UTAUT-based studies identify performance expectancy, facilitating conditions, government support, and trust as central drivers of digital rupee adoption (Desai, 2024; Kaur et al., 2024; Balasubramanian & Thirumaran, 2025). Chinese e-CNY research similarly emphasizes perceived security, awareness, and institutional support (Liu et al., 2024; Tronnier and Qiu, 2024).

However, CBDC adoption diverges from generic fintech acceptance in two critical respects. First, institutional trust functions as a necessary condition rather than a simple predictor. Søilen & Benhayoun show that performance expectations, social influence, and facilitating conditions increase adoption only when trust in the CBDC system is sufficiently high. Second, existing payment ecosystems fundamentally shape perceived value. Gupta et al. find that prior UPI experience in India negatively moderates the effect of performance expectancy on CBDC behavioral intention. Similarly, Fairweather et al. show that in Australia, where bank deposits are perceived as very safe, CBDC safety adds little marginal value.

2.1.2 Beyond Adoption: Governance as Design Space

A growing strand of research reconceptualizes CBDC not merely as a technology to adopt but as a governance architecture to evaluate. Agur et al. demonstrate that CBDC design involves trade-offs between privacy, financial stability, and monetary policy transmission. Privacy has become a central design concern as central banks move toward implementation (Jiang, 2023). Proposals for asymmetric privacy designs seek to protect consumer transaction data while preserving regulatory oversight capabilities

(Tinn & Dubach, 2021).

This governance-design perspective implies that users do not simply accept or reject CBDC, but evaluate bundles of attributes, including who can access transaction data, under what conditions, and with what recourse mechanisms. Yet most empirical work treats privacy, trust, security, and cost as separate latent constructs, underspecifying the attribute-level trade-offs users make when confronted with realistic CBDC configurations.

2.2 Privacy, Security, and Surveillance

2.2.1 Experimental Evidence on Privacy Valuation

A distinct literature examines how users trade off privacy against functionality in CBDC design. This body of work is particularly relevant for interpreting the findings, as it relies on experimental or quasi-experimental methods to isolate the causal role of privacy features in shaping adoption intentions and acceptance. Three experimental studies provide particularly relevant evidence for situating the present findings and illustrate both the causal role of privacy in CBDC acceptance and the context-dependency of that role.

Choi et al. show, through a nationally representative randomized experiment in Korea (N = 3,500), that stronger privacy protections causally increase willingness to use CBDC, with adoption intentions rising by as much as 64 percent for sensitive transactions such as medical payments. This finding establishes that privacy is not merely an attitudinal concern but a design feature that materially shifts behavioral intentions in specific use contexts. Fairweather et al. reach a complementary finding in Australia through a discrete choice experiment: in an environment where bank deposits are already perceived as very safe, incremental CBDC safety has no average value, but privacy configurations remain a salient margin of differentiation. A similar pattern emerges in Austria, where Abramova et al. find that respondents broadly satisfied with existing payment options express limited interest in a digital euro, and where Elsinger et al., using a discrete choice experiment with 1,421 Austrian residents, confirm that security and financial incentives are the strongest adoption drivers while privacy plays a secondary role. These Austrian findings converge with the present study in placing functional attributes above governance concerns, but diverge on the relative salience of privacy, which the present study finds to be the single most differentiating attribute between the US and India. The divergence likely reflects contextual differences, as Austria's high institutional trust and GDPR-backed privacy regime may render privacy less salient as a differentiating feature than it is in the more polarized US environment.

Mehlkop et al., in the only prior randomized design spanning both the United States and India, find that provider identity, whether the system is operated by the state or by a multinational technology company, is a first-order determinant of acceptance that often matters as much as abstract privacy levels. More recent conjoint and discrete choice experiments, including work on the Digital Shekel in Israel (Plato-Shinar and Maman, 2025) and on payment method preferences in Korea (Choi et al., 2026), confirm the viability of experimental approaches for eliciting CBDC design preferences and reinforce the finding that functional attributes dominate stated choices across diverse institutional settings.

Taken together, these experimental studies establish that privacy features can causally affect CBDC acceptance, but that privacy valuation is deeply context-dependent, shaped by existing infrastructure qual-

ity, provider identity, and institutional trust rather than by stable individual preferences. This literature motivates the focus on governance-laden design attributes and supports the expectation that privacy preferences cannot be understood independently of the institutional environment in which a CBDC would be deployed.

2.2.2 Trust and Privacy: A Non-Monotonic Relationship

A growing body of work examines whether stated privacy preferences translate into concrete CBDC design choices, often documenting what has been termed a "privacy paradox." The broader privacy literature has established that individuals frequently behave in ways inconsistent with their stated privacy concerns, a pattern shaped by contextual and situational factors rather than stable individual preferences (Acquisti et al., 2015). In the CBDC domain, this pattern recurs across multiple empirical settings: respondents express strong abstract concerns about privacy yet frequently select functional CBDC designs over more anonymous alternatives when faced with concrete trade-offs (Koziuk and Ivashuk, 2022; Koziuk, 2021). Trust in central banks does not reliably predict preferences for anonymity, and individuals who report high privacy concern often still favor designs emphasizing usability, convenience, or institutional oversight. A similar disconnection between trust and privacy behavior appears in China's e-CNY context, where Tronnier and Qiu (2024) find that soft and hard trust factors have limited explanatory power for either privacy concerns or actual usage patterns, suggesting that privacy perceptions operate independently of institutional trust.

Whether this paradox reflects genuine inconsistency in preferences or analytical limitations remains an open question. Koziuk et al. (2025) apply machine-learning techniques to CBDC privacy preferences and find complex, nonlinear interactions between privacy indices, institutional trust, and design choices that standard linear moderation models do not capture. When preferences depend on higher-order interactions or threshold effects, linear models may obscure meaningful patterns, raising the possibility that the privacy paradox partly reflects the limitations of conventional analytical approaches rather than true contradictions in user preferences.

However, an alternative interpretation is that the privacy paradox reflects measurement problems rather than genuine contradictions in user preferences. Some studies apply contextual integrity theory to CBDC-like systems, particularly in the context of the digital euro (Tronnier et al., 2023). Their work demonstrates that privacy acceptability is highly conditional on specific informational norms, including which actors (central banks, commercial banks, tax authorities, law enforcement, or foreign entities) access which types of data (identity, transaction amounts, metadata) and under what conditions (automatic access, AML-triggered access, or court-ordered disclosure). Simplified framings that contrast "high" versus "low" privacy obscure these distinctions and risk mischaracterizing user preferences.

From a CBDC design perspective, this literature implies that privacy preferences are multidimensional and context-specific rather than binary. Preferences hinge on the governance of information flows rather than on abstract notions of anonymity alone. This insight motivates the use of conjoint experiments that explicitly vary who accesses what information under what conditions, allowing researchers to map privacy preferences at the attribute level rather than relying on coarse attitudinal measures.

The relationship between institutional trust and privacy preferences further complicates standard theoretical expectations. In the context of China's e-CNY pilot, Tronnier and Qiu show that privacy concerns

significantly influence actual usage, while both soft trust (general confidence in institutions) and hard trust (belief in technical safeguards) have limited explanatory power for either privacy concerns or usage patterns. These findings diverge from models that treat institutional trust as a straightforward precondition for CBDC acceptance or privacy tolerance.

Cross-country evidence reinforces this non-monotonic relationship. Studies by Koziuk and Ivashuk show that high trust in central banks as privacy guarantors does not systematically predict preferences along the anonymity–functionality dimension (Koziuk and Ivashuk, 2022; Koziuk, 2021). Many respondents who express trust in central banks and concern about privacy nonetheless choose designs with greater functionality and institutional access over fully anonymous options. This pattern suggests that trust may enable willingness to engage with CBDC as a category, but does not anchor where individuals fall on the surveillance–autonomy continuum.

For this paper, I treat trust to operate as a background condition rather than a direct driver of specific privacy design preferences. Institutional trust may lower barriers to adoption by reducing perceived systemic risk, but it does not uniquely determine how users evaluate trade-offs between privacy, functionality, and oversight. This insight directly informs the empirical strategy, which treats trust as a potential moderator rather than assuming a monotonic relationship between trust and privacy demand.

2.3 Payment Ecosystems and Cross-Country Context

2.3.1 India: UPI as Baseline Infrastructure

India presents a distinctive and analytically important context for evaluating CBDC design preferences. The Unified Payments Interface (UPI) has achieved unprecedented penetration as a national payment rail: it is zero-fee for users, instant in settlement, fully interoperable across banks and apps, and ubiquitous across merchants as well as peer-to-peer transactions. As a result, UPI functions as baseline payment infrastructure rather than a niche innovation. Multiple studies document that Indian users often struggle to identify what additional value a CBDC would provide beyond UPI, particularly in everyday retail payments (Krishnamoorthy & Aggarwal, 2024; Bhatnagr, 2025). Evidence from the digital rupee pilot suggests that low uptake has been driven less by outright resistance and more by limited awareness and confusion about how CBDC differs functionally from existing UPI-based payments.

Gupta et al. provide a key theoretical mechanism explaining this pattern. They show that prior experience with UPI dampens the marginal perceived benefit of CBDC and weakens social influence effects that typically drive adoption of new payment technologies. In a mature instant-payment ecosystem, CBDC must offer clearly differentiated value in order to gain traction. Potential sources of such differentiation include offline capability, programmable or conditional transfers, distinctive privacy regimes, or cross-border functionality. Absent these features, CBDC risks being perceived as redundant rather than transformative.

India's broader digital public infrastructure further shapes CBDC evaluation by normalizing government involvement in financial data. Aadhaar-linked UPI operates with mandatory KYC requirements, transaction monitoring, and data-sharing frameworks embedded within state-branded payment rails. A growing literature emphasizes that government support, institutional legitimacy, and compatibility with existing state infrastructure are central drivers of digital payment adoption in India (Desai, 2024; Bala-

subramanian & Thirumaran, 2025; Dixit et al., 2025) even though the adoption remains minimal. In this context, trust in public institutions may condition not whether users are willing to engage with CBDC, but rather which governance features they treat as baseline expectations versus meaningful points of differentiation.

Data from this study's India sample illustrate this institutional environment concretely. Among Indian respondents, 94.6% trust the Reserve Bank of India to manage a digital currency, 90.6% express confidence in the RBI's ability to protect their financial privacy, and 96.5% report using digital payments daily or weekly. These figures are not merely moderator distributions that constrain statistical power in the heterogeneity analyses that follow; they describe a population whose relationship with state digital infrastructure is one of routine, functional engagement, in which government data collection is an embedded feature of everyday financial life rather than an exceptional intrusion. India nonetheless remains a deeply heterogeneous polity with respect to digital finance. While UPI penetration is near-universal in urban areas, cash still accounts for a substantial share of rural transactions, and over 300 million Indians remain outside the formal banking system. The India sample, which was recruited online, conducted in English, and drawn predominantly from urban and highly educated populations, captures the digitally engaged segment most likely to encounter CBDC in the near term but cannot speak for populations still navigating the transition to digital payments. As the results will show, the cash-versus-digital divide within India proves to be among the strongest sources of within-country preference heterogeneity, suggesting that this internal frontier is consequential for understanding where design preferences diverge.

2.3.2 United States: Fragmented Rails and Surveillance Contestation

The United States presents a sharply contrasting payment ecosystem. Retail payments are dominated by private card networks that impose merchant fees of approximately 1.5-2.5 percent, while bank-based ACH transfers remain relatively slow, typically settling over one to three business days. The overall landscape is fragmented across banks, fintech firms, and platform wallets, with no unified, zero-fee, instant public payment rail comparable to UPI. From a purely functional perspective, this fragmentation creates space for CBDC to fill genuine infrastructure gaps, including free instant settlement, a public-option payment rail, and more efficient delivery of government transfers.

At the same time, the US context is characterized by polarized and contested attitudes toward federal involvement in financial data. Empirical work shows that trust is fragmented across institutional actors: some users express greater trust in private platforms than in the federal government for data handling, while others exhibit the opposite pattern, often aligned with ideological or partisan orientations (Mehlkop et al., 2023). Privacy occupies contested political terrain, shaped by post-Snowden surveillance debates, high-profile data breaches, growing state-level privacy legislation, and ongoing discussions about financial monitoring and enforcement. The political environment surrounding CBDC in the United States has become distinctly adversarial in recent years. Executive Order 14178, signed in January 2025, explicitly prohibits federal agencies from establishing, issuing, or promoting a retail CBDC, while legislative developments including the GENIUS Act of 2025 position private-sector stablecoins as the politically preferred alternative to government-issued digital money. These policy actions are not merely background context for interpreting the survey results; they reflect and amplify an institutional environment in which government-issued digital currency is coded as a surveillance tool rather

than a public good.

Data from this study's US sample reflect this divided landscape. Trust in the Federal Reserve to manage a digital currency stands at 56.6%, and confidence in the Fed's ability to protect financial privacy at only 43.6%, a roughly even split that contrasts sharply with India's near-unanimity. This distributional variation creates the moderator heterogeneity needed for within-country analysis, but it also reflects a substantive institutional reality: attitudes toward government involvement in financial transactions are genuinely contested in the United States in a way that they are not in India.

2.3.3 Comparative Predictions

The Indian and U.S. payment ecosystems generate clear comparative predictions for how users evaluate CBDC governance features. In contexts where existing payment systems are already high-quality and widely trusted, such as India's UPI-dominated ecosystem, incremental CBDC features are likely to be valued less unless they offer distinctive governance or functionality advantages. In contrast, where payment systems are fragmented and costly, as in the United States, CBDC may be seen as filling genuine infrastructure gaps but also faces heightened scrutiny over privacy, government control, and surveillance.

The U.S. CBDC policy landscape reflects this tension. On January 23, 2025, the United States issued Executive Order 14178, titled "Strengthening American Leadership in Digital Financial Technology," which explicitly prohibits federal agencies from establishing, issuing, or promoting a retail CBDC within the United States and halts ongoing CBDC initiatives at agencies, effectively positioning the U.S. as an outlier among advanced economies on this front (Executive Order 14178, 2025). At the same time, federal policy has shifted toward promoting responsible growth of digital assets and establishing regulatory clarity for crypto and payment stablecoins as the preferred digital payment vehicles under new legislation such as the GENIUS Act of 2025, which creates a comprehensive regulatory framework for payment stablecoins (GENIUS Act, 2025).

Despite formal rejection of a U.S. CBDC, deliberations continue among policymakers, academics, and industry groups about the broader digital payment future, including stablecoin regulation, hybrid public–private monetary models, and the role of the Federal Reserve in supporting private digital money (Krause, 2025). This does not disregard the findings of the paper. These developments suggest that while a retail CBDC may be politically and legally constrained in the U.S., general attitudes toward digital payment infrastructure, privacy, and institutional trust remain central to user evaluations of any proposed payment innovation.

Comparative consumers' mental models of "state versus platform" and "public versus private money" are therefore anchored differently across contexts. Identical CBDC governance features will be evaluated against different reference points: in India, the benchmark is a trusted, government-branded rail normalized through everyday use; in the U.S., the benchmark is a fragmented, private-sector ecosystem supplemented by emerging stablecoin frameworks. This contextual framing motivates the cross-country experimental design and underpins expectations of systematic divergence in CBDC preference structures across institutional environments.

2.4 Theoretical Framework: Competing Perspectives

The preceding literature reveals a field with robust findings about what drives CBDC adoption, including performance, cost, and trust, but with limited understanding of how users navigate the governance trade-offs embedded in CBDC design. Four theoretical perspectives, organized here into two meta-categories, generate competing predictions about which attributes should dominate preferences, how individual characteristics should moderate those preferences, and what pattern of cross-country differences should emerge.

2.4.1 Individual-Level Perspectives

Three theoretical traditions predict that measurable individual characteristics should differentiate CBDC design preferences within countries. They disagree on which characteristics matter most but share an underlying assumption: that preferences are individually constructed from personal experiences and attitudes, and that identifying the relevant individual-level moderators is the key to understanding preference heterogeneity.

The security economics perspective (Anderson, 2001; Kahn & Roberds, 2009) predicts that users prioritize fraud protection, reliability, and cost above abstract governance concerns, and that users with direct fraud experience should weight security features more heavily, potentially accepting greater oversight if it enhances protection. This perspective aligns with the broad TAM/UTAUT finding that perceived usefulness, security, and cost consistently predict digital payment adoption across contexts.

The privacy and autonomy perspective (Solove, 2006; Nissenbaum, 2009) predicts that privacy protection should strongly increase adoption, that government control and surveillance should be viewed as costs rather than benefits, and that users who are confident in their ability to protect their privacy should differ meaningfully from those who are not. This perspective aligns with experimental evidence that privacy design causally shifts CBDC acceptance (Choi et al., 2026; Fairweather et al., 2024).

The institutional trust perspective (Levi & Stoker, 2000) predicts that trust in central banks should enable willingness to accept government-operated systems and data collection, such that high-trust users tolerate more government control while low-trust users demand architectural protections. However, the privacy-paradox literature suggests this relationship may not be monotonic, as trust in central banks does not reliably predict preferences along the anonymity-functionality dimension (Koziuk and Ivashuk, 2022).

2.4.2 Contextual and Institutional Perspectives

A contrasting set of traditions predicts that CBDC preferences are shaped by shared institutional experience rather than by individual variation. Economic sociology argues that economic behavior is embedded in social structures and cannot be understood as the product of atomized individual calculation (Granovetter, 1985). The moral economy tradition holds that communities converge on shared normative expectations about what constitutes legitimate economic practice, expectations that are collectively maintained and that govern evaluations of new economic arrangements (Thompson, 1971). Domestication theory in science and technology studies argues that technology meaning is constructed through situated practice within institutional contexts, such that the same technology carries different signifi-

cance depending on the infrastructure and social relationships into which it is introduced (Silverstone and Hirsch, 1992; Hyysalo, 2010). And comparative payment-systems research argues that CBDC is always evaluated relative to existing alternatives, against UPI in India and against cards and wallets in the United States, rather than against abstract standards (Gupta et al., 2024; Tronnier and Qiu, 2024).

Applied to CBDC, this contextual perspective predicts that cross-country differences in attribute weights should be larger than within-country heterogeneity by individual characteristics, because the institutional context that shapes baseline expectations is shared within countries but diverges between them. Attributes offering clear improvements over incumbent systems should dominate preferences, but what constitutes an improvement varies by context. Within any given country, shared ecosystem experience should produce convergent preferences regardless of individual characteristics.

2.5 Research Questions and Hypotheses

The inquiry is organized around five research questions that map onto the theoretical tensions outlined above. RQ5 (cross-country comparison) establishes baseline differences, while RQ1–RQ4 test within-country heterogeneity by individual characteristics.

1. **RQ1: Trust in Central Bank $\times$ Government Control**

    - **RQ1**: Do individuals with higher trust in central banks tolerate more government control in CBDC design?

    Institutional trust theory predicts that high-trust individuals should accept government oversight (H1a). However, privacy-paradox evidence suggests trust may not reliably predict control preferences once security and privacy are accounted for (H1b).

    - **H1a (Trust enables control)**: Individuals with high trust in central banks will exhibit stronger preferences for government-controlled CBDC designs.
    - **H1b (Trust is orthogonal)**: Trust in central banks will not systematically predict preferences for government control.

2. **RQ2: Trust in Digital Payments and Fraud Experience $\times$ Security**

    - **RQ2a**: Do individuals with higher trust in digital payments prioritize security less?
    - **RQ2b**: Do fraud victims prioritize security more?

    Security-economics theory predicts that fraud victims should weight security features more heavily (H2b) and that high-trust users should be less sensitive to incremental security improvements (H2a).

    - **H2a (Trust reduces security sensitivity)**: Individuals with high trust in digital payments will exhibit weaker preferences for high-security CBDC designs.
    - **H2b (Fraud increases security sensitivity)**: Fraud victims will exhibit stronger preferences for high-security CBDC designs.

3. **RQ3: Privacy Confidence $\times$ Privacy Features**

    - **RQ3**: Do individuals with higher confidence that the central bank will protect their financial

privacy demand more or less privacy-preserving CBDC features?

Institutional trust theory predicts that users who trust the central bank to protect their privacy should be more willing to accept conditional privacy or visibility designs relying on the institution to use the data responsibly. Privacy-skeptical users, by contrast, may demand architecturally private systems that do not depend on institutional trustworthiness (H3a). Alternatively, privacy-confident users may value privacy intrinsically and demand privacy-preserving features regardless of trust (H3b).

- **H3a (Distrust increases privacy demand)**: Individuals with low confidence in central bank privacy protection will exhibit stronger preferences for privacy-preserving CBDC features, substituting architectural guarantees for institutional trust.
- **H3b (Trust and privacy are orthogonal)**: Confidence in central bank privacy protection will not systematically predict preferences for privacy features.

4. **RQ4: Cryptocurrency Usage $\times$ Governance Preferences**
    - **RQ4a**: Do cryptocurrency users demand less government control in CBDC design?
    - **RQ4b**: Do cryptocurrency users demand more privacy-preserving features?

    Cryptocurrency users have revealed preferences for decentralized, pseudonymous payment systems. They may therefore demand more privacy and less government control in CBDC design (H4a, H4b). Alternatively, if crypto users are primarily technology adopters seeking financial access, they may accept any well-designed system regardless of governance structure.

    - **H4a (Crypto users demand autonomy)**: Cryptocurrency users will exhibit stronger preferences for limited government control.
    - **H4b (Crypto users demand privacy)**: Cryptocurrency users will exhibit stronger preferences for privacy-preserving features.

5. **RQ5: Cross-Country Divergence**
    - **RQ5**: Do CBDC design preferences differ systematically between the United States and India?

    Contextual theory predicts that baseline payment ecosystems and institutional norms should create large cross-country differences (H5a). India's UPI-saturated, state-branded infrastructure may normalize government involvement, while US card-centric systems and surveillance debates may heighten sensitivity to privacy and control. Alternatively, privacy and security may be universal concerns that dominate in both contexts (H5b).

    - **H5a (Divergence dominates)**: Differences in attribute preferences between the United States and India will be larger than within-country heterogeneity by trust, fraud, or demographics.
    - **H5b (Universal salience)**: Privacy and security will dominate preferences in both countries, with smaller cross-country differences than within-country heterogeneity.

3 Research Design and Methods

3.1 Study Design

As mentioned in the introduction, this study uses a conjoint survey experiment to examine user preferences over the design of a retail central bank digital currency. Conjoint analysis is particularly well suited for this research question because CBDC adoption decisions involve trade-offs across multiple, simultaneously bundled design features such as security, privacy, government control, cost, and usability. Rather than eliciting preferences for individual features in isolation, the conjoint design allows estimation of how users value specific design components when forced to choose between realistic policy configurations. This multidimensional choice structure avoids the artificial decomposition common to Likert scale surveys (Hainmueller et al., 2014) that elicit attitudes about privacy "in general" or trust "on average."

Randomization of attribute levels enables causal identification of the marginal effect of each feature on choice probabilities while keeping all other attributes constant (Hainmueller et al., 2014; Bansak et al., 2021). Unlike structural equation models, which infer preferences through correlational patterns in latent constructs, conjoint designs estimate average marginal component effects (AMCEs), the causal impact of shifting from one attribute level to another. This is particularly valuable for CBDC design because policymakers require evidence on the expected adoption consequences of specific design choices, for example, increasing state access to transaction data or tightening privacy protections.

Finally, conjoint experiments can mitigate social desirability bias by embedding sensitive considerations such as acceptance of government surveillance within multi attribute profiles, reducing respondents' ability to infer which feature is the focal treatment (Horiuchi et al., 2022). This is critical for CBDC governance preferences, where stated attitudes about privacy or state control may be shaped by perceived norms.

For this study, respondents were presented with repeated choice tasks in which they selected between two hypothetical digital currency alternatives. Each alternative was defined by seven attributes: transaction security, privacy, government control, cross-border usability, integration with existing payment platforms, transaction cost, and transaction speed. Each attribute had three substantively significant levels, designed to reflect policy-relevant design choices currently under discussion in debates around retail CBDCs.

Attribute levels were independently randomized between profiles and between choice tasks. Each respondent completed eight choice tasks, with two alternatives per task, yielding 16 profile evaluations per respondent. The order of choice tasks and the left-right positioning of alternatives were randomized to minimize ordering and presentation effects. Post-fielding checks confirmed that attribute levels were approximately evenly distributed across the sample, consistent with successful randomization.

Respondents were instructed to evaluate each pair of alternatives as hypothetical digital payment instruments and to select the option they would prefer to use. The conjoint task was framed as a choice between payment systems. So instead of testing a single adoption hypothesis, this study examines how users in different institutional contexts evaluate trade-offs across CBDC design features.

The design builds on emerging CBDC preference elicitation work. Fairweather et al. (2024) use discrete

choice experiments to value safety versus privacy trade offs in Australia and find that privacy configurations can matter more than incremental safety in high trust environments. Choi et al. (2026) randomize privacy levels and information treatments and show that privacy design causally shifts willingness to adopt. This paper extends this work by expanding the attribute space to include government control, welfare integration, offline functionality, and institutional intermediation; conducting parallel experiments in the United States and India to leverage ecosystem differences; and estimating heterogeneous effects by individual level trust, fraud experience, privacy confidence, and cryptocurrency usage.

3.2 Sampling and Data Collection

The survey was administered online to respondents in two countries: India and the United States. In India, respondents were recruited via Conjointly, yielding 600 completed responses. In the United States, respondents were recruited via Prolific, yielding an additional 600 completed responses. Data collection took place in November 2025.

Sample Characteristics:

- **United States:** $N = 600$ recruited, $N = 539$ final sample
- **India:** $N = 600$ recruited, $N = 608$ final sample
- **Total:** $N = 1{,}147$ respondents
- **Total observations:** 18,352 profile-level choices (1,147 $\times$ 16 profiles)

Each respondent completed the same conjoint experiment and accompanying survey questions. Respondents who did not complete the conjoint tasks were excluded from the final sample.

Quality Filters Applied:

- Completed all 8 conjoint tasks (no mid-survey dropouts)
- Response time within reasonable bounds (10–45 minutes; excludes speeders and interruptions)

Survey Structure: The survey consisted of three modules:

1. **Screening:** Age 18+, resident of US/India with a representative sample across gender, income, education, and states.
2. **Conjoint experiment:** 8 choice tasks, each presenting two CBDC profiles with randomized attributes. Forced choice between two profiles (no opt-out).
3. **Additional Questions:** Demographics, digital payment behavior, technological comfort, cross-border transaction experience, cryptocurrency familiarity, trust in financial institutions, attitudes toward government involvement in payments. Collected after conjoint tasks.

In addition to the conjoint tasks, the survey included questions measuring demographic characteristics, and a structured set of questions capturing respondents' digital payment behavior, technological comfort, cross-border transaction experience, familiarity with cryptocurrencies, trust in financial institutions, and attitudes toward government involvement in payments. These measures were designed to capture

individual-level constraints, institutional trust, and exposure to alternative payment systems. They are not treated as outcomes of the conjoint experiment, but are used to examine heterogeneity in preferences over CBDC design attributes.

3.3 CBDC Attributes and Levels

The designed attribute levels through an iterative process combining (1) review of CBDC pilots and proposals, including the e-CNY, the digital euro, and the digital rupee; (2) pilot testing with 50 respondents per country to ensure comprehension and attribute salience. The final design includes eight attributes spanning governance and functional dimensions.

Governance attributes (A1–A3) are theoretically motivated by competing frameworks and allow direct tests of security economics, privacy autonomy, and institutional trust predictions. Functional attributes (A4–A8) ground the experiment in realistic design debates and reduce the risk that respondents evaluate governance in isolation from cost and usability.See Table 1 for a full description.

3.4 Sample Characteristics

The two samples differ in ways consistent with the comparative payment ecosystems (See Table 2). A large share of respondents in India use UPI as their primary payment method, while the United States sample relies primarily on cards and wallets. Trust distributions also differ markedly, with a ceiling effect in India and substantially more distrust in the United States, motivating value based subgroup splits (Low: 1–2; Medium: 3; High: 4–5). Fraud exposure is broadly comparable across countries, providing leverage for testing security related heterogeneity. Both samples are more educated and digitally engaged than the general population, consistent with online recruitment, which limits generalizability to non users while capturing populations most likely to be early CBDC adopters.

3.5 Analytical Strategy

3.5.1 Primary Analysis: Marginal Means (MM)

The primary outcome variable is a binary indicator of profile selection. For each choice task, the selected alternative is coded as 1 and the non-selected alternative as 0. This binary choice outcome allows estimation of the causal effect of individual attribute levels on the probability that a given profile is chosen. The primary analysis uses Marginal Means (MM) to understand baseline preference levels across attribute configurations. Marginal means represent the average predicted probability that a profile containing a given attribute level is chosen, averaged over the distribution of all other attributes:

$$\overline{MM}_{k\ell} = \frac{\sum_{i,t} 1(\text{Attribute}_{k,it} = \ell) \times \text{Choice}_{it}}{\sum_{i,t} 1(\text{Attribute}_{k,it} = \ell)} \tag{1}$$

where i indexes respondents, t indexes choice tasks, k indexes attributes, and ℓ indexes levels.

$\overline{MM}_{k\ell}$ is the proportion of times level ℓ was chosen when presented. Unlike Average Marginal Component Effects (AMCEs), marginal means do not depend on the choice of a reference category and provide a measure of absolute preference levels. MM is particularly useful for assessing the relative importance of attributes in shaping baseline adoption preferences.

Table 1: CBDC Design Attributes and Levels

Attribute	Levels	Policy Interpretation
Security (A1)	L1: Very low chance of fraud with reversible transactions L2: Fraud risk similar to cash with reversible transactions L3: Fraud risk similar to cash with irreversible transactions	Captures fraud protection and transaction reversibility. L1 provides bank-grade protection; L3 mimics cash finality. Tests whether security concerns dominate governance preferences.
Privacy (A2)	L1: Full privacy with no third-party access L2: Identity and transaction details visible to recipient, accessible to government under specific conditions L3: Identity and transaction details accessible to banks, government, and third parties	Operationalizes surveillance vs. autonomy trade-off. L1 is cash-like anonymity; L3 is full financial transparency. Central to testing privacy-autonomy vs. institutional-trust theories.
Government Control (A3)	L1: Government cannot block or reverse transactions L2: Government can block specific transactions with law enforcement L3: Government can block or reverse any transaction at will	Degree of state oversight and intervention power. L1 is decentralized/crypto-like; L3 is maximal state control. Tests institutional trust predictions.
Cross-border Availability (A4)	L1: Works worldwide for international transactions L2: Limited to certain regions L3: Only usable domestically	Tests whether international usability drives adoption. Relevant for remittances, travel, cross-border commerce.
Integration with Existing Platforms (A5)	L1: Works with existing bank and payment apps L2: Requires downloading a new app with immediate use L3: Requires downloading a new app with a setup process	Compatibility with existing payment infrastructure. Critical in India (UPI integration) vs. US (bank app integration). Tests ecosystem/contextual theory.
Transaction Cost (A6)	L1: Free of charge L2: Flat 4.5 fee per transaction L3: One percent fee on transaction amount	Cost relative to existing methods. In India, UPI is zero-fee; in US, cards charge 1.5–2.5%. Tests whether cost dominates governance concerns.
Transaction Speed (A7)	L1: Instantaneous L2: Completed in a few hours L3: Takes one to two business days	Settlement speed. UPI in India is instant; US ACH takes days. Tests infrastructure expectations shaped by existing systems.

The unit of analysis is the profile-level choice: each respondent evaluates 16 profiles across 8 choice tasks, and each profile is coded as chosen (1) or not chosen (0). Standard errors are computed via clustered bootstrap with 1,000 replications, clustering at the respondent level to account for the non-independence of repeated choices by the same individual. Cross-country and subgroup comparisons are estimated via separate marginal mean calculations for each group, with post-hoc two-sample Z-tests for differences in proportions rather than pooled models with interaction terms. This approach was preferred because marginal means provide intuitive absolute preference levels for each subgroup, facilitating direct comparison of choice probabilities across countries and respondent types.

3.6 Cross-Country Comparisons (RQ5)

To test whether preferences differ systematically between the United States and India, country-specific marginal means is estimated and two-sample Z-tests are conducted for differences in proportions:

Table 2: Sample Characteristics by Country

Characteristic	USA (N=539)	India (N=608)	Interpretation
Demographics			
Age (mean)	38.2 years (SD=12.3)	31.5 years (SD=9.7)	India sample skews younger
Female (%)	52.1%	46.2%	Both roughly balanced
College degree+ (%)	68.4%	71.3%	Both highly educated samples
Digital Payment Ecosystem			
Use digital weekly+ (%)	87.6%	94.1%	India higher frequency usage
Primary: Card/Wallet	76.3%	18.2%	Stark ecosystem difference
Primary: UPI/Instant payment	8.4%	78.5%	India = UPI-dominant
Ever fraud victim (%)	30.4%	34.5%	Comparable victimization rates
Cryptocurrency Experience			
Own or use crypto (%)	51.0%	67.6%	India shows higher crypto adoption
Trust and Confidence (1–5 scale)			
Trust central bank (digital)	M=3.12 (SD=1.04)	M=4.58 (SD=0.71)	India dramatically higher
Trust digital payments	M=3.45 (SD=0.98)	M=4.38 (SD=0.76)	India more trusting of platforms
Privacy confidence	M=2.89 (SD=1.15)	M=4.41 (SD=0.82)	USA much less confident
Trust Distribution Details			
Trust CB: % max rating (5)	9.3%	66.6%	India severe ceiling effect
Trust CB: % low (1–2)	50.3%	5.9%	USA bimodal with plurality distrust
Trust CB: % medium (3)	31.2%	10.5%	USA has large "ambivalent" group

$$Z_{k\ell} = \frac{\overline{MM}_{k\ell}^{\text{USA}} - \overline{MM}_{k\ell}^{\text{India}}}{\sqrt{\text{SE}(\overline{MM}_{k\ell}^{\text{USA}})^2 + \text{SE}(\overline{MM}_{k\ell}^{\text{India}})^2}} \tag{2}$$

where standard errors are calculated via clustered bootstrap.

Multiple testing correction. With 21 attribute levels tested (7 attributes $\times$ 3 levels each), Bonferroni correction is applied to control the family-wise error rate at $\alpha = 0.05$. The corrected significance threshold is $\alpha/21 \approx 0.002$.

Statistical power. With $N = 539$ (USA) and $N = 608$ (India) and approximately 2,800–3,200 observations per attribute level, this design has 80% power to detect differences ≥ 3.2 percentage points at the Bonferroni-corrected α level.

3.7 Subgroup Analyses (RQ1–RQ4)

To explore preference heterogeneity, I examine how marginal means vary across respondent subgroups defined by theoretically relevant individual characteristics:

- **RQ1:** Trust in central bank (digital money) × Government control preferences
- **RQ2a:** Trust in digital payments × Security attribute preferences
- **RQ2b:** Fraud victimization × Security attribute preferences
- **RQ3:** Privacy confidence × Privacy attribute preferences
- **RQ4a:** Cryptocurrency usage × Government control preferences
- **RQ4b:** Cryptocurrency usage × Privacy preferences

Subgroup definition. For ordinal variables measured on five-point scales (trust, confidence), respondents are classified into value-based groups:

- **Low:** Scale values 1–2
- **Medium:** Scale value 3
- **High:** Scale values 4–5

This value-based approach (rather than median splits) is necessitated by the India trust ceiling effect, where 67% of respondents gave maximum ratings.

4 Findings

The results speak directly to RQ0. Individual-level characteristics overwhelmingly fail to differentiate CBDC design preferences within countries, while substantial cross-country differences persist across governance, security, and cost attributes. Of [462] within-country moderation tests, only [6.7%] reach even uncorrected significance, a rate barely above the 5% false-positive baseline expected by chance alone. This pattern supports the institutional-context account over individual-level theories and suggests that the forces shaping CBDC governance preferences operate at the level of shared ecosystem experience rather than individual psychology. The remainder of this section documents these patterns in detail.

First, it begins with cross-country comparisons (RQ5) to establish baseline differences, then examining within-country heterogeneity by individual characteristics (RQ1-RQ4). The analysis reveals a striking pattern: dramatic cross-country divergence coexists with near-universal within-country consensus.

4.1 RQ5: Payment Ecosystem Maturity Shapes Baseline Preferences (cross-country comparison)

Before examining cross-country differences and heterogeneity, respondents' overall choices are described to identify which CBDC design features matter most.

As observed in Figure 1, transaction cost exhibits the largest preference gradient in both countries.

In the United States, respondents select free CBDC designs 61.5% of the time versus only 36.3% for percentage-fee designs—a 25.2 percentage point range representing the single largest attribute effect observed. Indian respondents also prefer free designs (56.7%) over percentage fees (49.9%), though with a narrower 13.3 percentage point gradient. This finding aligns with payment economics predictions that cost is a primary driver of payment method choice.

Speed emerges as the second-strongest driver in both countries. Respondents in both contexts strongly prefer instantaneous transactions (USA: 57.7%, India: 55.1%) over delayed settlement (USA: 40.7%, India: 43.3%). The resulting 14–17 percentage point preference for speed is remarkably consistent across contexts, suggesting that transaction speed is a near-universal priority in digital payment systems.

Governance attributes matter substantially in the United States but not in India. American respondents exhibit large preference gradients for privacy (18.1 percentage points), security (17.1 percentage points), and government control (17.0 percentage points)—nearly as large as their sensitivity to transaction costs. Indian respondents, by contrast, display relatively flat preferences across governance attributes: privacy (3.2 percentage points), security (5.0 percentage points), and government control (2.4 percentage points). This divergence foreshadows the cross-country patterns examined in RQ5.

4.1.1 Directional Agreement, Intensity Differences

Despite divergence in intensity, respondents in both countries prefer the same attribute levels directionally. Both samples select CBDC profiles with free transactions, instantaneous speed, high security, full privacy, minimal government control, worldwide usability, and integration with existing applications more than 50% of the time. The cross-country disagreement is therefore not about what constitutes a desirable CBDC design, but about how much these features matter when users are forced to make trade-offs.

This shared baseline is an important context for the analyses that follow. American and Indian respondents agree that free, fast, secure, private, and autonomous CBDCs are preferable, yet they diverge substantially in the intensity with which governance features shape their choices.

4.1.2 Main Cross-Country Result

American and Indian respondents exhibit fundamentally different preference intensities across governance, security, and functional attributes. Of 21 attribute levels tested, 13 show statistically significant differences at the Bonferroni-corrected threshold of $\alpha = 0.002$, with effect sizes ranging from 3 to 13 percentage points. These differences align strongly with contextual predictions: India's UPI-saturated ecosystem produces acceptance of government oversight, while the United States' fragmented payment system produces stronger demand for privacy protection and autonomy.

Three clear patterns emerge from the cross-country comparison.

Pattern 1: USA demands privacy and autonomy; India accepts oversight. American respondents prefer full privacy (L1) by 7.8 percentage points more than Indian respondents (59.8% vs. 52.0%, $p < .001$) and reject full visibility (L3) by 7.2 percentage points more (41.7% vs. 48.9%, $p < .001$). Similarly, Americans prefer designs in which the government cannot block transactions (L1) by 6.6

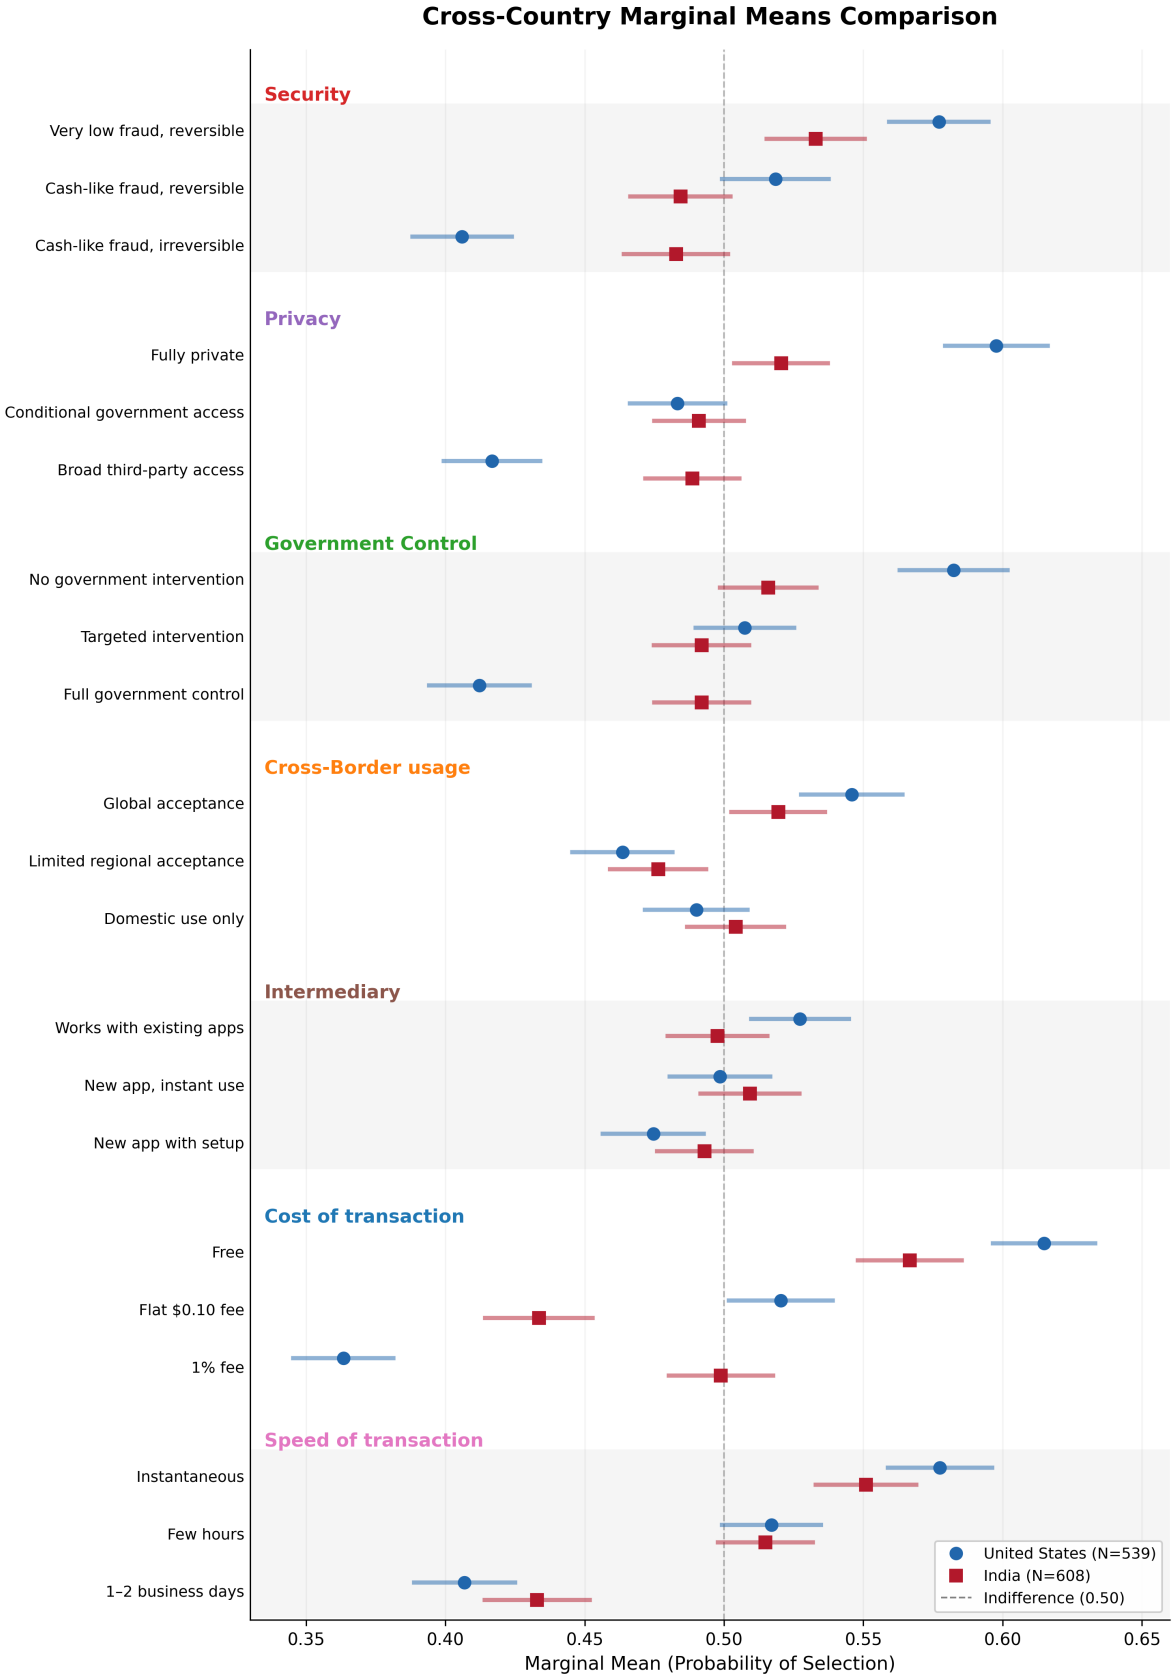

Figure 1: Marginal means by country for each CBDC attribute level. The baseline is marked with a dashed reference line at 0.50.

percentage points (58.2% vs. 51.6%, $p < .001$) and reject maximal government control (L3) by 8.0 percentage points more than Indians (41.2% vs. 49.2%, $p < .001$). These patterns align with contextual predictions: India's UPI operates with embedded government oversight, KYC requirements, and data-sharing arrangements, yet achieves 78.5% adoption as the primary payment method. Indian users appear adapted to surveillance-intensive digital infrastructure.

Pattern 2: Americans are more cost-sensitive; Indians shows slight acceptance of fee-based models. The single largest cross-country effect in the entire study involves transaction cost. American respondents strongly prefer free transactions (L1) and sharply reject percentage-fee models (L3), exhibiting a 13.6 percentage point lower preference for L3 compared to Indian respondents (36.3% vs. 49.9%, $p < .001$). Indian respondents display greater willingness to accept fee-based CBDC designs. This may reflect different baseline expectations shaped by existing ecosystems: while UPI's zero-fee model anchors Indian expectations, Indians may view fees as acceptable if other features such as speed and reliability are guaranteed. By contrast, Americans accustomed to free consumer-facing digital payments (e.g., Venmo, Zelle) exhibit strong resistance to explicit transaction fees.

Pattern 3: Security preferences diverge on reversibility. American respondents prefer the most secure design featuring very low fraud risk and reversible transactions (L1) by 4.4 percentage points relative to Indian respondents (57.7% vs. 53.3%, $p < .001$). Indian respondents, in contrast, are 7.7 percentage points more accepting of cash-like irreversible designs (L3) than Americans (48.3% vs. 40.6%, $p < .001$). These differences reflect ecosystem calibration: Indian users, accustomed to UPI's irreversible once-confirmed transaction model, tolerate finality, whereas American users, socialized into credit card chargebacks and dispute resolution, demand recourse options.

These findings strongly support H5a (the divergence hypothesis): cross-country differences in payment ecosystem maturity generate systematically different CBDC preference structures. They reject H5b (the universal preferences hypothesis). There is no single "optimal" CBDC design; preferences are fundamentally ecosystem-dependent.

Figure 1 displays marginal means by country for all 21 attribute levels. Table 3 (in Appendix) reports Z-test statistics for cross-country differences.

4.2 Within-Country Heterogeneity: The Consensus Effect

Having established that CBDC preferences diverge sharply between the United States and India, this section examines whether individual-level variation explains preferences within countries. The finding is that preferences are shared within each payment ecosystem to a striking degree. Across [462] moderator-by-attribute-level tests, the observed significance rate is [6.7%], barely above the 5% rate expected by chance, and only [2] tests survive Bonferroni correction. This within-country consensus spans every theoretically motivated moderator tested: trust in central banks, trust in digital payments, fraud victimization, privacy confidence, and cryptocurrency usage. This pattern, termed here the consensus effect, indicates that CBDC governance preferences are shaped by shared ecosystem experience rather than differentiated by individual characteristics.

4.2.1 RQ1: Trust in Central Bank x Government Control Preferences

Government control preferences are uniformly held within each country regardless of trust level. In the United States, both high-trust and low-trust respondents prefer CBDC designs where government cannot block transactions (L1) and reject maximal government control (L3), with no statistically significant differences across any of the three attribute levels (all p > .25). In India, both high-trust and low-trust respondents exhibit near-flat preferences across all government control configurations, although the India result is substantially constrained by the severe ceiling effect in trust (only 5.9% of Indian respondents fall in the low-trust category, yielding just 72 profile-level observations).

This finding carries a meaningful distinction. Trust may well matter for whether individuals are willing to adopt CBDC at all, consistent with its well-documented role as a gateway variable in adoption research (Søilen & Benhayoun, 2021). But trust does not determine which governance design respondents prefer once they are evaluating specific configurations. Citizens who trust the central bank and citizens who distrust it arrive at the same design preferences within their institutional context, suggesting that trust in central banks as monetary authorities does not translate into acceptance of central banks as transaction surveillance authorities. These findings reject H1a and are consistent with H1b.

Table 3: RQ1 Results—Trust × Government Control

Country	Attribute Level	Low Trust Mean	High Trust Mean	Z	p-value
USA ($n_{\text{low}} = 819$, $n_{\text{high}} = 1{,}361$)	Cannot block (L1)	0.592	0.567	1.14	.252
	Block w/ law enforcement (L2)	0.502	0.518	−0.74	.461
	Can block any (L3)	0.404	0.417	−0.59	.555
India ($n_{\text{low}} = 72$, $n_{\text{high}} = 3{,}074$)	Cannot block (L1)	0.514	0.514	0.00	.997
	Block w/ law enforcement (L2)	0.485	0.493	−0.13	.895
	Can block any (L3)	0.500	0.493	0.11	.910

Note: Trust measured on a 5-point scale. Low = 1–2; High = 4–5. India exhibits a severe ceiling effect (67% rated trust at 5; only 5.9% fall in the Low trust category). Observation counts reflect profile-level choices rather than individual respondents.

4.2.2 RQ2: Security Preferences Unaffected by Trust or Fraud Experience

Security preferences are shared within countries regardless of trust in digital payment platforms or personal fraud victimization. In both the United States and India, fraud victims and non-victims converge on the same security preference ordering: reversible, fraud-protected designs (L1) are preferred, cash-like reversible designs (L2) are intermediate, and irreversible designs (L3) are penalized. The confidence intervals for fraud and non-fraud groups overlap substantially at every security level in both countries, indicating no statistically meaningful divergence.

This convergence suggests that security operates as a collectively held baseline expectation, a threshold requirement shaped by ecosystem norms, rather than as a dimension along which personal experience generates differentiated demand. The security economics prediction that fraud victims should weight protection more heavily assumes that personal experience recalibrates individual risk preferences. The present data instead suggest that the relevant calibration happens at the ecosystem level: Americans expect chargebacks and reversibility because credit card infrastructure has established that norm, while Indians tolerate transaction finality because UPI's irreversible-once-confirmed model has established a different one. Individual fraud experience does not override these collectively held standards. These findings reject both H2a and H2b.

Table 4: RQ2a Results—Trust in Digital Payments × Security (USA Only)

Attribute Level	Low Trust Mean	High Trust Mean	Z	*p*-value
Very low fraud (L1)	0.590	0.576	0.41	.684
Cash-like, reversible (L2)	0.523	0.520	0.06	.951
Cash-like, irreversible (L3)	0.388	0.407	−0.58	.565

Note: USA only ($n_{\text{low}} \approx 230$, $n_{\text{high}} \approx 2{,}100$). Trust in digital payments measured on a 5-point scale. Low = 1–2; High = 4–5. India RQ2a tests are severely underpowered due to an extreme ceiling effect (95% rated trust in digital payments as 4–5; India Low trust $n \approx 32$). Tests were conducted but are not interpretable.

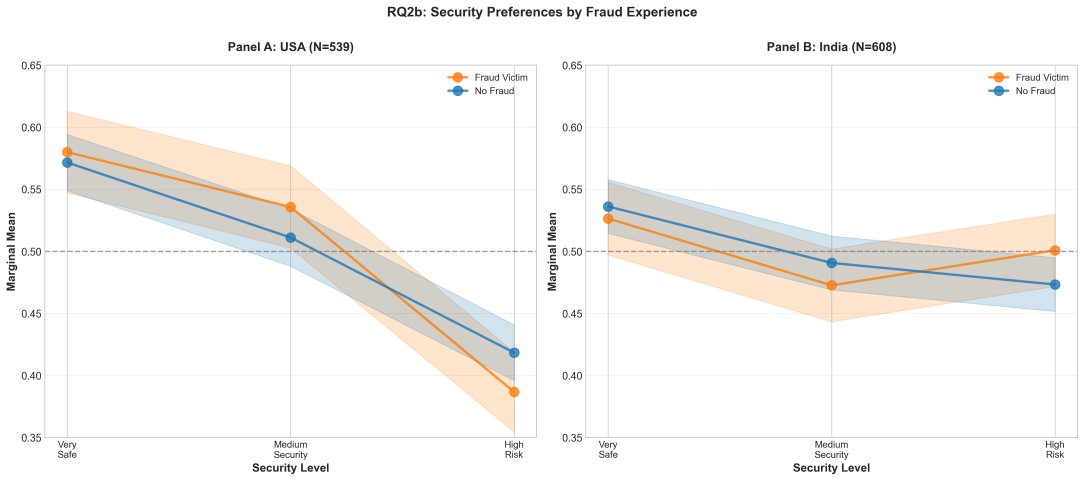

Figure 2: Marginal means for security attribute levels by fraud victimization status (RQ2b). The figure compares respondents who report having experienced digital payment fraud with those who have not, separately by country. Points indicate marginal means, defined as the proportion of times a given security level was chosen when presented. Error bars represent 95% confidence intervals. Both Americans and Indians converge on preferring reversible, fraud-protected designs (L1 > L2 > L3). Differences between fraud victims and non-victims are small and statistically insignificant.

Figure 2 visualizes the absence of heterogeneity in security preferences by fraud experience. In both the United States and India, respondents who report prior fraud victimization do not exhibit systematically stronger preferences for higher-security CBDC designs than those who have not experienced fraud. The confidence intervals for fraud and non-fraud groups overlap substantially at each security level, indicating no statistically meaningful divergence. This pattern reinforces the interpretation that security operates as a baseline requirement rather than a dimension along which individual experience generates differentiated preferences. Even respondents with direct exposure to fraud do not demand additional security beyond what is already perceived as adequate, suggesting that marginal improvements in security features do not translate into higher choice probabilities once a minimum threshold is met.

These findings reject H2a (trust reduces security sensitivity) and H2b (fraud increases security sensitivity).

4.2.3 RQ3: Privacy Confidence

In the United States, confidence in central bank privacy protection significantly moderates privacy attribute preferences, the only statistically significant within-country heterogeneity effect in the study. Low-confidence Americans, those who doubt the Federal Reserve will protect their financial privacy, prefer full privacy 4.9 percentage points more than high-confidence Americans (62.7% vs. 57.7%, p = .015). Both groups prefer full privacy over alternatives; the effect is a difference in the strength of that

Table 5: RQ3 Results — Privacy Confidence × Privacy Attributes

Country	Attribute Level	Low Confidence Mean	High Confidence Mean	Z	p-value
USA ($n_{\text{low}} = 1{,}047$, $n_{\text{high}} = 1{,}270$)	Full privacy (L1)	0.627	0.577	2.42	.015
	Conditional privacy (L2)	0.460	0.486	−1.23	.217
	Full visibility (L3)	0.410	0.436	−1.21	.227
India ($n_{\text{low}} = 52$, $n_{\text{high}} = 2{,}923$)	Full privacy (L1)	0.481	0.518	−0.53	.598
	Conditional privacy (L2)	0.544	0.495	0.74	.461
	Full visibility (L3)	0.471	0.488	−0.24	.807

Note: Confidence was measured for "How confident are you that the [Federal Reserve/RBI] will protect your financial privacy when using a CBDC?" on a 5-point scale. Low = 1–2; High = 4–5. The USA shows a statistically significant difference for the fully private design (L1), significant at $\alpha = .05$ and marginally surviving Holm–Bonferroni correction for three tests. India exhibits no significant moderation due to a severe ceiling effect (82% rated privacy confidence as 4–5; India Low confidence $n = 52$). Observation counts reflect profile-level choices rather than individual respondents.

preference rather than a reversal of direction.

This finding aligns with the distinction between privacy-by-design and privacy-by-promise (Cavoukian, 2011): low-confidence Americans demand architectural guarantees that do not depend on institutional behavior, while high-confidence Americans accept designs where privacy is contingent on institutional discretion. Three mechanisms are consistent with this pattern. First, low-confidence Americans may distrust that the Federal Reserve will honor privacy commitments, fearing data sharing with other agencies, policy changes, or mission creep, and therefore prefer full privacy (L1) because it provides structural protection independent of institutional behavior. Second, high-confidence Americans may believe the Fed will protect their privacy even under conditional designs (L2), making them willing to accept designs where privacy depends on institutional discretion. Third, and most consequentially for the broader argument, even the sole within-country exception to the consensus effect reflects an institutional dynamic rather than individual psychology: the relevant moderator is confidence in the institution's behavior, not a personal privacy attitude divorced from institutional context.

4.2.4 RQ4: Cryptocurrency Users Show No Governance Preference Differences

Cryptocurrency users exhibit the same ecosystem-shaped preference structure as non-users in both countries. American crypto holders share the baseline American demand for privacy and resistance to government control, while Indian crypto holders share the baseline Indian indifference across governance configurations. This pattern holds across all twelve tests (3 levels × 2 attributes × 2 countries), none of which reaches statistical significance.

This finding challenges the common assumption in CBDC policy debates that cryptocurrency adoption signals libertarian or anti-surveillance governance preferences. With crypto penetration at 51% in the United States and 68% in India, cryptocurrency participation has moved well beyond early cypherpunk or ideologically motivated communities, and many users likely engage with crypto for speculative investment, cross-border transfers, inflation hedging, or platform experimentation rather than as an expression of opposition to state money. Ecosystem-level norms dominate individual technological experience: crypto users in the United States exhibit baseline American preferences, and crypto users in India mirror baseline Indian preferences, because cryptocurrency exposure does not override the institutional socialization produced by national payment ecosystems. These results reject both H4a and H4b.

Additionally, respondents may conceptualize CBDC and cryptocurrency as serving *distinct and com-*

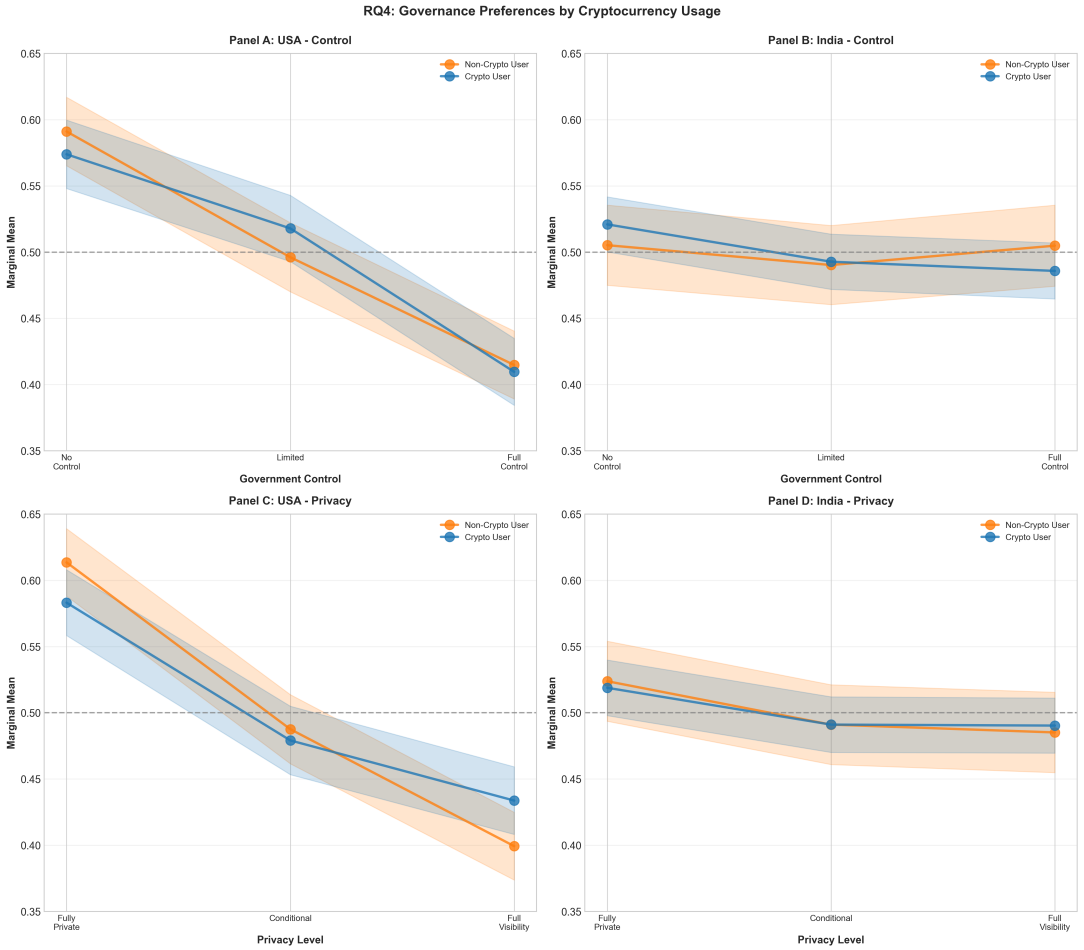

Figure 3: Cryptocurrency usage and CBDC governance preferences (RQ4). The figure displays marginal means for government control attributes (Panels A and B) and privacy attributes (Panels C and D), comparing respondents who report cryptocurrency ownership with those who do not, separately for the United States and India. Points indicate marginal means. Error bars represent 95% confidence intervals. Across both countries and both governance dimensions, preference patterns are substantively similar for crypto and non-crypto users, indicating no systematic moderation by cryptocurrency experience.

plementary functions rather than as substitutes. CBDC is framed as fiat-backed, reversible, and institutionally guaranteed, while cryptocurrency is risk-bearing, volatile, and often irreversible. From this perspective, crypto users may value CBDC precisely because it offers features that crypto lacks, such as legal recourse, consumer protection, and monetary stability. Consequently, they do not demand that CBDC replicate the governance or privacy model of Bitcoin or other decentralized systems.

Finally, crypto usage may already satisfy autonomy-seeking preferences outside the CBDC domain. Users who value decentralization, anonymity, or censorship resistance may view cryptocurrency itself as the appropriate vehicle for those goals, while evaluating CBDC on different criteria altogether. In this sense, CBDC governance preferences may reflect a form of domain separation: autonomy is sought in crypto, while stability, reliability, and usability are sought in state-backed digital money.

It would be unwise to treat cryptocurrency users as a distinct political or governance constituency in CBDC debates especially in regions with higher crypto usage. Crypto adoption alone does not predict resistance to state involvement, demand for anonymity, or preference for decentralized control in public digital money. For policymakers, this implies that targeting CBDC design toward crypto users as a pro- or anti-surveillance bloc is unlikely to be effective, which is one of the big policy narratives. Instead,

CBDC preferences appear to be shaped far more strongly by national payment infrastructure, institutional trust environments, and shared social expectations than by participation in alternative monetary systems.

5 Discussion

This study employed conjoint experiments across two institutionally distinct payment ecosystems, the United States ($N = 539$) and India ($N = 608$), to examine how users evaluate governance laden design features of retail central bank digital currencies. The findings challenge conventional assumptions about CBDC adoption, while offering insights into the relationship between institutional context, individual level variation, and design preferences. The consensus effect thus suggests a substantial cross-country divergence coexists with near-universal within-country agreement.

5.1 Core Findings

5.1.1 What Matters Overall

The central empirical contribution of this paper is documenting that CBDC governance preferences are collectively constructed within institutional contexts. Individual-level characteristics, including trust in central banks, trust in digital payments, fraud victimization, privacy confidence, and cryptocurrency usage, do not differentiate design preferences within countries because those preferences are shaped by shared ecosystem experience rather than personal attitudes. Of 36 within-country moderation tests, only one reaches statistical significance (2.8%), a rate that the individual-level theoretical frameworks dominating CBDC adoption research would not predict. The institutional trust, security economics, and privacy-autonomy literatures all expect meaningful heterogeneity by individual characteristics, yet the data show convergence instead.

Yet this within-country consensus coexists with substantial cross-country divergence. Of 21 attribute levels tested for cross-country differences, 13 show statistically significant differences at Bonferroni-corrected thresholds, with effect sizes ranging from 3 to 13 percentage points. Americans and Indians evaluate identical CBDC features differently—not because of individual psychology but because of ecosystem context.

The consensus effect has a natural interpretation through the lens of economic sociology. Finance is experienced socially rather than individually, and users do not evaluate CBDC governance features through private cost-benefit calculations but rather assess them relative to shared institutional baselines established through collective participation in existing payment infrastructure. The moral economy framework (Thompson, 1971) offers a useful lens here: within each payment ecosystem, there exists a shared normative understanding of what constitutes legitimate state involvement in financial transactions. In India, where government-branded digital infrastructure delivers tangible everyday benefits, government oversight falls within the boundaries of the moral economy of digital payments. In the United States, where federal involvement in financial transactions activates surveillance anxieties rooted in post-Snowden politics and contested federal authority, that same oversight falls outside those boundaries. These boundaries are collectively maintained, which is precisely why individual variation in trust or privacy attitudes fails to predict where individuals fall within their own country's consensus.

This interpretation carries a methodological implication. The individualist cost-benefit framing that dominates technology acceptance research, and that structures much of the CBDC adoption literature, may itself reflect a culturally specific assumption about how people relate to monetary infrastructure. The present findings suggest that collective institutional experience, rather than individual rational calculation, is the more relevant unit of analysis for understanding CBDC governance preferences.

5.1.2 Cross-Country Patterns

Two patterns characterize the cross-country divergence:

Privacy and autonomy versus acceptance of oversight. Americans prefer full privacy 7.8 percentage points more than Indians (59.8% vs. 52.0%) and reject full government visibility 7.2 percentage points more (41.7% vs. 48.9%). Americans prefer CBDC designs where government cannot block transactions by 6.6 percentage points (58.2% vs. 51.6%) and reject maximal government control by 8.0 percentage points more (41.2% vs. 49.2%). These differences align with contextual predictions: India's UPI-saturated ecosystem, with embedded KYC requirements and government data-sharing, has normalized oversight; the US fragmented system and surveillance debates produce demands for privacy and autonomy.

Security and reversibility. Americans prefer reversible, fraud-protected designs more strongly, while Indians tolerate irreversible transactions at higher rates (48.3% vs. 40.6% for cash-like irreversible). United States respondents exhibit steeper security gradients, while India respondents show flatter patterns. This reflects ecosystem calibration: UPI's irreversible-once-confirmed model has conditioned Indian expectations, while American credit card chargebacks might have established reversibility as baseline. Security is more differentiated in the United States than in India.

5.1.3 Institutional Privacy Distrust

The sole within-country heterogeneity effect emerges in US privacy preferences. Low-confidence Americans—those who doubt the Federal Reserve will protect their financial privacy—prefer full privacy 4.9 percentage points more than high-confidence Americans (62.7% vs. 57.7%, $p = .015$). This finding is intuitive: users who distrust the central bank's commitment to privacy protection demand architecturally private systems that do not depend on institutional trustworthiness. They seek privacy-by-design rather than privacy-by-promise.

This finding has significant implications and reinforces the paper's broader theme. Even the one significant within-country effect is about institutional trust—specifically, distrust of the central bank translating into demand for structural protections. Users who doubt the Fed will honor privacy commitments want CBDC designs where privacy is guaranteed by architecture (full privacy, L1) rather than contingent on institutional behavior (conditional privacy, L2).

5.1.4 Trust as a Contextual Moderator

Government control is uniformly rejected in the United States and comparatively non differentiating in India. This pattern transcends measured individual differences. Even high trust respondents in the United States prefer minimal government oversight, while low trust respondents in India do not display

sharp rejection of state involvement.

Trust does not significantly moderate preferences for government control or security. High trust and low trust respondents show similar gradients, jointly preferring minimal government control and stronger security. The main exception is privacy. Respondents with low confidence that the central bank would protect their privacy show stronger preferences for fully private CBDC designs ($p = 0.031$), suggesting that privacy skepticism translates into differentiated feature preferences.

The India sample exhibits a strong ceiling effect on trust measures, reducing power to detect trust based heterogeneity. More substantively, the absence of interaction effects is consistent with trust becoming background rather than foreground, that is, a taken for granted feature of the payment ecosystem rather than an actively evaluated dimension during design trade offs.

5.2 Theoretical Implications

5.2.1 For CBDC Research

The findings challenge several assumptions common in CBDC adoption research.

First, privacy concerns are not a universal barrier to CBDC adoption. Prior work often treats privacy as a broadly binding constraint (Agur et al., 2022; Tronnier et al., 2023). The present results suggest privacy salience varies across contexts. Privacy related attitudes shape feature preferences in the United States but not in India, implying that privacy preserving design may be essential in some jurisdictions and less central in others.

Second, CBDC acceptance appears endogenous to institutional embedding rather than solely technical features. The strong within country consensus across moderators suggests users evaluate CBDC through existing institutional relationships rather than finely individualized preference functions. In India, success of state led digital public goods may supply legitimacy that transfers to CBDC. In the United States, the absence of analogous public payment infrastructure may limit the ability of a federal CBDC to inherit legitimacy.

Third, the meaning of CBDC attributes is not fixed but constructed through context. “Government control” is evaluated differently in a developmental state context where digital infrastructure has delivered tangible benefits than in a liberal market context where government surveillance has distinct ideological valence. Design features therefore cannot be evaluated in abstraction from institutional arrangements.

A further consideration is that constructs such as "trust" and "privacy" may carry different connotative weight across institutional contexts. "Government access to transaction data" is evaluated against a backdrop of UPI-embedded KYC requirements and Aadhaar-linked identity verification in India, and against a backdrop of post-Snowden surveillance debates, IRS enforcement controversies, and contested federal authority in the United States. While the conjoint design mitigates this concern by presenting concrete design features rather than abstract attitudinal prompts, the institutional meaning of those features is inevitably context-dependent, which is precisely the point the contextual ecosystem perspective makes.

5.2.2 For Security Economics

The findings contribute to understanding how users reason about security in digital payment contexts.

Security preferences are not individually calibrated. The prediction that fraud victims should weight security more heavily (H2b) assumes that fraud experience creates lasting preference shifts. the null results suggest either that fraud victims already select into high-security payment methods (a selection effect), or that security operates as a threshold requirement rather than a continuous preference—everyone wants "enough" security, but marginal improvements beyond that threshold do not differentiate choices.

Security is a non-negotiable baseline. Both Americans and Indians converge on preferring reversible, fraud-protected designs regardless of individual risk exposure. This cross-country consensus on security suggests that security is perceived as a governance baseline rather than a feature to be traded against other attributes.

5.2.3 For Privacy Research

The institutional privacy distrust finding offers a nuanced account of privacy preferences in digital infrastructure contexts.

Users who distrust the central bank's commitment to privacy protection demand architecturally private systems. This is not privacy pessimism about personal capabilities but rather rational skepticism about institutional behavior—a demand for structural guarantees that do not depend on institutional trustworthiness. High-confidence users, by contrast, are willing to accept conditional privacy designs because they trust the institution to exercise discretion responsibly.

This has implications beyond CBDC. In any context where privacy depends on institutional behavior, data retention policies, access controls, sharing agreements, users who distrust the institution will demand privacy-by-design rather than privacy-by-policy. Stated commitments and privacy policies are insufficient for institutionally distrustful users; only architectural guarantees (encryption, data minimization, no central visibility) will suffice. This suggests that privacy credibility requires demonstrated commitment through structural design, not merely announced intentions.

The distinction between privacy-by-design and privacy-by-promise gains additional specificity from recent technical work mapping the landscape of privacy-enhancing technologies for digital payments (Auer et al., 2025). Zero-knowledge proofs, anonymous credentials, and other cryptographic tools offer structural privacy guarantees that do not depend on institutional discretion, providing precisely the kind of architectural protection that institutionally distrustful users in this study demand. The empirical finding and the technical possibility converge: privacy-by-design is not merely a theoretical aspiration but a technically feasible design choice that addresses a documented demand among users who lack confidence in institutional commitments to data protection.

5.3 Policy Implications

The pattern of results suggests that CBDC is evaluated as an infrastructure extension, where ecosystem experience sets baseline expectations for what digital money should deliver, rather than as a policy proposition, where ideological commitments about privacy, surveillance, or government authority would

drive choices. If CBDC were evaluated as a policy proposition, attitudinal moderators such as trust in government, privacy attitudes, and crypto ideology should dominate within-country heterogeneity. Instead, the few within-country effects that emerge reflect operational expectations calibrated by prior ecosystem experience, and the most reliably moderated attribute is transaction speed, a functional performance dimension, rather than privacy or government control. This has meaningful learnings that can be drawn for the policy environment.

5.3.1 For Central Banks Generally

Ecosystem context precedes feature design. The within-country consensus suggests that CBDC acceptance depends less on matching designs to diverse individual preferences and more on whether the ecosystem context permits CBDC to inherit legitimacy from existing institutions. India's digital public goods success (UPI, Aadhaar) creates a foundation that may transfer to the digital rupee; contexts lacking such foundations may find CBDC adoption challenging regardless of technical sophistication.

Privacy credibility requires demonstrated commitment. Privacy features are valuable only if users believe them. The institutional privacy distrust finding suggests that central banks cannot simply announce privacy protections—they must establish credibility through structural design (architecturally private systems) or institutional commitments (legislative backing, independent oversight) rather than policy statements alone.

Security is table stakes. The convergence on security preferences across countries and subgroups indicates that robust security is a necessary but not sufficient condition for adoption. Central banks cannot trade off security for other features; security must be assured before other design considerations become relevant.

5.3.2 For the United States

The uniform rejection of government control across all trust levels—including among high-trust respondents—suggests that CBDC skepticism reflects structural concerns about federal financial surveillance rather than addressable trust deficits. Any US CBDC design would need credible structural commitments (congressional authorization, independent oversight, strict data minimization) to address governance concerns that cross attitudinal lines.

The institutional privacy distrust finding suggests that privacy-preserving design is especially important for the US context, where low-confidence populations translate their concerns into differentiated feature preferences. Surveillance-intensive architectures would likely face adoption barriers among users who distrust the Federal Reserve's commitment to privacy protection.

5.3.3 For India

UPI's success creates both opportunity and constraint. The opportunity lies in high baseline confidence in government-backed digital payment infrastructure and a normalized acceptance of oversight. The constraint is extremely high expectations: any CBDC that is slower, more expensive, or less convenient than UPI will struggle to achieve adoption.

The strong demand for dispute resolution suggests that Indian users expect consumer protection mecha-

nisms that match or exceed existing UPI standards. As a result, CBDC designs in India should prioritize functional improvements—such as enhanced dispute resolution, offline usability, or programmability over governance differentiation, which appears comparatively less salient in the Indian context.

5.4 Limitations

It is important to acknowledge several limitations that should be kept in mind when reading the findings.

5.4.1 Sample Size and Statistical Power

The sample sizes across both countries are modest for detecting interaction effects. While adequate for estimating main effects and cross-country differences, the within-country moderation tests may be underpowered to detect small-to-medium interaction effects. The predominance of null results in the heterogeneity analyses could reflect true absence of moderation or insufficient statistical power. Future research with larger samples could determine whether the consensus effect reflects genuine preference homogeneity or Type II error.

The India sample's severe distributional skew compounds this concern. With 67% of Indian respondents expressing maximum trust in central banks and only 5.9% in the low-trust category, the "Low trust" India subgroup contains only 72 observations—far below conventional power thresholds for detecting medium-sized effects. The apparent within-country consensus in India may partially reflect measurement limitations rather than true preference homogeneity.

5.4.2 Sample Composition

Both samples were recruited through online platforms and the survey was conducted in English. In the Indian context, where hundreds of millions of citizens transact in regional languages and engage with digital payments through vernacular interfaces, the English-language restriction limits generalizability to urban, educated, English-proficient populations. Geographic distribution within countries is not reported, and regional variation, between digitally saturated southern Indian states and less-connected northern states, or between coastal and interior US populations, may introduce preference heterogeneity that is not captured by the national-level analysis. The high rates of cryptocurrency familiarity in both samples (51% US, 68% India) further suggest that these respondents are more financially engaged than general populations, which should be considered when interpreting the scope of the findings.

5.4.3 Hypothetical Choices and External Validity

Conjoint experiments measure stated preferences over hypothetical designs, not actual adoption decisions. When real money, real privacy risks, and real government oversight are at stake, users might reason differently. The extensive literature on intention and behavior gaps in technology adoption counsels caution in translating stated preferences directly into adoption predictions.

Furthermore, respondents evaluated CBDC profiles without experiencing actual system performance. Trust, security, and privacy attributes were described textually; actual CBDC deployment would involve experienced performance that might shift preferences. The relatively high marginal means for preferred attribute levels (often 0.55–0.60) suggest respondents engaged meaningfully with choice tasks, but we

cannot assess how preferences would evolve with experience.

Additionally, the design forced choices between two CBDC profiles without an opt-out option. Respondents who would reject both profiles in favor of cash or existing digital payments could not express this preference. This design choice maximizes information about relative preferences but may overstate absolute CBDC acceptance. In real deployment contexts, competition with existing payment methods may matter more than the survey design captures.

5.4.4 Cross-Sectional Design

The data capture preferences at a single point in time (November 2025). CBDC attitudes are evolving as public awareness increases, pilot programs generate track records, and political debates reshape information environments. The preference structures documented here may shift with changing contexts. Longitudinal designs tracking preference evolution would strengthen causal interpretation.

Finally, while the US-India comparison reveals meaningful cross-national variation, two countries cannot identify which specific ecosystem features drive differences. Payment infrastructure maturity, government digital services track record, political culture around surveillance, and many other factors differ simultaneously. Additional countries with intermediate profiles, the EU, China, smaller economies, would help isolate mechanisms and test generalizability.

5.5 Future Research

Mechanism tests. While the findings document that individual characteristics largely fail to moderate CBDC design preferences, we cannot definitively explain why this consensus effect arises. Experimental manipulations of information environments, institutional framing, or prior belief activation could test competing mechanisms, such as social normalization, expectation anchoring to existing payment systems, or cognitive simplification in complex governance decisions.

Disaggregating trust and privacy. More refined measurement of trust components, distinguishing competence, benevolence, and integrity, and of privacy concerns, separating surveillance, secondary data use, and breach risk, may reveal heterogeneity that the coarser measures obscure. Multi-dimensional scales and latent-class approaches could uncover subgroups that standard moderation tests fail to detect.

Privacy distrust interventions. The institutional privacy distrust finding raises questions about what would satisfy skeptical users. Testing whether technical measures (zero-knowledge proofs), institutional measures (independent oversight, legislative guarantees), or structural measures (decentralized architecture) differentially affect institutionally distrustful populations would inform design strategy.

6 Conclusion

This study examined how users in two institutionally distinct contexts, the United States and India, evaluate governance-laden features of retail CBDC design. The central finding is the consensus effect: within each payment ecosystem, users converge on shared preference structures shaped by institutional context rather than individual psychology. This convergence holds across every tested moderator, including trust in central banks, fraud experience, privacy attitudes, and cryptocurrency usage, and is

documented through [462] within-country moderation tests where the significance rate barely exceeds the chance false-positive baseline. Yet this within-country homogeneity coexists with substantial cross-country divergence, as 13 of 21 attribute levels differ significantly between the United States and India after Bonferroni correction, with effect sizes reaching 8 to 14 percentage points on governance and cost attributes.

The consensus effect documented here finds independent support in experimental CBDC research from other institutional contexts as well. These converging results across methodologically similar studies strengthen the conclusion that CBDC design preferences are shaped primarily by ecosystem context rather than by individual characteristics.

For central banks, the implication is that CBDC acceptance depends less on accommodating diverse individual preferences, which turn out to be surprisingly homogeneous within countries, and more on whether the institutional context permits CBDC to inherit legitimacy from existing infrastructure. In India, where government-led digital payment infrastructure has achieved near-universal trust and adoption, governance features are accepted by default and the design challenge is functional performance. In the United States, where government involvement in financial transactions is politically contested, no feature configuration overcomes governance skepticism. The problem is not design but institutional context.

The single exception to within-country is around Americans who doubt the Federal Reserve will protect their financial privacy demand architecturally private systems that do not depend on institutional trustworthiness. This finding implies that privacy-by-design is not merely a preference to accommodate, but a structural necessity for building credibility among institutionally skeptical populations.

This research states that CBDC is not merely a technical innovation but a governance-laden infrastructure that inherits and activates institutional relationships. The individual-level frameworks that dominate CBDC adoption research predict preference heterogeneity that does not exist in these data. Understanding the institutional relationships that do shape preferences, and the payment ecosystems that sustain them, may be the most important contribution research can offer as central banks worldwide navigate the uncertain terrain of digital currency design.

References

Agur, I., Anil, R., & Dell'Ariccia, G. (2022). Designing central bank digital currencies. *Journal of Monetary Economics*, 125, 62–79.

Anderson, R. (2001). Why information security is hard: An economic perspective. In *Proceedings of the 17th Annual Computer Security Applications Conference*.

Balasubramanian, S. A., & Thirumaran, P. (2025). Understanding enablers and inhibitors of digital rupee: A dual-factor theory perspective. *Digital Policy, Regulation and Governance*, 27(5), 507–522.

Bansak, K., Hainmueller, J., Hopkins, D. J., & Yamamoto, T. (2021). Beyond the breaking point? Survey satisficing in conjoint experiments. *Political Science Research and Methods*, 9(1), 53–71.

Bhatnagr, P. (2025). Enhancing digital currency adoption: Examining user experiences. *Management Decision*, 63(7), 2292–2316.

Bijlsma, M., van der Cruijsen, C., Jonker, N., & Reijerink, J. (2021). What triggers consumer adoption of CBDC? *DNB Working Paper No. 709*.

Desai, A. (2024). Analysing consumer switching behaviour and financial antecedents to determine the intention to use central bank digital currency: Evidence from the Indian economy. *Digital Policy, Regulation and Governance*.

Dixit, V., Shailesh, A., Mishra, S., & Verma, R. (2025). Adoption of central bank digital currency in India: A structural model using ISM. *NMIMS Management Review*, 33(4), 276–288.

Fairweather, M., et al. (2024). Discrete choice experiment on CBDC vs. deposits: Safety and privacy attributes. *Reserve Bank of Australia Discussion Paper*.

GENIUS Act. (2025). Guiding and Establishing National Innovation for U.S. Stablecoins Act. https://www.congress.gov/bill/119th-congress/senate-bill/1582/text.

Gupta, R., et al. (2024). Does previous experience with the Unified Payments Interface (UPI) affect the usage of central bank digital currency (CBDC)? *Digital Finance*.

Hainmueller, J., Hopkins, D. J., & Yamamoto, T. (2014). Causal inference in conjoint analysis: Understanding multidimensional choices via stated preference experiments. *Political Analysis*, 22(1), 1–30.

Horiuchi, Y., Markovich, Z., & Yamamoto, T. (2022). Does conjoint analysis mitigate social desirability bias? *Political Analysis*, 30(4), 535–549.

Jiang, J. (2023). Privacy implications of central bank digital currencies. *Seton Hall Law Review*, 54, 69.

Kahn, C. M., & Roberds, W. (2009). Why pay? An introduction to payments economics. *Journal of Financial Intermediation*, 18(1), 1–23.

Kaur, P., et al. (2024). Exploring factors affecting central bank digital currency adoption: A perspective from Generation Z. *International Journal of Bank Marketing*.

Koziuk, V. and Ivashuk, Y. (2022). Does it matter for CBDC design? Privacy–anonymity preferences from the side of hierarchies and egalitarian cultural patterns. *Economics – Innovative and Economics Research Journal*, 10(1).

Koziuk, V. (2021). Confidence in digital money: Are central banks more trusted than age is matter? *Investment Management and Financial Innovations*, 18(1), 12–32.

Koziuk, V. (2020). Confidence to digital currencies of central banks: Institutional paradox or age matters. *World of Finance*, 2(63), 08–23.

Krause, D. (2025). A roundtable discussion on the Trump administration's policy on private stablecoins over a U.S. CBDC. SSRN Working Paper. https://ssrn.com/abstract=5134818.

Krishnamoorthy, B., & Aggarwal, V. (2024). Digital rupee for retail adoption and challenges. In *Proceedings of the 10th International Conference on Smart Computing and Communication (ICSCC)* (pp. 171–175). IEEE.

Leeper, T. J., Hobolt, S. B., & Tilley, J. (2020). Measuring subgroup preferences in conjoint experiments. *Political Analysis*, 28(2), 207–221.

Levi, M., & Stoker, L. (2000). Political trust and trustworthiness. *Annual Review of Political Science*, 3, 475–507.

Liu, Z., et al. (2024). Determinants of individuals' intentions to use central bank digital currency: Evidence from China. *Technological Forecasting and Social Change*.

Mehlkop, G., et al. (2023). Cross-national survey experiment on privacy and actor identity in CBDC-like payments. *Journal of Behavioral Finance*.

Nissenbaum, H. (2009). *Privacy in Context: Technology, Policy, and the Integrity of Social Life*. Stanford University Press.

Ogunmola, A., & Das, S. (2024). Technology acceptance model for digital rupee adoption in India. *Journal of Financial Services Research*.

Søilen, K. S., & Benhayoun, L. (2021). Household acceptance of central bank digital currency: The role of institutional trust. *Technology Analysis & Strategic Management*, 34(9).

Solove, D. J. (2006). A taxonomy of privacy. *University of Pennsylvania Law Review*, 154(3), 477–564.

Tinn, K., & Dubach, C. (2021). Central bank digital currency with asymmetric privacy. *Available at SSRN 3787088*.

Tronnier, F., Biker, P., Baur, E., and Löbner, S. (2023). An evaluation of information flows in digital euro transactions using contextual integrity theory. In *Proceedings of the 2023 European Interdisciplinary Cybersecurity Conference*, 7–12.

Tronnier, F., Harborth, D., and Biker, P. (2023). Applying the extended attitude formation theory to central bank digital currencies. *Electronic Markets*, 33(1), 13.

Tronnier, F. and Qiu, W. (2024). How do privacy concerns impact actual adoption of central bank digital currency? An investigation using the e-CNY in China. *Quantitative Finance and Economics*, 8(1), 126–152.

Koziuk, V., Dziubanovska, N., and Tsegelnyy, I. (2025). Can artificial intelligence help to identify the privacy paradox: The case of CBDC.

Executive Order 14178. (2025). Strengthening American Leadership in Digital Financial Technology. https://www.presidency.ucsb.edu/documents/executive-order-14178-strengthening-american-leadership-digital-financial-technology.

Granovetter, M. (1985). Economic action and social structure: The problem of embeddedness. *American Journal of Sociology*, 91(3), 481–510.

Thompson, E.P. (1971). The moral economy of the English crowd in the eighteenth century. *Past & Present*, 50, 76–136.

Silverstone, R. and Hirsch, E. (1992). *Consuming Technologies: Media and Information in Domestic Spaces*. Routledge, London.

Hyysalo, S. (2010). *Health Technology Development and Use: From Practice-Bound Imagination to Evolving Impacts*. Routledge, London.

Acquisti, A., Brandimarte, L., and Loewenstein, G. (2015). Privacy and human behavior in the age of information. *Science*, 347(6221), 509–514.

Abramova, S., Böhme, R., Elsinger, H., Stix, H., and Summer, M. (2022). What can CBDC designers learn from asking potential users? Results from a survey of Austrian residents. *Oesterreichische Nationalbank Working Paper No. 241*.

Elsinger, H., Stix, H., and Summer, M. (2025). Consumer preferences for a digital euro: Insights from a discrete choice experiment in Austria. *BIS Working Papers No. 1302*, Bank for International Settlements.

Cavoukian, A. (2011). *Privacy by Design: The 7 Foundational Principles*. Information and Privacy Commissioner of Ontario, Canada.

Auer, R., Böhme, R., Clark, J., and Demirag, D. (2025). Privacy-enhancing technologies for digital payments: Mapping the landscape. *BIS Working Papers No. 1242*, Bank for International Settlements.

Plato-Shinar, R. and Maman, L. (2025). Public attitudes towards CBDC and the role of trust in the central bank. *Bar Ilan University Faculty of Law Research Paper*, forthcoming.

Choi, S., Kim, B., Kim, Y.S., Kwon, O., and Park, S. (2026). Predicting the payment preference for CBDC: A discrete choice experiment. *Bank of Korea Working Paper No. 26-01*.

A Appendix: Survey Instrument

A.1 Survey Questions and tasks

This appendix reproduces the conjoint choice task screen as presented to respondents in each survey. Respondents in both countries evaluated 16 randomly generated paired CBDC profiles and selected the profile they would prefer to use based on the following prompt: *Out of the two options on the page, please select the option you would most likely use for everyday transactions.*

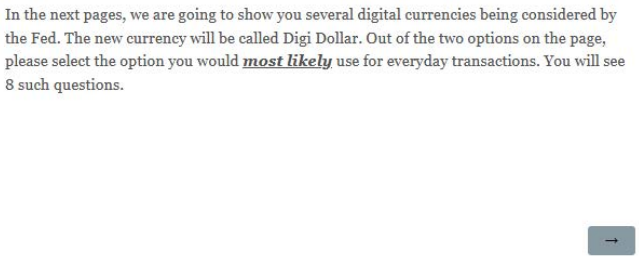

Figure 4: Selection Prompt

A.1.1 Conjoint Task Examples

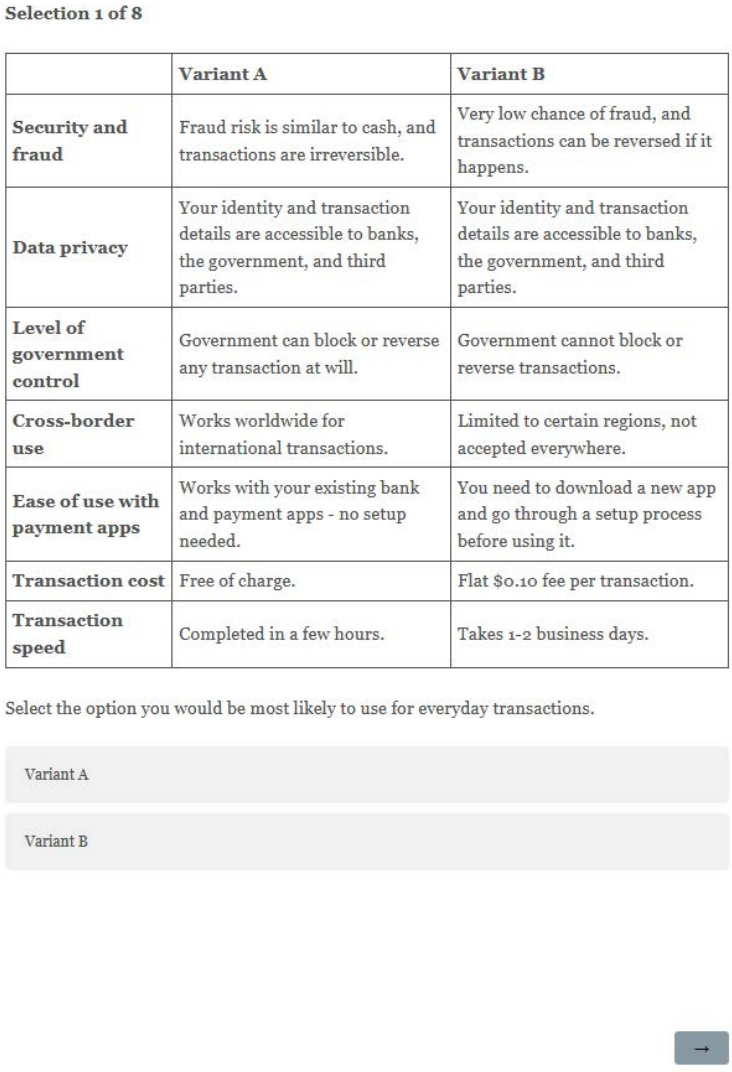

Selection 1 of 8

	Variant A	Variant B
Security and fraud	Fraud risk is similar to cash, and transactions are irreversible.	Very low chance of fraud, and transactions can be reversed if it happens.
Data privacy	Your identity and transaction details are accessible to banks, the government, and third parties.	Your identity and transaction details are accessible to banks, the government, and third parties.
Level of government control	Government can block or reverse any transaction at will.	Government cannot block or reverse transactions.
Cross-border use	Works worldwide for international transactions.	Limited to certain regions, not accepted everywhere.
Ease of use with payment apps	Works with your existing bank and payment apps - no setup needed.	You need to download a new app and go through a setup process before using it.
Transaction cost	Free of charge.	Flat $0.10 fee per transaction.
Transaction speed	Completed in a few hours.	Takes 1-2 business days.

Select the option you would be most likely to use for everyday transactions.

Variant A

Variant B

→

Figure 5: Conjoint Task 1

Note: Each respondent in India and the US completed 16 forced-choice tasks. Attribute levels were randomized independently across profiles and tasks, with no restrictions on level combinations. The Conjointly platform's data export displays "Federal Reserve" in certain question text fields (Q22–Q24); however, respondents saw correctly localized "Reserve Bank of India" text during the actual survey administration. This is a documented platform export artifact, not a survey administration error.

A.1.2 Post-Conjoint Questions

1. **Digital payment frequency:** Frequency of digital payment method use such as mobile payment apps and online banking (US: Q16; India: Q8). Response options: Daily, Weekly, Monthly, Rarely, Never.
2. **Technology comfort:** Comfort adopting new technology such as mobile apps, digital wallets, and new forms of currency (US: Q18; India: Q9). Five-point scale from "Extremely uncomfortable"

to "Extremely comfortable."

3. **Cash vs. digital usage:** Relative frequency of cash versus digital/card payments (US: Q19; India: Q10). Five ordered options from "I mostly only use cash" to "I mostly use digital payments."
4. **Cross-border payment behavior:** International purchase frequency, remittance behavior, and international travel frequency (US: Q20–Q22; India: Q11–Q13).
5. **Cryptocurrency experience:** Binary—whether respondent has invested in, traded, or used cryptocurrency such as Bitcoin or Ethereum (US: Q23; India: Q14).
6. **Cryptocurrency ban support:** Whether the respondent supports, opposes, or has no opinion on a ban on cryptocurrency (US: Q25; India: Q15).
7. **Trust in digital payments:** General trust in the security of digital payment systems and methods (US: Q29; India: Q17). Five-point scale from "I do not trust them at all" to "I trust them completely."
8. **Fraud victimization:** Whether the respondent has been a victim of online financial fraud (US: Q32; India: Q18). Three options: Yes, No, Not Sure.
9. **CBDC adoption intent:** Likelihood of adopting CBDC if government offered incentives such as cashback or discounts (US: Q33; India: Q19). Five-point Likert scale from "Very unlikely" to "Very likely."
10. **Trust in central bank (financial system):** Trust in the Federal Reserve/RBI to manage the country's financial system and monetary policies (US: Q36; India: Q22). Five-point scale from "I distrust it completely" to "I trust it completely."
11. **Trust in central bank (digital currencies):** Trust in the Federal Reserve/RBI to securely manage digital currencies (US: Q37; India: Q23). Five-point scale.
12. **Privacy confidence:** Confidence that the Federal Reserve/RBI will protect financial privacy when using a CBDC (US: Q38; India: Q24). Five-point scale from "Not confident at all" to "Very confident."
13. **CBDC concerns:** Multi-select of specific concerns about CBDC (US: Q42; India: Q25).
14. **Demographics:** Age, gender, education, employment status, and annual household income (US: Q47–Q51; India: Q27–Q31).

A.2 Demographic Comparison

This appendix presents the demographic composition for the US and India. The figures below summarize the distribution of respondents across five demographic dimensions.

Age

Bars represent the percentage of respondents within each age bracket. The India sample skews substantially younger: 66.2% of Indian respondents are under 35, compared to 28.7% in the United States. Conversely, 37.8% of US respondents are 55 or older, versus just 2.9% in India. This age gap reflects

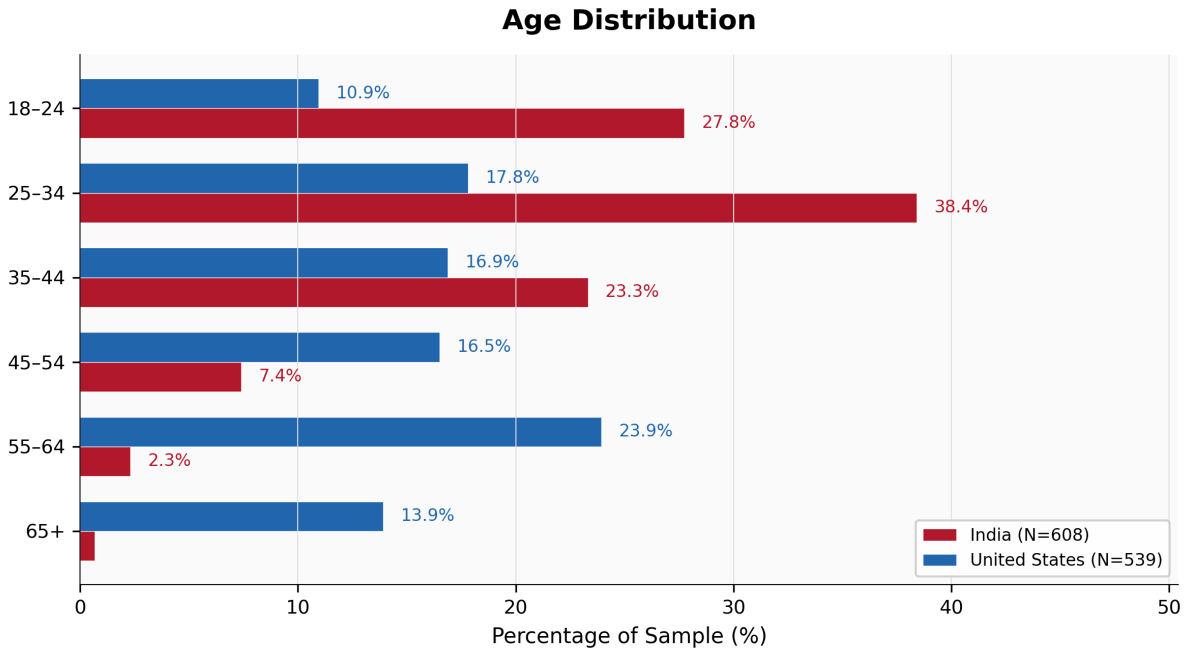

Figure 6: Age Distribution

both actual population demographics — India's median age is approximately 28 years compared to 38 in the United States — and the recruitment characteristics of each platform. Conjointly's India panel draws heavily from digitally active younger adults, while Prolific's US panel achieves broader age representation. Critically, despite this substantial age asymmetry, the moderator analysis in Section **??** finds no evidence that age-correlated individual characteristics (such as technology comfort or digital payment frequency) differentiate CBDC design preferences within either country, consistent with the infrastructure experience framework's prediction that ecosystem context dominates individual demographic variation.

A.2.1 Education

Education categories were harmonized across the two survey instruments. Both samples are well-educated relative to their national populations: approximately 40% hold bachelor's degrees in each country. However, the India sample has a substantially higher share of postgraduate degree holders (44.1% vs. 20.0% in the US), while the US has a larger "some college / associate degree" segment (26.0% vs. 3.1% in India). This pattern reflects both India's emphasis on postgraduate credentialing in urban professional populations and the broader distribution of US tertiary education pathways. The high education levels in both samples suggest respondents are well-positioned to evaluate the technical and governance trade-offs presented in the conjoint task, though they may not be representative of the broader populations that would ultimately use a retail CBDC.

A.2.2 Employment

The India sample reports substantially higher full-time employment (69.6% vs. 49.9%). The US sample shows greater dispersion across self-employment (14.3%), unemployment (10.0%), and other categories including students, homemakers, and retirees (13.2%). The higher full-time employment rate in India

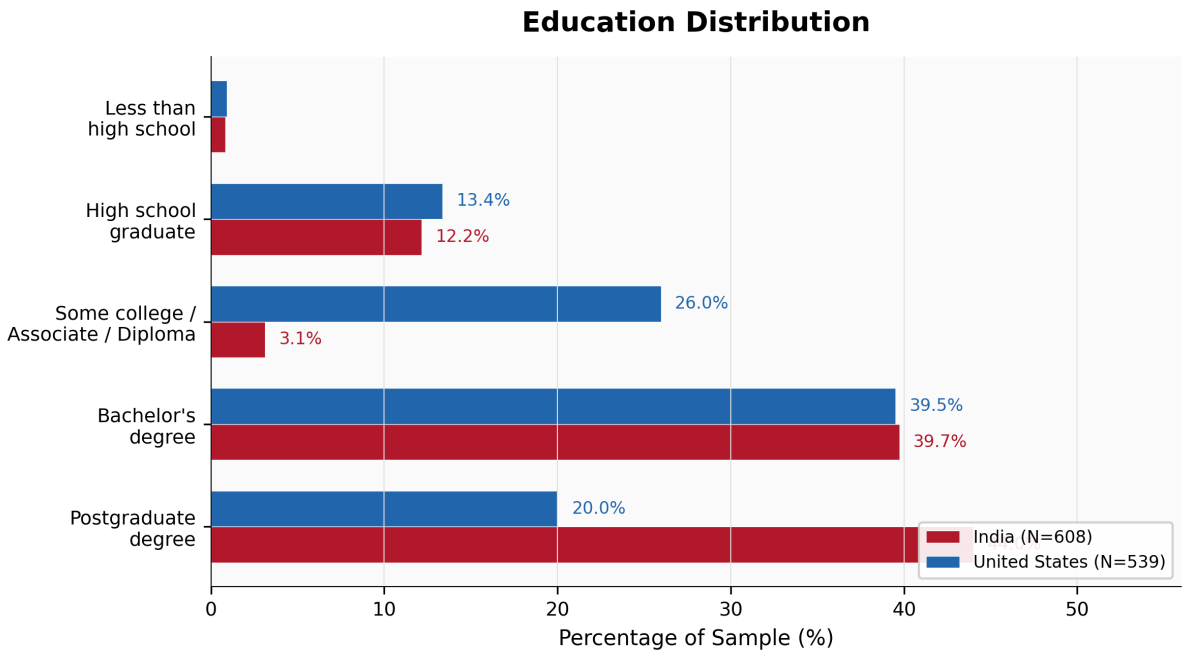

Figure 7: Educational profile

is consistent with recruiting digitally literate urban professionals — the population segment most likely to have direct experience with UPI and other digital payment infrastructure. This experiential base is precisely what the infrastructure experience framework identifies as the mechanism shaping CBDC evaluation: Indian respondents assess CBDC relative to a mature, government-facilitated digital payment ecosystem they use daily.

A.2.3 Gender

The US sample is approximately balanced (47.3% male, 51.2% female, 1.5% other/non-binary). The India sample skews male (59.2% male, 40.8% female), with no non-binary category available through the Conjointly platform. The male overrepresentation in India is consistent with documented gender gaps in Indian digital payment adoption and online survey panel composition. Gender is not included as a moderator variable in the present analysis, as the theoretical framework focuses on technology engagement, institutional trust, and financial experience rather than demographic identity per se.

A.2.4 Income

Income is presented in local currency (USD and INR respectively) because purchasing power parity conversion would obscure the domestic economic context within which respondents evaluate financial products. The US distribution is roughly centered on the $60,000–$100,000 range, with a median bracket of $60,001–$80,000 (15.4%). The India distribution shows concentration at the lower bracket (less than 800,000 26.0%) with a substantial high-income tail (greater than 7,000,000 15.6%). Both distributions span the full income range, providing adequate variation for detecting income-related preference heterogeneity. Seven Indian respondents (1.2%) selected "Prefer not to say" and are excluded from the income figure; percentages are calculated over the remaining 601 respondents.

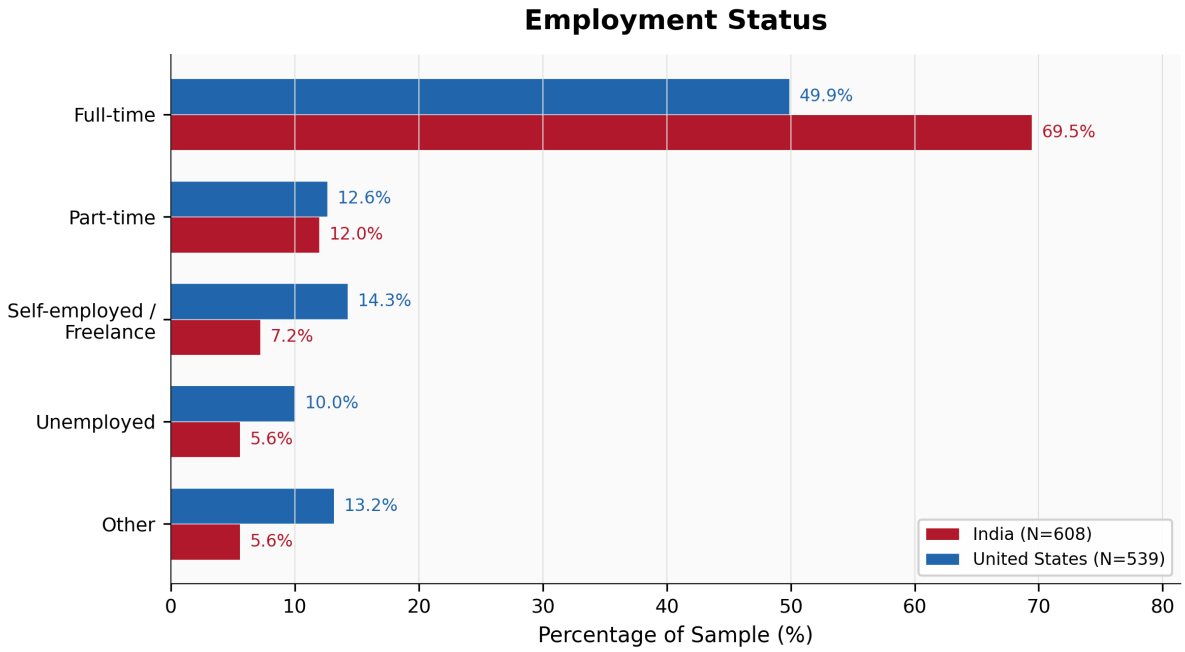

Figure 8: Employment Status

A.2.5 Ethical Considerations

The study received IRB approval from *Georgia Institute of Technology*.[1] All respondents provided informed consent before participating. The survey included no embedded attention checks beyond the consent screen and completion time threshold; quality filtering relied on these two criteria. The conjoint experiment presents hypothetical CBDC designs and involves no deception, sensitive personal disclosures beyond standard demographic and attitudinal questions, or vulnerable populations. Respondents were compensated through the panel providers' standard compensation structures (Qualtrics for US, Conjointly for India). Data are stored and analyzed in de-identified form in accordance with institutional data management protocols.

A.3 Why marginal means in the main results and AMCEs in the appendix?

Marginal means are presented in the main results and report Average Marginal Component Effects (AMCEs) in the appendix for four reasons.

- Marginal means are more intuitive for policy audiences, as they can be interpreted as simple choice proportions (e.g., "X% of profiles with full privacy were chosen").
- Marginal means show absolute preference levels, whereas AMCEs capture relative changes from a baseline category.
- Under random assignment of attribute levels, both marginal means and AMCEs convey the same causal information; however, marginal means are pedagogically clearer for interpreting governance preferences.
- AMCE results are provided in the appendix for readers familiar with standard conjoint estimation

[1] *IRB protocol number: IRB2025-600; approved on August 3, 2025.*

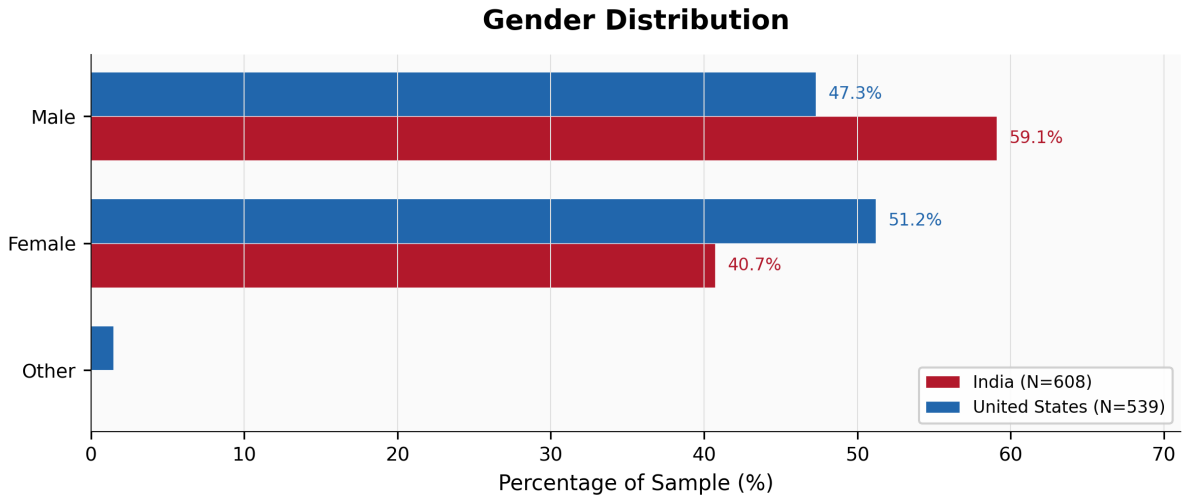

Figure 9: Gender Distribution

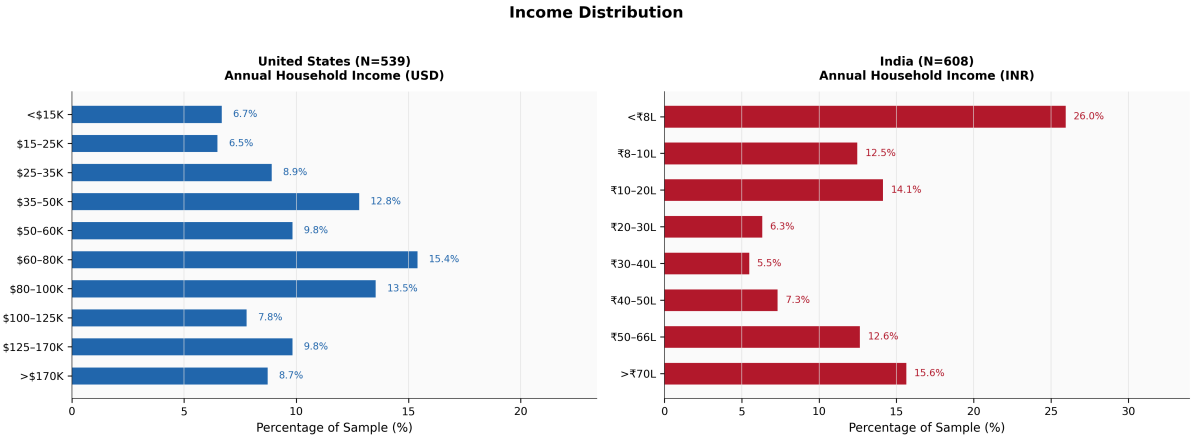

Figure 10: Income Distribution

and for comparability with the existing literature.

Inference. Standard errors for marginal means are calculated via bootstrap (1,000 iterations) clustered at the respondent level to account for repeated observations per individual.

A.3.1 Supplementary Analysis: Average Marginal Component Effects (AMCE)

Average Marginal Component Effects (AMCEs) is also reported, which are estimated using linear probability models regressing the binary choice outcome on indicator variables for each attribute level:

$$\text{Choice}_{ijt} = \alpha + \sum_{k=1}^{7} \sum_{\ell \neq \text{baseline}} \beta_{k\ell} \cdot 1(\text{Attribute}_{k,ijt} = \ell) + \varepsilon_{ijt} \tag{3}$$

where i indexes respondents, j indexes profiles within choice tasks, t indexes choice tasks, and standard errors are clustered at the respondent level.

All attributes are dummy-coded, with the middle level (L2) serving as the omitted reference category. This coding choice allows direct comparison between more permissive and more restrictive policy de-

signs relative to an intermediate baseline.

AMCE interpretation. $\hat{\beta}_{k\ell}$ represents the change in probability that a profile is chosen when attribute k is set to level ℓ instead of the baseline level L2, holding all other attributes constant.

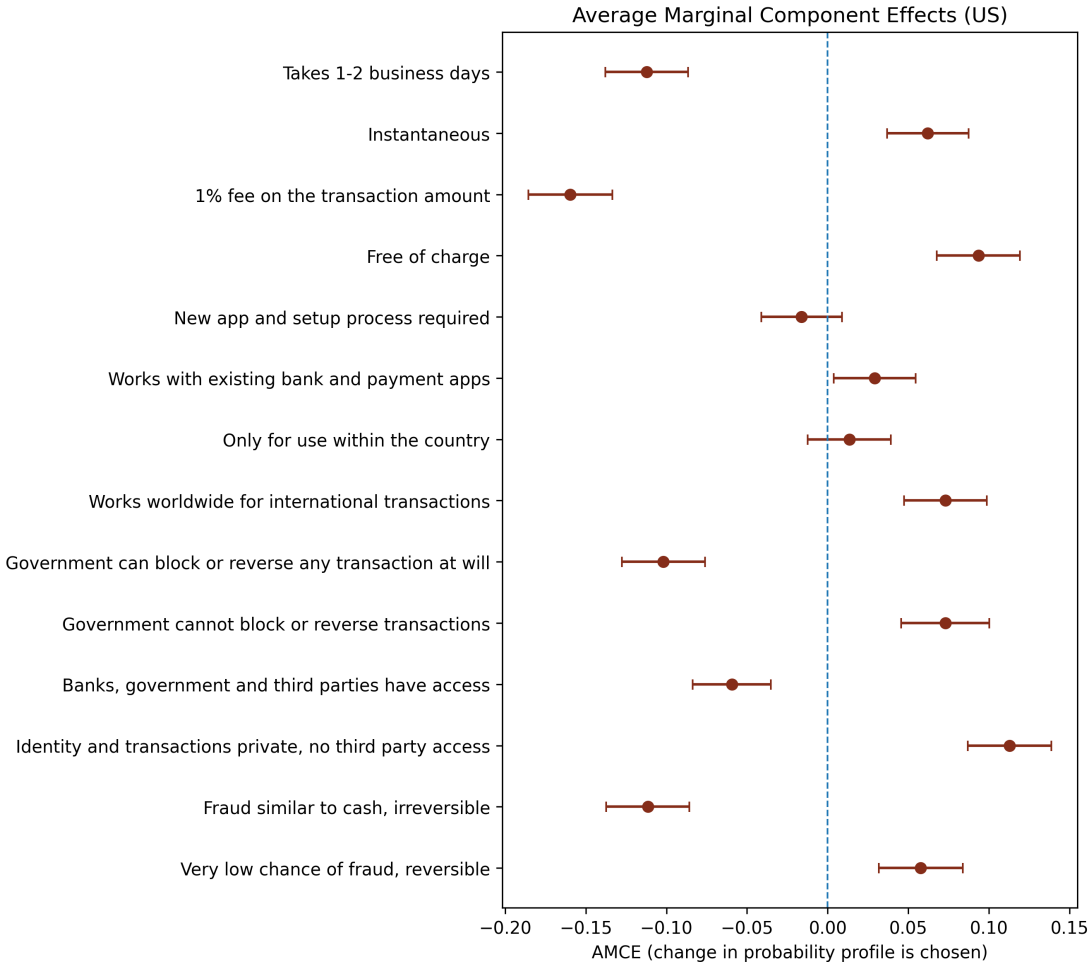

Figure 11: AMCE plot for US with L2 as baseline

A.3.2 Robustness Checks

Several robustness checks are conducted:

- Alternative model specifications (logistic regression, inclusion of demographic controls)
- Sensitivity to subgroup definitions (tertile splits, continuous interaction terms)
- Sample restrictions by response quality (completion time)

IRB details intentionally omitted for double-blind review (to be added later).

.

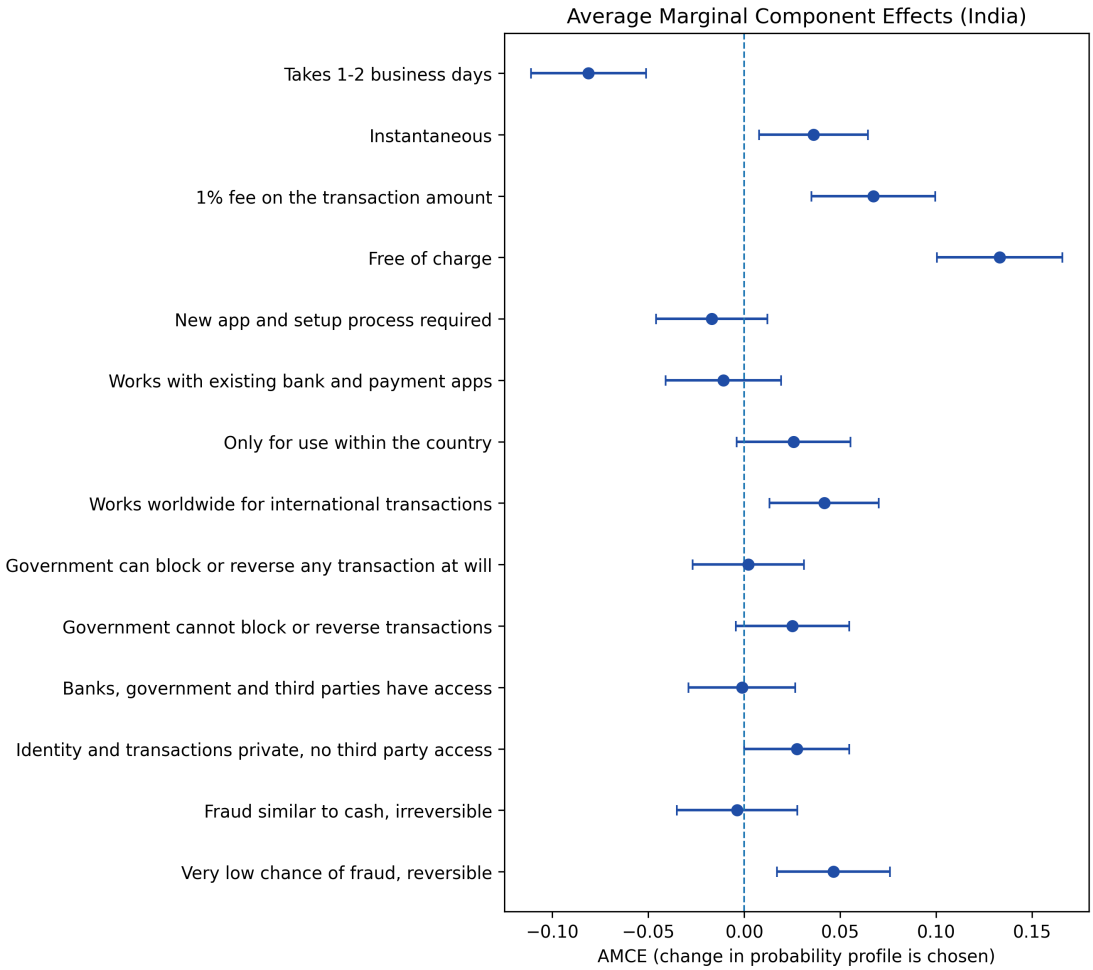

Figure 12: AMCE plot for India with L2 as baseline

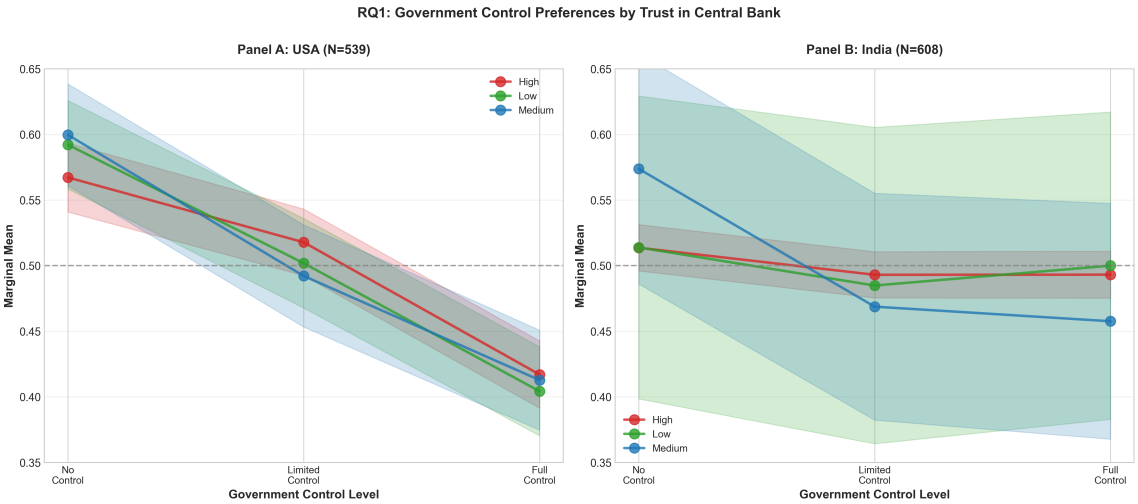

Figure 13: Marginal means for government control preferences by trust subgroup. Lines represent Low, Medium, and High trust groups within each country. Error bars represent 95% confidence intervals.

Table 6: Cross-Country Differences in CBDC Preferences

Attribute	Level	USA Mean	India Mean	Z-stat	p-value
Governance Attributes					
Privacy (A2)	Full privacy (L1)	0.598	0.520	6.08	$< .001$
Privacy (A2)	Conditional (L2)	0.483	0.491	-0.65	.518
Privacy (A2)	Full visibility (L3)	0.417	0.489	-5.84	$< .001$
Government Control (A3)	Cannot block (L1)	0.582	0.516	5.01	$< .001$
Government Control (A3)	Block w/ law enforcement (L2)	0.507	0.492	1.23	.217
Government Control (A3)	Can block any (L3)	0.412	0.492	-6.30	$< .001$
Security (A1)	Very low fraud, reversible (L1)	0.577	0.533	3.46	$< .001$
Security (A1)	Cash-like, reversible (L2)	0.518	0.484	2.54	.011
Security (A1)	Cash-like, irreversible (L3)	0.406	0.483	-5.83	$< .001$
Functional Attributes					
Cost (A6)	Free (L1)	0.615	0.567	3.61	$< .001$
Cost (A6)	Low fixed fee (L2)	0.520	0.434	6.34	$< .001$
Cost (A6)	Percentage fee (L3)	0.363	0.499	-10.24	$< .001$
Cross-border (A4)	Worldwide (L1)	0.546	0.519	2.09	.037
Cross-border (A4)	Limited regions (L2)	0.464	0.476	-1.01	.314
Cross-border (A4)	Domestic only (L3)	0.490	0.504	-1.09	.277
Integration (A5)	Works with existing apps (L1)	0.527	0.498	2.31	.021
Integration (A5)	New app, immediate use (L2)	0.499	0.509	-0.83	.405
Integration (A5)	New app, setup required (L3)	0.475	0.493	-1.45	.148
Speed (A7)	Instantaneous (L1)	0.577	0.551	2.01	.045
Speed (A7)	Few hours (L2)	0.517	0.515	0.17	.865
Speed (A7)	1–2 business days (L3)	0.407	0.433	-1.96	.050

Note: Entries report marginal means (proportion chosen) by country and two-sample Z-tests for differences in proportions. Bonferroni-corrected significance threshold is $\alpha = 0.05/21 \approx 0.002$. Uncorrected p-values are reported.

Figure 14: Marginal means for security preferences by trust in digital payments subgroup. Points indicate marginal means. Error bars represent 95% confidence intervals.

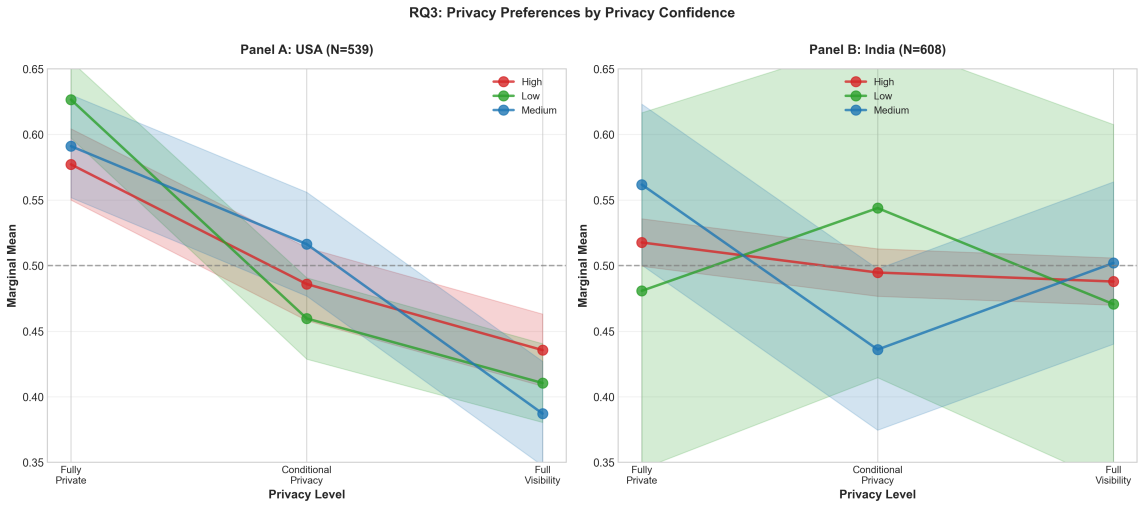

Figure 15: Marginal means for privacy attribute preferences by privacy-confidence subgroup. Points indicate marginal means. Error bars represent 95% confidence intervals.

Table 7: RQ4a Results — Cryptocurrency Use × Government Control

Country	Attribute Level	Non-Crypto Mean	Crypto User Mean	Z	p-value
USA ($n_{\text{non}} = 1{,}389$–$1{,}420$, $n_{\text{crypto}} = 1{,}408$–$1{,}537$)	Cannot block (L1)	0.591	0.574	−0.92	.356
	Block w/ law enforcement (L2)	0.496	0.518	1.18	.237
	Can block any (L3)	0.415	0.410	−0.28	.778
India ($n_{\text{non}} = 1{,}026$–$1{,}077$, $n_{\text{crypto}} = 2{,}149$–$2{,}219$)	Cannot block (L1)	0.505	0.521	0.84	.401
	Block w/ law enforcement (L2)	0.490	0.493	0.13	.893
	Can block any (L3)	0.505	0.486	−1.01	.315

Note: Cryptocurrency usage defined as self-reported ownership or use of any cryptocurrency. Observation counts reflect profile-level choices rather than individual respondents.

Table 8: RQ4b Results — Cryptocurrency Use × Privacy Attributes

Country	Attribute Level	Non-Crypto Mean	Crypto User Mean	Z	p-value
USA ($n_{\text{non}} = 1{,}405$–$1{,}410$, $n_{\text{crypto}} = 1{,}436$–$1{,}516$)	Full privacy (L1)	0.614	0.583	−1.68	.094
	Conditional privacy (L2)	0.488	0.479	−0.45	.651
	Full visibility (L3)	0.399	0.434	1.87	.062
India ($n_{\text{non}} = 1{,}043$–$1{,}059$, $n_{\text{crypto}} = 2{,}176$–$2{,}221$)	Full privacy (L1)	0.524	0.519	−0.26	.791
	Conditional privacy (L2)	0.491	0.491	0.00	.999
	Full visibility (L3)	0.485	0.490	0.28	.782

Note: No statistically significant differences are observed across any attribute level in either country. Marginal effects that approach conventional significance in the USA do not survive multiple-testing correction. Observation counts reflect profile-level choices rather than individual respondents.