Delhi Metropolitan Education
Affiliated to GGSIP University, New Delhi & Approved by Bar Council of India

RESEARCH ARTICLE
doi: 10.53361/dmejc.v7i01.05

Trust, Privacy, and Transparency in Multi-Service Platforms: A Study of India's Super Apps

Prakriti Tripathi*, Muktesh Rohilla, Ramvati

Master of Design (M.Des), NIFT Panchkula, India

ARTICLE INFO

***Correspondence:**
prakriti.tripathi@nift.ac.in
Master of Design (M.Des), NIFT Panchkula, India

Dates:
Received: 18-04-2026
Accepted: 21-05-2026
Published: 30-06-2026

Keywords:
Super Apps, Trust, Privacy, Transparency, Trust Spillover, AI, DPDP, REFLECT Framework, Dark Patterns, User Experience, India

How to Cite:
Tripathi, P., Rohilla, M., Ramvati. (2026) Trust, Privacy, and Transparency in Multi-Service Platforms: A Study of India's Super Apps. MediaSpace: DME Journal of Communication,7(1),54-63. doi: 10.53361/dmejc.v7i01.05

Abstract

India's super apps are growing fast. Platforms like PhonePe, Paytm, Tata Neu, and Jio have layered payments, shopping, travel, and entertainment into a single application, building directly on the scale of UPI adoption. As AI features chatbots, personalised recommendations, fraud detection get embedded into these platforms, questions of trust, privacy, and transparency become harder to sidestep.

This study looks at how Indian users aged 18–45 perceive trust, privacy, and transparency within AI-enabled super apps and specifically, whether trust in one service bleeds over to others on the same platform. Using a structured questionnaire (N = 50) and the REFLECT Ethical UX Framework, four hypotheses were tested. Three findings stand out: a moderate-to-strong trust spillover effect across all AI feature pairs (Pearson r = 0.533–0.742, all $p < 0.001$); a privacy paradox where AI convenience (M = 3.82) and data concern (M = 3.90) co-exist; and a clear user demand for transparency, with 80.0% of respondents saying they would share more data if given honest explanations and genuine control. The study also finds serious failures in the Legibility and Consent dimensions of current super app interfaces, and proposes three design interventions aligned with India's Digital Personal Data Protection (DPDP) Act 2023 and the 13 dark patterns prohibited by the Central Consumer Protection Authority (CCPA).

Introduction

The Rise of Super Apps in India

India's digital ecosystem has shifted quickly. UPI crossed 18 billion monthly transactions in 2024, building a large base of habitual digital payment users (NPCI, 2024). That payment habit is the foundation on which super apps now sit. A super app functions like a digital operating system payments, shopping, travel, food, entertainment all inside one application (Hasselwander, 2024).

India is still working out what its super app model looks like. PhonePe and Paytm are payments-first. Tata Neu pulls together retail, groceries, travel, and financial products. Jio sits across telecom, e-commerce, and entertainment. Amazon India merges shopping, Prime Video, and Amazon Pay under one login. None of them are as tightly consolidated as WeChat or Grab, but the convergence

MediaSpace:
DME Journal of Communication
e-ISSN: 2583-035X

trajectory is clear (Hasselwander & Weiss, 2025). The draw is simple: one login, one app, many services. Consumer research shows users gravitate toward super apps when the experience feels seamless particularly in finance and e-commerce and that they will adopt new services within a trusted platform as long as they see compatibility and competence (Fang et al., 2024).

Why Trust Becomes the Central Problem

Bundling services means bundling data. A single super app session can touch location data, purchase history, browsing behaviour, biometric credentials, and financial records at the same time. AI raises the stakes further: chatbots require conversation histories, recommendation engines run on behavioural logs, fraud detection depends on transaction patterns. The more an AI does, the more it needs and the more trust becomes the platform's most exposed variable (Cai et al., 2025).

Research on AI trust points to three consistent factors: personal experience with the system, perceived reliability, and the context in which it operates (Ng & Zhang, 2025). Trust erodes when users cannot tell how their data is being used, when AI decisions feel opaque, or when one bad experience in a connected service sours their view of the whole platform. Prior work has mostly studied trust within single platforms. What is missing is quantitative evidence on how trust moves across multiple bundled services inside a super app that is the gap this study addresses.

The 2026 Regulatory and Security Context

Early 2026 brought a clear regulatory shift. The Digital Personal Data Protection (DPDP) Rules 2026, issued under the DPDP Act 2023, require platforms to replace blanket consent permissions with clear, specific, purpose-limited notices. Data Fiduciaries must inform users what data is collected, why it is collected, and how consent can be withdrawn at any time.

In the same period, NPCI released security upgrades for UPI in April 2026 mandatory two-factor biometric authentication, AI-powered anomaly detection on transactions, and stronger device-binding protocols a direct response to rising UPI fraud in 2025. This study was conducted at that regulatory and technological inflection point, when the benefits and the risks of AI-powered financial systems were both visible to ordinary users.

Literature Review

The Growth of Super App Ecosystems

Super apps follow a recognisable growth logic: a company starts with a high-frequency service payments or messaging then layers additional services on top to improve retention and capture more data (Hasselwander, 2024). A multidisciplinary review of the field places super apps in their own analytical category, one that requires frameworks drawing on platform economics, security assessment, and user behaviour together (Hasselwander et al., 2025).

Structurally, super apps resemble mobile operating systems: a host application runs mini-apps from third parties, all sharing a central data layer (Yang et al., 2023). The architecture is efficient for users but creates real security exposure data leakage between mini-apps, API abuse, insufficient service isolation failures that can hit millions of users at once. On the demand side, e-commerce, banking, and social media are consistently the services users most want bundled, with payments serving as the gateway through which they first enter the broader platform (Hasselwander & Weiss, 2025).

AI Integration and Its Effect on User Trust

AI is no longer an add-on in super apps it is baked into how these platforms work. Personalised recommendations, predictive analytics, chatbots, automated fraud detection, smart search: these are baseline expectations for users of PhonePe, Paytm, and Amazon India, not premium extras. Studies show that AI does improve engagement, but only when the system is seen as reliable (Bayashot, 2025; Chhetri et al., 2025).

The factors that shape AI trust are relatively well-established. Bach et al. (2022) identify four pillars transparency, explainability, fairness, and perceived reliability across human-computer interaction research. In the Indian context specifically, Ahire (2025) finds that transparency and data privacy safeguards are the strongest trust predictors for generative AI, and argues these cannot be assessed outside sociocultural frameworks. Antonio and Kurniawan (2025) take this further, showing that trust mediates the link between AI capability and long-term super app use a finding with direct commercial relevance.

Privacy, the Privacy Paradox, and Structural Risk

AI personalisation and privacy exposure come as a pair. Cai et al. (2025) show that super apps can infer sensitive attributes health conditions, financial stress, personal relationships purely from interaction histories across mini-apps, without users ever disclosing that information directly. The platform's centralised architecture makes this worse: a vulnerability in one mini-app or API can propagate across the entire system (Yang et al., 2023).

The privacy paradox where users report concern about data but continue using data-intensive platforms anyway (Norberg et al., 2007) sits at the heart of this study. Users do not disengage when worried; they keep using while staying quietly concerned. This study examines that behaviour empirically. Cvetkovic et al. (2026), drawing on data from twelve countries, find that trust in major tech companies is a significant predictor of trust in their AI systems, with national regulatory context acting as a moderator a dynamic directly relevant to India's evolving DPDP framework.

Research Gap

Research on super apps, AI features, and user trust has mostly developed in parallel rather than in conversation. The missing piece is quantitative evidence on how AI integration shapes trust across multiple bundled services, specifically in India. No published study has tested whether trust in AI features of one service on a platform influences trust in other services on the same platform what this study calls the 'trust spillover effect.' The four hypothesis tests here are designed to fill that gap.

The REFLECT Ethical UX Framework

To evaluate how super app interfaces perform against ethical design standards, this study applies the REFLECT Ethical UX Framework as an analytical lens. REFLECT is an author-applied framework developed for this study, built around seven dimensions: Respect (protecting user autonomy), Ethical Design (avoiding manipulative patterns), Fairness (equitable AI behaviour), Legibility (comprehensible disclosures), Empathy (acknowledging user discomfort), Consent (informed and voluntary data agreement), and Transparency (visible AI logic and data flows). The framework was chosen because it connects regulatory requirements what platforms are obligated to do with actual design practice how users encounter those requirements on screen. Applied to survey data, it lets each interface failure be mapped to a specific ethical dimension rather than left as a vague complaint about complexity or privacy.

Research Question and Objectives

Research Question

How does AI integration within super apps influence user trust and privacy perceptions among Indian users?

The study pursues three objectives

- To review privacy and security concerns related to AI integration in super apps.
- To examine whether trust in one service influences trust in another within a super app.
- To examine the effect of data transparency on user trust in AI-enabled features within super apps.

Methodology

Research Design

The study used a quantitative cross-sectional design. A structured questionnaire collected data from Indian users aged 18 to 45 who regularly use super apps or multi-service platforms. It covered demographics, app usage patterns, frequency of AI feature use, and five-point Likert scale ratings for trust, convenience, privacy concern, transparency demand, and multi-service comfort (1 = strongly disagree; 5 = strongly agree). Open-ended responses were collected alongside and used to give texture to the quantitative patterns. Convenience sampling was used appropriate for exploratory work in contexts where probability sampling is not operationally realistic.

Sample and Data Cleaning

Eighty-four people responded to the survey. Three exclusion criteria were applied: unfamiliarity with super apps, no use of any multi-service platform, and age outside the 18–45 range. After cleaning, the final sample comprised N = 50 active super app users. The 40.5% exclusion rate (n = 34) is itself informative: India's super app category is still forming, and many people use individual platforms within it without recognising them under that label. Among those retained, 98.0% were aged 18–34, matching the demographic profile of India's active digital payment and e-commerce user base.

Analytical Approach

Quantitative data were analysed using descriptive statistics (mean, frequency, percentage) and Pearson's correlation coefficient. Each hypothesis was tested with the method best suited to its variable structure: one-way ANOVA for grouped comparisons (H1, H4), independent samples t-test for two-group comparisons (H3), and pairwise Pearson correlation for continuous variable relationships (H2, H3). Cronbach's alpha for the four-item privacy subscale was $\alpha = 0.521$, indicating that the items tap distinct dimensions of privacy concern rather than a single underlying construct. This justified using the single-item personal data concern measure as the independent variable for H1. Open-ended responses were reviewed thematically to contextualise the numbers and illustrate how users articulate their own behaviour.

Results

Platform Usage and AI Feature Adoption

Amazon was the most used platform at 66.0% (n = 33), followed by Paytm at 62.0% (n = 31), PhonePe at 48.0% (n = 24), MyJio at 36.0% (n = 18), and Tata Neu at 18.0% (n = 9). Amazon is included on functional grounds: it integrates shopping, Prime Video, and Amazon Pay within a single login, meeting Hasselwander's (2024) definition of a super app. Service use ranked in this order: payments, shopping, travel and bookings, utility services, entertainment, and food delivery confirming that transactional interactions form the trust base on which users assess all other bundled services.

AI feature use was high: 80.0% of respondents (n = 40) use AI features regularly within their super apps, 14.0% (n = 7) use them occasionally, and 10.0% (n = 5) not at all. The most-used features were personalised recommendations, chatbots and virtual assistants, smart search, image recognition, and voice search. Mean trust scores were similar across all four AI features tested Personalised Recommendations (M = 3.54, SD = 1.28), Chatbots (M = 3.54, SD = 1.30), Smart Notifications (M = 3.48, SD = 1.42), and Fraud Detection (M = 3.40, SD = 1.31) pointing to a uniformly moderate level of trust rather than any single feature standing out. Table 1 presents the demographic profile of the final sample (N = 50). The sample skewed young, with 48.0% aged 18–24 and 50.0% aged 25–34, and was fairly balanced by gender (56.0% male, 40.0% female).

Table 1: Demographic profile of the final sample (N = 50).

Variable	*Category*	*n*	*%*
Age group	18–24 years	24	48.0%
	25–34 years	25	50.0%
	35–44 years	1	2.0%
Gender	Male	28	56.0%
	Female	20	40.0%
	Prefer not to say	2	4.0%
Occupation	Student	25	50.0%
	Working Professional	20	40.0%
	Self-Employed / Business	7	14.0%

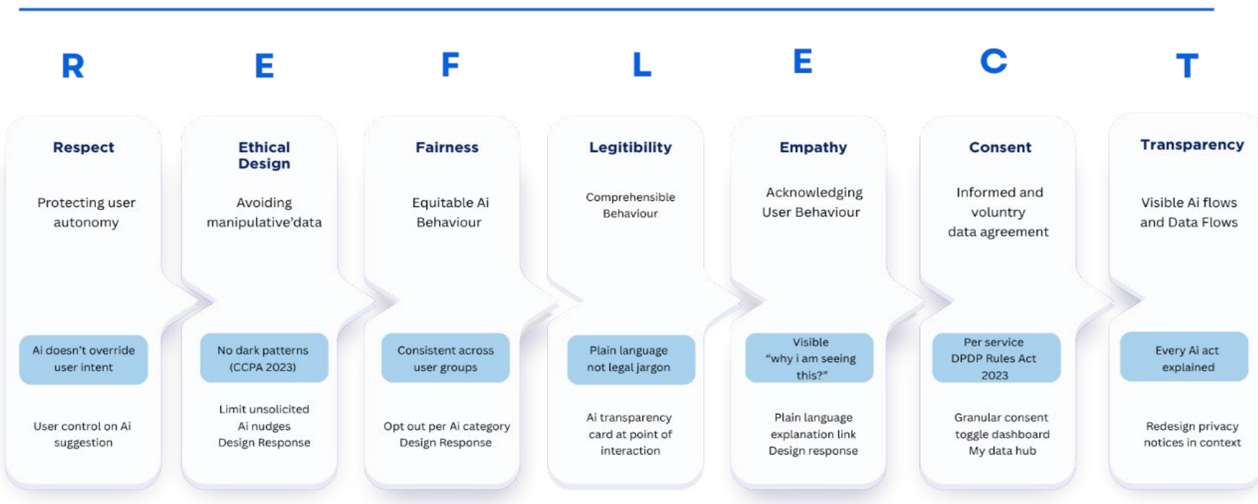

Figure 1: REFLECT Framework

Core Attitudinal Findings

Table 2 presents mean Likert scores for the five core attitudinal variables measured in this study.

The single highest mean in the entire dataset was for the statement that clear AI and data usage information would increase trust (M = 4.00) the most consistent signal across all respondents. Data concern (M = 3.90) also edged above convenience satisfaction (M = 3.82). That narrow gap is the privacy paradox in numeric form, and it is what the hypothesis tests examine formally.

Table 2: Descriptive statistics for core trust and privacy variables (N = 50).

Statement	*Mean (1–5)*	*Reading*
AI features make the app more convenient	3.82	Moderately positive
I trust AI features in these apps	3.44	Moderate / near neutral
I am concerned about how much personal data these apps collect	3.90	High concern
Clear information about AI and data usage would increase my trust	4.00	Strong agreement
I feel comfortable using multiple services in one super app	3.70	Moderately positive
AI-based fraud detection is acceptable to me	3.40	Cautious acceptance

Hypothesis Testing

Four hypotheses were tested using statistical methods matched to the scale and distribution of each variable pair. Results are summarised in Table 3.

H1 Privacy Concerns and AI Trust

Privacy concern scores were grouped into three categories (Low: 1–2, n = 7; Medium: 3, n = 6; High: 4–5, n = 37), with AI trust as the dependent variable in a one-way ANOVA. The result was significant: $F(2, 47) = 3.516$, $p = 0.038$. The direction was unexpected. The low concern group actually reported the lowest trust (M = 2.29), while the high concern group reported the highest (M = 3.59). A Pearson correlation confirmed a weakly positive relationship ($r = 0.304$, $p = 0.032$). H1 null rejected. Privacy concern significantly affects AI trust and the pattern suggests that higher concern and higher trust co-exist among more digitally experienced users, which is precisely what the privacy paradox predicts.

H2 Trust Spillover Across AI Services

Pairwise Pearson correlations were run across all six combinations of the four AI features. Table 4 presents the full pairwise Pearson correlation results for all six AI feature combinations. Every pair was statistically significant (all $p < .001$), with correlations

Table 3: Hypothesis test summary (N = 50).

	Hypothesis	*Test*	*Key Statistic*	*p-value*	*Decision*
H1	Privacy concerns affect AI trust	One-way ANOVA	F(2,47) = 3.516	0.038	Reject null ✓
H2	Trust spillover across AI services	Pearson r (6 pairs)	r = 0.533–0.742	<0.001	Reject null ✓
H3	Data transparency increases AI trust	t-test + Pearson r	t=3.068 / r=0.513	0.004 / <0.001	Reject null ✓
H4	Security perception increases AI trust	ANOVA + Pearson r	F=0.279 / r=0.086	0.757 / 0.550	Fail to reject ✗

Table 4: Pairwise Pearson correlation results for H2 Trust Spillover (N = 50).

Service Pair	*Pearson r*	*r²*	*p-value*	*Strength*
Fraud Detection ↔ Recommendations	0.562	0.316	<0.001	Moderate-Strong
Fraud Detection ↔ Chatbots	0.603	0.364	<0.001	Strong
Fraud Detection ↔ Notifications	0.642	0.412	<0.001	Strong
Recommendations ↔ Chatbots	0.718	0.515	<0.001	Very Strong
Recommendations ↔ Notifications	0.742	0.551	<0.001	Very Strong
Chatbots ↔ Notifications	0.533	0.284	<0.001	Moderate-Strong

ranging from r = .533 to r = .742. The two strongest relationships were between Recommendations and Notifications (r = 0.742) and Recommendations and Chatbots (r = 0.718), which suggests personalisation features sit at the centre of the platform's trust network.

H2 null rejected across all six pairs. Trust does not stay contained within individual features it moves across the platform, with personalisation features sitting at the centre of the network.

5.3.3 H3 Data Transparency and Trust

Respondents were split by median transparency demand score (High Demand: score ≥ 4, n = 38; Low Demand: score < 4, n = 12). The t-test yielded t(48) = 3.068, p = 0.004, with a mean trust difference of 1.18 points and a large effect size (Cohen's d = 0.97). Pearson correlation supported this: r = 0.513, p < 0.001, accounting for 26.3% of variance in AI trust. The behavioural data reinforces the picture: 80.0% of respondents (n = 40) said they would share more data if the app explained how it was used and gave them real control. H3 null rejected. Transparency demand is a strong predictor of AI trust, with a large effect size.

5.3.4 H4 Security Perceptions and Trust

Security was operationalised through trust in AI-based fraud detection. Three groups were formed: Low (n = 12), Medium (n = 12), and High (n = 26) fraud detection trust. One-way ANOVA returned F(2, 47) = 0.279, p = 0.757 non-significant and Pearson correlation confirmed r = 0.086, p = 0.550. A supplementary binary analysis of security breach concern also failed to reach significance (t = 0.445, p = 0.659). H4 null was not rejected. This is a null finding worth reading carefully. The questionnaire measured trust in a security feature, not security risk perception two distinct psychological constructs that should not be conflated. The result also points to a communication gap: the April 2026 UPI security upgrades are objectively significant, but they have not yet become visible to users at the interface level, and security that users cannot see does not translate into measurable trust.

Discussion

The Integrated Picture: Privacy Paradox Meets Spillover Effect

Taken together, the three rejected hypotheses form a coherent account of how AI trust works in India's super apps. H1 shows that privacy concern and AI trust are not opposites they co-exist, because the most concerned users are also the most digitally experienced. H2 shows that trust is platform-wide rather than feature-specific: correlations of r = 0.533–0.742 across all six AI feature pairs mean trust in one service is a reasonable proxy for trust in all others on the same platform. H3 gives the practical lever: transparency demand explains 26.3% of the variance in AI trust, and 80.0% of users say they would share more data for clearer explanations and genuine control.

The spillover architecture (H2) changes the economics of transparency investment. Improving transparency for one AI feature does not stop at that feature given how trust moves across the platform, it will likely lift trust in chatbots, fraud alerts, and notifications that the same user encounters later. A feature-by-feature cost-benefit analysis would miss this. The multiplier effect is already built into the platform's trust structure.

The H1 finding adds a wrinkle. Users who score 4 or higher on privacy concern but remain active AI feature users are not being inconsistent they exhibit what this study calls transaction fatigue: a learned resignation to blanket consent architectures that are legally acceptable but practically unreadable. One respondent put it plainly: 'For every superapp, I would not want to go through all the nitty and grittys. It is just like everything is written in T&C and we just check the box without reading.' Pushing transparency (the H3 direction) is more likely to change behaviour than waiting for privacy concern to become a deterrent concern without consequence is already the established equilibrium.

REFLECT Framework Assessment

Applying REFLECT to the hypothesis findings surfaces three interface failures with direct data support.

Transparency and Legibility are where the failures run deepest. 74.0% of users are highly concerned about data collection, yet most continue using AI features without understanding what their data is actually used for. That is a Legibility failure: the interface does not show AI logic at the point of interaction. The Transparency failure sits at the institutional level: privacy policies exist, but they are written in legal language at lengths that make genuine reading unrealistic. The H2 spillover finding raises the cost of both failures when a dark pattern surfaces or a privacy failure occurs in one AI feature, trust damage spreads platform-wide. A correlation structure of r = 0.533–0.742 means a trust crisis in one service is functionally a platform crisis.

Dark Patterns: The Regulatory Dimension

The CCPA's November 2023 notification identified 13 prohibited dark patterns in Indian digital commerce. Three of them show up directly in this study's findings. Interface Interference using design to steer attention away from privacy-protective options is visible in the near-universal practice of burying privacy settings deep in menus while personalisation defaults sit prominently. Forced Action requiring users to accept unnecessary data collection as the price of service access is inherent in the bundled consent model, where opting out of location-based advertising is structurally equivalent to refusing the payment service. Basket Sneaking adding implicit data-sharing consents inside primary service agreements is embedded in most super app onboarding flows. The DPDP Rules 2026 create enforcement pressure, but the trust risk arrives first: given the spillover architecture, a user who loses trust in one AI feature is already likely to disengage from multiple services simultaneously.

Design Recommendations

Three UI/UX interventions follow from the data. Each is grounded in the hypothesis results, mapped against REFLECT dimensions, and designed to be compliant with India's DPDP regulatory framework.

AI Transparency Cards

Every AI-generated output a recommendation, a chatbot response, a fraud alert should carry a collapsible AI Transparency Card. The card shows in plain language: what data produced this output, why it was surfaced, and a one-tap option to opt out of this type of AI interaction. It directly targets the Legibility failure and speaks to the 80.0% of users who said transparency would raise their trust. There is also a commercial argument here: users who understand why a recommendation appeared are more likely to act on it and more willing to share the data that improves future ones. Transparency and personalisation quality are not in tension they reinforce each other.

My Data Hub

Rather than presenting all consent decisions at onboarding and moving on, platforms should give users a persistent, accessible My Data Hub a dashboard where they can see which data categories are shared with which services, toggle permissions on or off, review a log of AI decisions made on their behalf, and submit deletion requests for specific data types. Two taps from the main menu. Not buried in settings. My Data Hub targets the Consent and Control failures in the REFLECT assessment directly. It also brings platforms in line with the DPDP Act 2023's requirement for specific, purpose-limited consent and the right to withdraw consent at any time (Section 6(4)). The deeper point is simpler: it gives users the interface tools to move from resigned acceptance to something closer to informed agreement.

Trust Momentum Onboarding

Instead of front-loading all permission requests at registration when users have the least context and the lowest trust platforms should time data requests to the user's natural service journey. When someone first opens the payment section, ask for payment-relevant permissions and briefly explain the exchange. When they navigate to travel bookings for the first time, ask for location access and explain specifically how it improves travel suggestions. The H2 spillover finding supports this directly: users who trust the AI in their primary domain (typically payments) are measurably more open to sharing data in adjacent domains. Incremental permission architecture converts that natural trust progression into a structured, ethically sound onboarding flow rather than a single risk-loaded moment at sign-up.

Conclusion

Three findings from this study have direct implications for how super apps are designed and regulated in India.

First, trust is platform-wide, not feature-specific. The trust spillover effect (Pearson $r = 0.533–0.742$ across all six AI feature pairs, all $p < 0.001$) means every trust gain and every trust loss moves across the entire platform. Getting AI trust right in one feature is not a UX nicety it is a growth multiplier with measurable statistical force.

Second, the privacy paradox is a design problem, not an inherent contradiction. AI convenience ($M = 3.82$) sits alongside high data concern ($M = 3.90$) not because users are being inconsistent, but because they have accepted opaque consent as the price of access. Eighty percent of them said the same thing: clear explanations and genuine control would change their behaviour. The gap between the convenience users already enjoy and the trust they still withhold is precisely the gap that transparent interface design can close.

Third, current super app interfaces fail the REFLECT framework's Transparency, Legibility, and Consent dimensions. The DPDP Rules 2026 require purpose-limited, plain-language consent notices a necessary regulatory floor. The H3 result ($r = 0.513$, Cohen's $d = 0.97$) suggests that floor is also sufficient when implemented well: users who already receive transparency report significantly higher trust. The central design challenge this study points to is closing the gap between what the regulation now requires and what users are actually encountering on screen.

Limitations

Several limitations apply. The final sample of $N = 50$ fell well short of the intended $N = 200$, which constrains statistical power for sub-group analyses by platform or service category. Convenience sampling limits how far the findings can be generalised. The cross-sectional design captures one regulatory and technological moment in a sector that is moving fast, and Likert self-reports may not map cleanly onto actual user behaviour. The sample skews heavily toward younger, urban, educated users aged 18–34 (98.0%); how these findings translate to Tier 2 and Tier 3 cities or older demographics remains an open question. The privacy subscale returned a Cronbach's α of 0.521, indicating that the four items tap distinct privacy dimensions rather than a single construct analytically defensible, but worth revisiting in future instrument design. The H4 gap is the most tractable to fix: a direct item such as 'I am worried about security breaches in super apps' would capture security risk perception as its own construct rather than conflating it with trust in a specific security feature.

Future Scope

Three directions are worth pursuing. A pan-India study covering Tier 2 and Tier 3 cities would test whether the trust spillover architecture holds at different levels of digital literacy, infrastructure quality, and AI familiarity. A longitudinal study tracking the same cohort over 12–24 months would show whether the April 2026 UPI security upgrades and the DPDP Rules 2026 once they become part of users lived experience actually shift trust perceptions, and which communication design approaches get there faster. Platform-specific studies comparing PhonePe, Paytm, and Tata Neu directly, using structural equation modelling on larger stratified samples, could move the field beyond correlation and toward identifying the actual causal paths between payment trust, AI trust, and cross-service adoption.

References

Ahire, V. Y. (2025). Exploring user perception and trust in generative AI applications: A primary study among Indian consumers. Social Science Journal for Advanced Research, 5(4), 114–120. https://doi.org/10.5281/zenodo.16791916

Antonio, G. R., & Kurniawan, M. P. (2025). Modelling user behavior in super apps: Integrating trust and personal innovativeness into the UTAUT3 framework for Gojek. Journal of Entrepreneurship, 4(3), 16–44. https://doi.org/10.56943/joe.v4i3.808

Bach, T. A., Khan, A., Hallock, H., Beltrão, G., & Sousa, S. (2022). A systematic literature review of user trust in AI-enabled systems: An HCI perspective. International Journal of Human-Computer Interaction, 40(5), 1251–1266. https://doi.org/10.1080/10447318.2022.2138826

Bayashot, Z. A. (2025). AI-powered personalization in super apps: International case studies on user engagement. Journal of Software Engineering and Applications, 18(7), 268–285. https://doi.org/10.4236/jsea.2025.187016

Cai, Y., Zhang, Z., Yao, M., Liu, J., Zhao, X., Fu, X., Li, R., Li, Z., Chen, X., Guo, Y., & Li, D. (2025). I can tell your secrets: Inferring privacy attributes from mini-app interaction history in super-apps (arXiv:2503.10239). arXiv. http://arxiv.org/abs/2503.10239

Central Consumer Protection Authority. (2023). Guidelines for prevention and regulation of dark patterns. Ministry of Consumer Affairs, Government of India.

Chandra, S., Sikdar, P. K., & Kumar, V. (1995). Dynamic PCU estimation of capacity on urban roads. Journal of Indian Roads Congress, 23, 17–28.

Chandra, S. (2004). Capacity estimation procedure for two-lane roads under mixed traffic conditions. Journal of Indian Roads Congress, 165(1), 139–170.

Chhetri, V., Upadhyay, K., Siddique, A. B., & Farooq, U. (2025). What users value and critique: Large-scale analysis of user feedback on AI-powered mobile apps (arXiv:2506.10785). arXiv. http://arxiv.org/abs/2506.10785

Cvetkovic, A., Savolainen, I., Celuch, M., Heiskari, M., Soares Ruokosuo, E., Arriaga, P., Koike, M., & Oksanen, A. (2026). User trust in AI and major tech companies in twelve countries. Behaviour & Information Technology. Advance online publication. https://doi.org/10.1080/0144929X.2026.2619648

Digital Personal Data Protection Act, 2023, No. 22 of 2023, Ministry of Electronics and Information Technology, Government of India. (2023).

Digital Personal Data Protection Rules, 2026 [Draft rules under the DPDP Act 2023], Ministry of Electronics and Information Technology, Government of India. (2026).

Fang, Y. H., Liao, C. H., & Li, C. Y. (2024). Super app on demand: Exploring the impact of service synergy on willingness to use a new service. Electronic Commerce Research and Applications, 67, Article 101430. https://doi.org/10.1016/j.elerap.2024.101430

Hasselwander, M. (2024). Digital platforms' growth strategies and the rise of super apps. Heliyon, 10(5), Article e25856. https://doi.org/10.1016/j.heliyon.2024.e25856

Hasselwander, M., Kriswardhana, W., Esztergar-Kiss, D., Lah, O., & Steinberg, M. (2025). One app for everything: A multidisciplinary review of super apps [Working paper]. SSRN. https://doi.org/10.2139/ssrn.5772697

Hasselwander, M., & Weiss, D. (2025). Consumer preferences for super app services: E-commerce, social media, and banking dominate. European Research on Management and Business Economics, 31(2), Article 100284. https://doi.org/10.1016/j.iedeen.2025.100284

National Payments Corporation of India. (2024). UPI product statistics. https://www.npci.org.in/what-we-do/upi/product-statistics

National Payments Corporation of India. (2026). UPI security enhancement guidelines – April 2026 circular.

Ng, S. W. T., & Zhang, R. (2025). Trust in AI chatbots: A systematic review. Telematics and Informatics, 97, Article 102240. https://doi.org/10.1016/j.tele.2025.102240

Norberg, P. A., Horne, D. R., & Horne, D. A. (2007). The privacy paradox: Personal information disclosure intentions versus behaviors. Journal of Consumer Affairs, 41(1), 100–126. https://doi.org/10.1111/j.1745-6606.2006.00070.x

Yang, Y., Wang, C., Zhang, Y., & Lin, Z. (2023). SoK: Decoding the super app enigma: The security mechanisms, threats, and trade-offs in OS-alike apps (arXiv:2306.07495). arXiv. http://arxiv.org/abs/2306.07495