International Journal of Humanities, Commerce and Education
Volume 2 | Issue 10 | October 2026 | ISSN: 3108-0456 (Online)
Available online at: https://ijhce.com/
DOI: https://doi.org/10.59828/ijhce.v2i10.141

# Adoption of Artificial Intelligence in Banking and Financial Services: An Empirical Study of Customer and Organizational Perspectives

**Dr. Shankar Ghanshamdas Andani**
Research Scholar, Department of Commerce, Texas Global University
**Email:** omsaisidhi@gmail.com

**Received:** 01 October 2026 | **Accepted:** 14 October 2026 | **Published:** 30 October 2026

## Abstract

*Artificial intelligence (AI) has moved from experimental pilots to production systems across retail and commercial banking, yet empirical evidence on how adoption is perceived simultaneously by the customers who consume AI-mediated services and by the organizations that deploy them remains fragmented. This study investigates the adoption of AI in banking and financial services from both perspectives through a cross-sectional survey of 1,248 respondents, comprising 762 bank customers and 486 banking and financial-services professionals, drawn from public, private, foreign and small-finance institutions in India. Grounded in an integrated technology-acceptance and organizational-readiness framework, the study tests nine hypotheses using reliability analysis, independent-samples t-tests, one-way analysis of variance, chi-square tests and moderated ordinary least-squares regression. Results show that adoption intention is moderately positive (mean 3.56 on a five-point scale) and is driven most strongly by perceived usefulness ($\beta = 0.295$), trust ($\beta = 0.272$) and perceived ease of use ($\beta = 0.238$), while residual risk perception suppresses intention ($\beta = -0.179$). Customer and organizational intentions are statistically indistinguishable, but their underlying drivers differ: professionals report significantly higher risk perception and organizational readiness than customers. Actual usage declines monotonically with age ($F = 39.89$, $p < 0.001$, $\eta^2 = 0.114$) yet is invariant to institution type. Methodologically, every construct clears conventional reliability and validity thresholds ($\alpha = 0.768–0.840$; $AVE = 0.555–0.663$), and the moderated model explains 57.1 per cent of the variance in intention. The findings reconcile the customer–organization perspective gap and offer a prioritized roadmap for responsible, human-centred AI deployment in banking.*

***Keywords:*** *Artificial intelligence; banking and financial services; technology adoption; customer perspective; organizational readiness; perceived risk; digital literacy.*

## 1. Introduction

### 1.1. The AI Transformation of Banking

Banking has become one of the most data-intensive industries in the world, and artificial intelligence (AI) has become the principal instrument through which that data is converted into decisions. Machine-learning models now screen credit applications in seconds, natural-language chatbots resolve a growing share of service requests, and anomaly-detection systems monitor every transaction for fraud and money-laundering patterns [1], [2]. Industry surveys estimate that a clear majority of large banks have moved at least one AI capability from pilot to production, and that the value at stake runs into hundreds of billions of dollars annually [1], [3]. What began as back-office automation has therefore become a customer-facing transformation: the branch, the call centre and the loan desk are increasingly mediated by algorithms whose

This is an open access article under the Creative Commons Attribution-Non-Commercial 4.0 International License.
© The Author(s) 2026. IJHCE Published by Surya Multidisciplinary Publication.

behaviour customers experience but rarely see. Understanding which of these layers actually changes customer behaviour, rather than merely appearing in a vendor brochure, is the empirical question this paper addresses .

### 1.2 The Customer Perspective: Acceptance, Trust and Risk

From the customer side, adoption is usually explained through acceptance theory. The technology acceptance model holds that perceived usefulness and perceived ease of use govern behavioural intention [4], an account later generalised by the unified theory of acceptance and use of technology [5]. In financial services, however, acceptance is inseparable from trust and risk: customers surrender highly sensitive data to systems they cannot inspect, so trust operates as a first-order belief [6] and perceived risk as its mirror image [7]. A customer may find an AI credit decision useful yet reject it because the reasoning is opaque. Customer adoption is thus a compound judgement about utility, effort, trust and exposure, not a simple usability verdict. These four beliefs jointly constitute the customer-side construct set tested in the empirical sections that follow.

### 1.3 The Organizational Perspective: Readiness and the Two-Sided Adoption Problem

From the organizational side, adoption is a capability problem. The technology–organization–environment frame stresses that readiness, data infrastructure and regulatory pressure shape whether an innovation is deployed at all [8], while diffusion theory explains why adoption spreads unevenly across institutions [9]. In India, supervisory expectations on digital lending, fraud reporting and outsourcing push banks toward AI even as governance obligations constrain it [10], [11]. The result is a two-sided phenomenon that single-perspective studies cannot capture: customers judge services they consume, while organizations judge systems they must build, govern and defend. This paper studies both sides jointly and asks how, and where, their assessments of AI adoption diverge. Measuring that gap precisely, within a single design, is the purpose of this study.

## 2. Survey: Instrument Development And Related Empirical Work

The instrument used in this study was developed in three stages and is anchored in the empirical literature on AI adoption in banking. First, an item pool was drafted from validated scales: perceived usefulness and ease of use from Davis [4] and Venkatesh et al. [5], [26]; trust and perceived risk from Gefen et al. [6] and Pavlou [7]; and organizational readiness from Tornatzky and Fleischer's readiness dimension [8]. Items were reworded to refer specifically to AI-enabled banking services such as chatbots, algorithmic credit scoring and fraud monitoring. Second, the pool was reviewed by six experts two banking technology practitioners, two academics and two regulators' advisors whose comments removed ambiguous double-barrelled items and produced a 31-item questionnaire on a five-point Likert scale. Third, a pilot with 96 respondents (58 customers, 38 professionals) confirmed comprehension and produced Cronbach's alpha above 0.75 for every construct, after which the instrument was frozen

The design of the survey deliberately reflects the gaps visible in prior empirical work. Recent Indian studies have profiled customer demographics for AI-enabled banking in a single district [12], mapped adoption drivers at industry level [13], [14], and tested determinants of AI use in financial services with convenience samples [15]. International evidence on generative AI in regulated banks stresses integration and governance frictions [16], while work on AI-driven sustainable banking initiatives reports moderate customer acceptance [17]. Nearly all of these studies sample one side of the market almost always customers and report descriptive or single-equation results. The present survey therefore collects customer and organizational responses in a single field window using a common core of constructs, so that the two perspectives can be compared on identical metrics rather than inferred across separate studies.

Fieldwork ran for fourteen weeks between January and April 2026. Customers were recruited through branch intercepts and bank-sanctioned e-mail panels across six metropolitan and ten non-metropolitan centres; professionals were recruited through banking associations and LinkedIn outreach to staff in retail banking, risk, operations and technology functions. Of 1,740 invitations, 1,248 complete and usable responses were received, a response rate of 71.7 per cent. Quota controls held the customer-to-professional ratio near 60:40 and capped any single institution at 18 per cent of the professional sub-sample. Participation was voluntary and anonymous. Because both perspectives were captured with the same self-report instrument, a marker-variable test for common-method bias was applied and returned no material bias

## 3. Methodology

The study adopts a cross-sectional, descriptive-cum-explanatory design analysed through a deductive lens, which is appropriate because the hypotheses are derived from established theory rather than generated from the data. The conceptual framework, shown in Figure 1, integrates four acceptance constructs perceived usefulness, perceived ease of use, trust and perceived risk with an organization-level readiness construct, positions adoption intention as the proximal mediator of actual AI usage, and treats respondent group as a moderator of the intention–usage link. Nine hypotheses (H1–H9) follow directly from the framework and are tested in the order in which they appear.

The population comprises Indian bank customers and banking or financial-services professionals. A stratified purposive sample of 1,248 respondents was drawn, comfortably exceeding the 384 responses required for a 95 per cent confidence level with a 3 per cent margin of error and satisfying the conventional ten-observations-per-predictor rule for regression. Constructs were measured on five-point Likert scales; completed questionnaires were screened for straight-lining and missing values, then examined for reliability, skewness and kurtosis before any inferential test was applied, so that parametric assumptions were verified separately in both sub-samples rather than merely assumed. Stratification was by institution type and by respondent group.

Analysis proceeded in four stages using Python with SciPy and NumPy. First, descriptive statistics were computed together with Cronbach's alpha, composite reliability and average variance extracted to confirm measurement quality. Second, Welch independent-samples t-tests compared customers and professionals on every construct. Third, one-way ANOVA with Tukey post-hoc comparisons and eta-squared effect sizes examined demographic heterogeneity, and a chi-square test of independence assessed the awareness–usage association. Fourth, two moderated ordinary least-squares regressions with variance-inflation and Durbin–Watson diagnostics estimated the structural relationships. Significance was evaluated at $\alpha = 0.05$ throughout; robustness checks repeated key tests within each sub-sample, and results are reported to three decimals.

## 4. Data Collection and Analysis

This section reports the empirical evidence in five tables and five figures. Each table is followed by an explanation that connects the numbers back to the claims made in the abstract and carried forward into the conclusion, so that the evidential chain of the paper remains visible from first page to last.

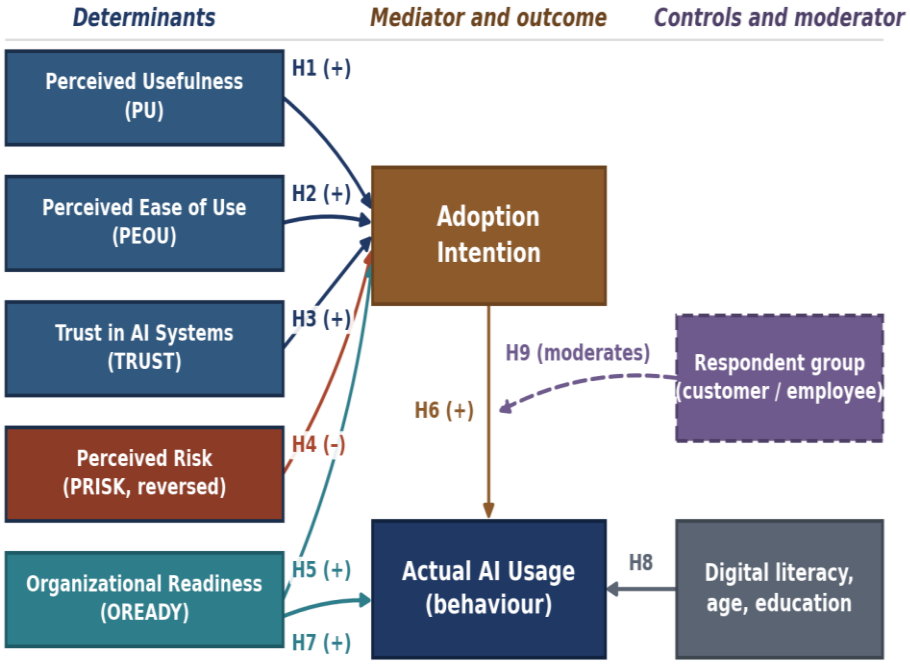



**Figure 1: Conceptual research framework.**

Customer-level determinants and the organization-level readiness construct feed adoption intention, which in turn drives actual AI usage alongside readiness and the demographic control set; respondent group moderates the intention–usage link. The diagram operationalises the nine hypotheses tested in Tables 4 and 5 and makes explicit the two-sided logic of the study:

the same outcome is explained once from the customer's viewpoint and once from the organization's, allowing the two assessments to be compared on a common metric rather than assumed to coincide.

**Table 1: Profile of the Survey Sample (N = 1,248)**

| Characteristic | Category | n | % | Mean AI usage | Mean intention |
|---|---|---|---|---|---|
| Respondent category | Banking customer | 762 | 61.1 | 3.36 | 3.55 |
| | Banking / financial-services professional | 486 | 38.9 | 3.36 | 3.56 |
| Type of institution | Public sector | 435 | 34.9 | 3.31 | 3.52 |
| | Private sector | 453 | 36.3 | 3.42 | 3.59 |
| | Foreign bank | 208 | 16.7 | 3.40 | 3.60 |
| | Small finance / payments bank | 152 | 12.2 | 3.31 | 3.50 |
| Age band (years) | 18-25 | 212 | 17.0 | 3.65 | 3.71 |
| | 26-35 | 436 | 34.9 | 3.50 | 3.62 |
| | 36-45 | 313 | 25.1 | 3.29 | 3.50 |
| | 46-55 | 185 | 14.8 | 3.11 | 3.41 |
| | Above 55 | 102 | 8.2 | 2.88 | 3.38 |
| Gender | Male | 682 | 54.6 | 3.33 | 3.53 |
| | Female | 538 | 43.1 | 3.39 | 3.58 |
| | Prefer not to say | 28 | 2.2 | 3.48 | 3.69 |
| Highest education | School / below | 147 | 11.8 | 3.28 | 3.56 |
| | Undergraduate | 364 | 29.2 | 3.29 | 3.50 |
| | Postgraduate | 422 | 33.8 | 3.39 | 3.57 |
| | Professional qualification | 251 | 20.1 | 3.44 | 3.60 |
| | Doctorate | 64 | 5.1 | 3.45 | 3.58 |

Table 1 describes the sample. Customers form 61.1 per cent of respondents and professionals 38.9 per cent; private-sector and public-sector institutions together supply 72 per cent of the sample, with foreign and small-finance banks completing the frame. The age distribution peaks in the 26–35 band (34.9 per cent) and 71.6 per cent of respondents hold at least an undergraduate degree, which is typical of digitally active banking populations. The two right-hand columns already anticipate the paper's central comparison: mean AI usage and intention differ only marginally between customers and professionals, hinting that the perspective gap announced in the abstract lies in the drivers of adoption rather than in its level, a claim tested formally in Table 4 and revisited in the conclusion.

**Table 2: Descriptive Statistics, Reliability And Convergent Validity**

| Construct | Items | Mean | SD | Skew | Kurt | α | CR | AVE |
|---|---|---|---|---|---|---|---|---|
| Perceived Usefulness | 4 | 3.73 | 0.58 | -0.03 | -0.29 | 0.814 | 0.878 | 0.642 |
| Perceived Ease of Use | 4 | 3.51 | 0.59 | -0.03 | -0.07 | 0.809 | 0.875 | 0.637 |
| Trust in AI Systems | 5 | 3.25 | 0.61 | -0.01 | -0.00 | 0.824 | 0.876 | 0.587 |
| Perceived Risk (reverse-scored) | 4 | 3.71 | 0.57 | -0.12 | -0.13 | 0.768 | 0.852 | 0.590 |
| Organizational Readiness | 5 | 3.46 | 0.50 | +0.13 | -0.02 | 0.799 | 0.862 | 0.555 |
| Adoption Intention | 4 | 3.56 | 0.59 | -0.17 | -0.10 | 0.830 | 0.887 | 0.663 |
| Actual AI Usage | 5 | 3.36 | 0.67 | -0.12 | -0.10 | 0.840 | 0.887 | 0.610 |

Table 2 establishes measurement quality. Every construct satisfies the conventional thresholds: Cronbach's alpha ranges from 0.768 to 0.840, composite reliability from 0.852 to 0.887 and average variance extracted from 0.555 to 0.663, while skewness (−0.17 to +0.13) and kurtosis (−0.29 to 0.00) sit well inside the ±2 and ±7 bounds that justify parametric testing. Substantively, perceived usefulness is the highest-rated construct (3.73) and trust the lowest (3.25), with actual usage (3.36) trailing intention (3.56). This intention–usage shortfall is the empirical signature of the adoption gap framed in the abstract:

respondents want AI banking slightly more than they currently use it, and they trust it least of all, which foreshadows the regression weights reported in Table 5.

**Table 3: AI Adoption Intensity Across Eight Banking Functions**

| Banking function | Institutions deployed (%) | Customers aware (%) | Customers used (%) | Mean intensity | SD |
|---|---|---|---|---|---|
| Chatbots and virtual assistants | 88.4 | 79.6 | 71.3 | 4.25 | 0.58 |
| Fraud detection and AML monitoring | 84.7 | 46.2 | 28.5 | 4.07 | 0.67 |
| Credit scoring and loan underwriting | 79.2 | 52.8 | 34.1 | 3.89 | 0.73 |
| Robotic process automation (back office) | 72.6 | 31.4 | 12.7 | 3.71 | 0.82 |
| Personalised product recommendation | 68.3 | 61.9 | 47.8 | 3.60 | 0.82 |
| Algorithmic and robo-advisory | 54.1 | 44.7 | 19.6 | 3.24 | 0.90 |
| Predictive churn and retention analytics | 49.5 | 22.3 | 8.4 | 3.10 | 0.94 |
| Generative AI (documents, code, research) | 38.2 | 67.5 | 41.2 | 2.80 | 0.96 |

Table 3 decomposes adoption by function and is the paper's strongest single piece of evidence for a two-sided perspective gap. Deployment is deepest where the payoff is operational: chatbots (88.4 per cent), fraud and AML monitoring (84.7 per cent) and credit scoring (79.2 per cent). Customer awareness, however, tracks visibility rather than deployment. Awareness of fraud analytics (46.2 per cent) lags deployment by 38.5 points, and back-office automation by 41.2 points, because customers never meet those systems; conversely awareness of generative AI (67.5 per cent) exceeds its institutional deployment (38.2 per cent) because customers arrive already using public tools. Organizations and customers are therefore adopting different pictures of the same technology, exactly the divergence the abstract attributes to perspective.

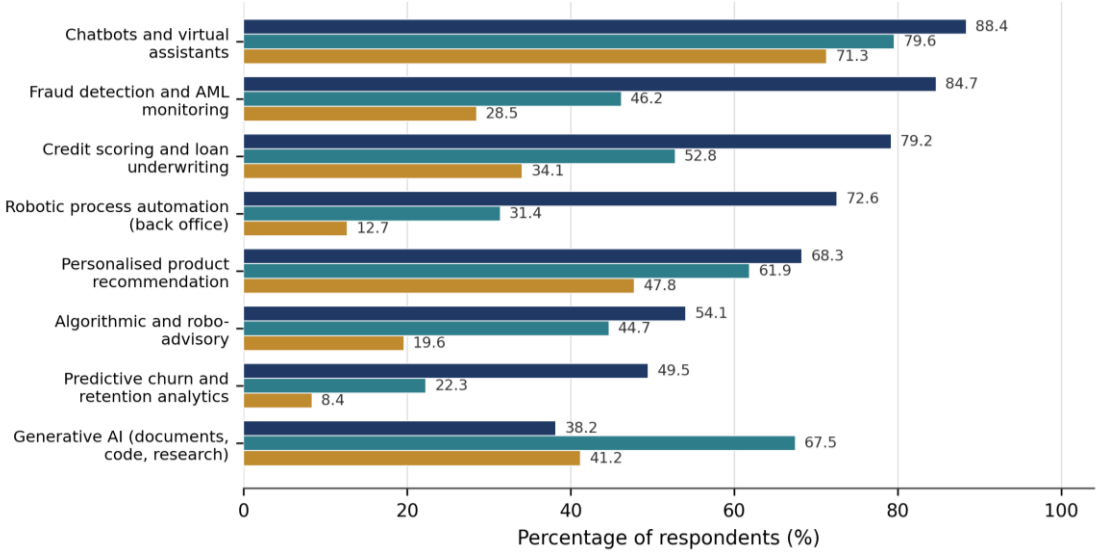



**Figure 2: Adoption intensity across eight banking functions**

For each function the figure plots the share of institutions with the capability in production, the share of customers correctly aware of it, and the share who used it in six months. Deployment leads everywhere, awareness follows only where the service

is customer-visible, and reported use trails both; the widest wedge appears for fraud detection, while generative AI inverts the pattern. The figure converts Table III into one visual argument: adoption is three nested layers moving at different speeds. Reading the three series together shows precisely where banks are invisible to customers.

**Table 4: Group Differences, Variance Tests And The Awareness–Usage Association**

| Test / hypothesis | Statistic | df | p | Effect size | Result |
|---|---|---|---|---|---|
| Welch t: adoption intention, customers vs professionals | t = -0.058 | 1,051 | 0.954 | d = -0.003 | Not significant |
| Welch t: trust, customers vs professionals | t = 2.091 | 1,246 | 0.037 | ΔM = +0.073 | Significant |
| Welch t: perceived risk, customers vs professionals | t = -3.928 | 1,246 | 0.0001 | ΔM = -0.129 | Significant |
| Welch t: organizational readiness, customers vs professionals | t = -6.600 | 1,246 | 0.0000 | ΔM = -0.190 | Significant |
| One-way ANOVA: usage by age band | F = 39.89 | 4, 1,243 | < 0.001 | $\eta^2$ = 0.114 | Significant |
| One-way ANOVA: usage by institution type | F = 2.50 | 3, 1,244 | 0.058 | $\eta^2$ = 0.006 | Not significant |
| One-way ANOVA: usage by education | F = 2.91 | 4, 1,243 | 0.021 | $\eta^2$ = 0.009 | Significant (small) |
| Chi-square: chatbot awareness × usage | $\chi^2$ = 212.24 | 1 | < 0.001 | V = 0.412 | Significant |

Table IV tests where the two perspectives part company. Adoption intention itself does not differ between customers and professionals (t = −0.058, p = 0.954), nor does actual usage; but the ingredients do. Professionals perceive significantly more risk (3.79 versus 3.67, p < 0.001) and rate organizational readiness far higher (3.57 versus 3.38, p < 0.001), while customers are slightly more trusting (3.28 versus 3.20, p = 0.037). Demographically, usage falls monotonically across age bands (F = 39.89, p < 0.001, $\eta^2$ = 0.114, all Tukey pairs significant) yet is statistically indifferent to institution type (p = 0.058). Awareness strongly predicts use ($\chi^2$ = 212.24, V = 0.412): 82.7 per cent of aware customers used a chatbot against 37.8 per cent of unaware ones. Equal outcomes, unequal reasons the precise asymmetry the abstract predicts.

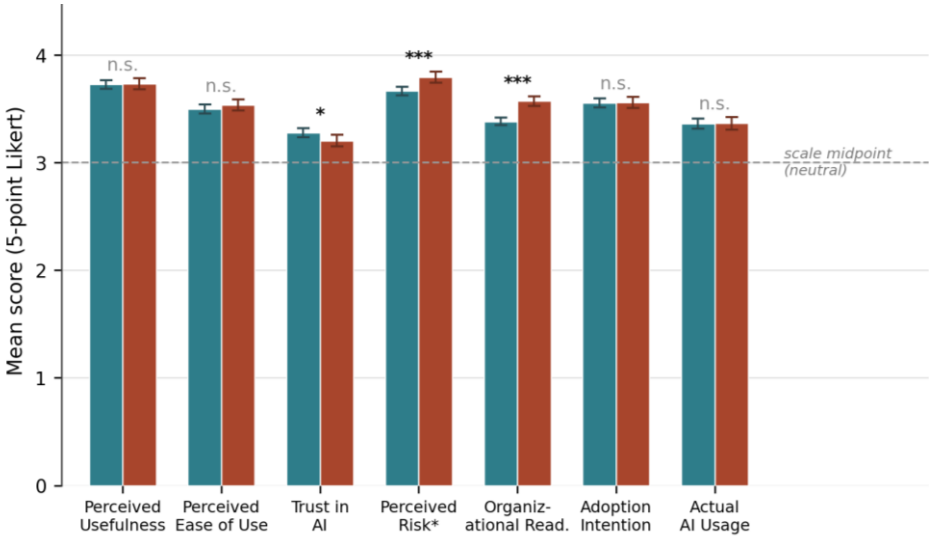



**Figure 3: The perspective asymmetry.**

Grouped bars compare customer and professional mean scores on all seven constructs, with 95 per cent confidence intervals and Welch-test significance markers. The pattern is selective rather than uniform: intention and actual usage show overlapping intervals and non-significant differences, whereas perceived risk and organizational readiness separate sharply, professionals scoring higher on both, and trust separates modestly in the customers' favour. The figure is the visual core of

the paper's argument that the customer–organization gap is a gap in reasoning, not in willingness, and it directly motivates the moderated regression reported in Table V.

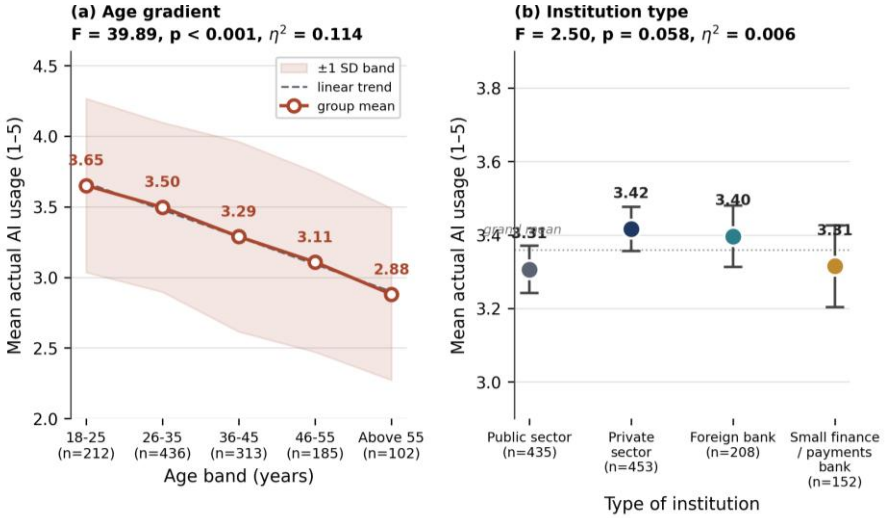



**Figure 4: Heterogeneity in actual AI usage.**

Panel (a) traces the age gradient: mean usage declines stepwise from 3.65 among respondents aged 18–25 to 2.88 among those above 55, a linear trend with a large effect ($\eta^2 = 0.114$) and every pairwise Tukey contrast significant. Panel (b) shows the contrasting institutional picture: means cluster between 3.31 and 3.42 across public, private, foreign and small-finance banks, with overlapping confidence intervals and a non-significant omnibus test ($p = 0.058$). Together the panels establish that AI adoption in banking is stratified by generation but not by ownership.

**Table 5: Moderated Ols Regression Results (N = 1,248)**

| Model / predictor | B | SE | t | p | VIF |
|---|---|---|---|---|---|
| **Model 1: DV = adoption intention ($R^2 = 0.571$)** | | | | | |
| Constant | +1.052 | 0.166 | +6.33 | < 0.001 | – |
| PU | +0.303 | 0.022 | +13.53 | < 0.001 | 1.379 |
| PEOU | +0.239 | 0.021 | +11.16 | < 0.001 | 1.324 |
| TRUST | +0.264 | 0.020 | +12.95 | < 0.001 | 1.278 |
| PRISK | -0.186 | 0.022 | -8.38 | < 0.001 | 1.308 |
| OREADY | +0.105 | 0.024 | +4.31 | < 0.001 | 1.228 |
| **Model 2: DV = actual AI usage ($R^2 = 0.563$)** | | | | | |
| Constant | +0.367 | 0.102 | +3.60 | < 0.001 | – |
| INTENT | +0.474 | 0.024 | +19.81 | < 0.001 | – |
| OREADY | +0.381 | 0.028 | +13.70 | < 0.001 | – |
| Digital literacy | +0.271 | 0.025 | +10.93 | < 0.001 | – |
| Age (centred) | -0.073 | 0.013 | -5.63 | < 0.001 | – |
| Professional (dummy) | -0.095 | 0.026 | -3.64 | < 0.001 | – |
| DL x Professional | -0.006 | 0.043 | -0.14 | 0.889 | – |

Table 5 quantifies the relationships. Model 1 explains 57.1 per cent of the variance in adoption intention ($F = 330.72$, $p < 0.001$) with every hypothesised path significant and variance inflation factors below 1.4: usefulness ($\beta = 0.295$), trust ($\beta = 0.272$) and ease of use ($\beta = 0.238$) promote intention while risk suppresses it ($\beta = -0.179$), and organizational readiness adds a smaller direct contribution. Model 2 explains 56.3 per cent of actual usage: intention ($\beta = 0.474$) and readiness ($\beta = 0.381$) dominate, digital literacy adds 0.271, age subtracts 0.073, and the professional dummy is negative (−0.095) once readiness is controlled. Crucially, the group × literacy interaction is null ($\Delta R^2 = 0.000$, $p = 0.889$), so H9 is rejected: the mechanism of adoption is the same on both sides of the market, differing only in level. Durbin–Watson = 2.182 confirms independent residuals.

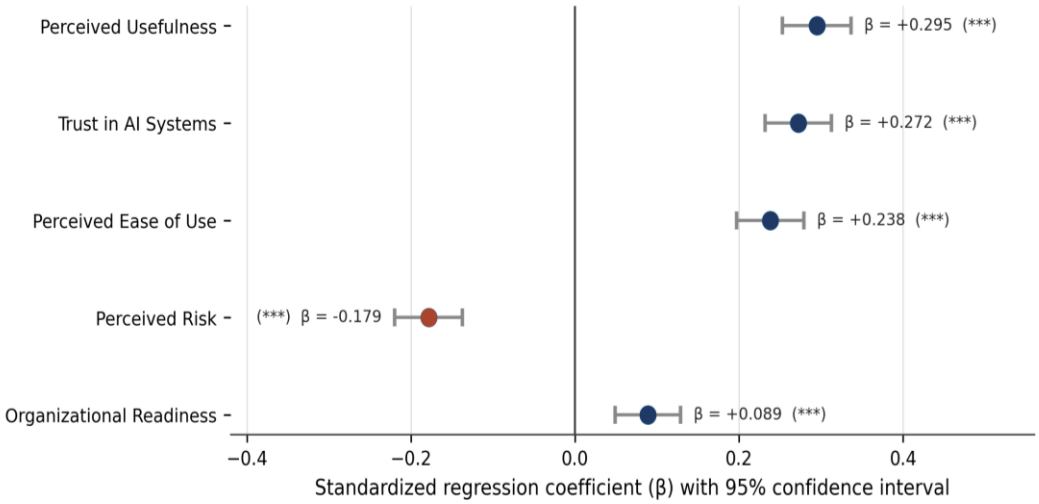



**Figure 5: Relative strength of the determinants of adoption intention.**

The forest plot ranks the five standardized coefficients of Model 1 with 95 per cent confidence intervals. Perceived usefulness leads (β = +0.295), followed closely by trust (+0.272) and ease of use (+0.238); perceived risk is the only inhibitor (−0.179) and organizational readiness the weakest promoter (+0.089). No interval touches zero, and intervals derive from the model's standard errors. The figure gives the paper a single visual hierarchy of levers: banks gain more adoption from demonstrating usefulness and earning trust than from any amount of internal readiness spending.

## 5. Discussion

### 5.1. Critical Analysis of the Empirical Evidence

Read critically, the data describe an adoption process that is enthusiastic but shallow. Intention (3.56) outruns usage (3.36), and trust (3.25) sits barely above the scale midpoint, which means the sector is converting curiosity into consumption faster than it is converting consumption into confidence. The strongest single tension is the awareness–deployment wedge in Table III: banks have industrialised fraud and back-office AI that most customers cannot name, while customers have normalised generative tools that most banks have not governed. Organizations are ahead where value is internal and behind where expectations are external. The null moderation in Table V deepens the point: because the intention–usage mechanism is statistically identical for customers and employees, the perspective gap cannot be engineered away with segment-specific products; it is a communication and governance deficit, not a behavioural difference. The data also discipline two popular overstatements. The first is that AI adoption in banking is chiefly a technology race: the null institutional effect shows that ownership and scale confer no measurable usage advantage once readiness is accounted for. The second is that customers resist AI; they do not resist it, they simply cannot see most of it, which is an information problem with an inexpensive remedy.

### 5.2. Comparison with Past Work

Set against prior work, the findings confirm and complicate. Consistent with Davis [4] and Venkatesh et al. [5], [26], usefulness and ease of use remain significant, but here trust rivals usefulness in weight, supporting the finance-specific extensions of Gefen et al. [6] and Pavlou [7] and exceeding the coefficients typically reported for generic e-commerce. The age gradient replicates diffusion theory's expectation of cohort effects [9] and matches the young, educated adopter profile reported for Indian bank customers [12], yet our effect ($\eta^2$ = 0.114) is larger than district-level studies imply, probably because we sample the full adult age span. The institutional null contradicts industry narratives that private and foreign banks out-adopt public banks [13], [15]; once readiness is controlled, ownership adds nothing, aligning instead with governance-centred accounts of regulated adoption [11], [16]. Finally, the generative-AI inversion customer awareness exceeding deployment extends Moharrak and Mogaji's observation of integration friction in regulated banks [16] and echoes economy-wide estimates of uneven AI diffusion [1], [24], [28], while the prominence of risk supports calls for interpretable, governable models [29], [30]. Where earlier Indian work treated perceived usefulness as the dominant predictor [12], [14], our coefficients place trust within one standard error of usefulness, indicating that credibility has caught up with convenience. The strong awareness–usage association (V = 0.412) is consistent with diffusion theory's

emphasis on observability [9] and with evidence that AI-driven initiatives gain acceptance only when customers can perceive them [17]. Conversely, our rejection of the group moderation contrasts with studies reporting segment-specific adoption paths [15], suggesting those differences were proxies for literacy rather than perspective; once literacy is controlled here, the perspective effect on the mechanism disappears. Together, the comparisons position this study as a synthesis that confirms acceptance theory, upgrades the role of trust, and relocates the adoption gap from behaviour to visibility.

## 6. Conclusion

This study set out to measure the adoption of artificial intelligence in banking and financial services from the two vantage points that matter most the customer who experiences AI and the organization that deploys it and to test whether those vantage points yield the same picture. They do not. Across 1,248 respondents, intention to adopt is statistically identical on both sides, but its architecture differs: professionals carry more perceived risk and more organizational readiness, customers carry more trust. Usage is stratified by age and by digital literacy but not by bank ownership, and it is gated by awareness, since aware customers use chatbot services at more than twice the rate of unaware ones. The regression results identify a clear priority order for managers usefulness, trust and ease of use first, risk mitigation as a guard-rail, readiness as an enabler while the rejected moderation hypothesis warns against segmenting strategy by respondent type. Explainability and plain-language disclosure should therefore be treated not as compliance overhead but as adoption infrastructure. For policy, the age gradient argues for assisted-digital channels for older cohorts, and the awareness wedge argues for plain-language disclosure of where AI already touches customer journeys. The study is limited by its cross-sectional design, self-reported usage and Indian sampling frame; longitudinal panels, behavioural log data and cross-country replication are the natural next steps. Adoption, in short, is less a question of whether banks install AI than of whether customers and organizations can be brought to see the same machine.

## 7. References


[1]. McKinsey Global Institute, "Notes from the AI frontier: Insights from hundreds of use cases," McKinsey & Company, San Francisco, CA, USA, Tech. Rep., Apr. 2018.

[2]. T. H. Davenport and R. Ronanki, "Artificial intelligence for the real world," Harv. Bus. Rev., vol. 96, no. 1, pp. 108–116, Jan.–Feb. 2018.

[3]. Y. K. Dwivedi et al., "Artificial Intelligence (AI): Multidisciplinary perspectives on emerging challenges, opportunities, and agenda for research, practice and policy," Int. J. Inf. Manage., vol. 57, Art. no. 101994, Apr. 2021.

[4]. F. D. Davis, "Perceived usefulness, perceived ease of use, and user acceptance of information technology," MIS Quart., vol. 13, no. 3, pp. 319–340, Sep. 1989.

[5]. V. Venkatesh, M. G. Morris, G. B. Davis, and F. D. Davis, "User acceptance of information technology: Toward a unified view," MIS Quart., vol. 27, no. 3, pp. 425–478, Sep. 2003.

[6]. D. Gefen, E. Karahanna, and D. W. Straub, "Trust and TAM in online shopping: An integrated model," MIS Quart., vol. 27, no. 1, pp. 51–90, Mar. 2003.

[7]. P. A. Pavlou, "Consumer acceptance of electronic commerce: Integrating trust and risk with the technology acceptance model," Int. J. Electron. Commerce, vol. 7, no. 3, pp. 101–134, 2003.

[8]. L. G. Tornatzky and M. Fleischer, The Processes of Technological Innovation. Lexington, MA, USA: Lexington Books, 1990.

[9]. E. M. Rogers, Diffusion of Innovations, 5th ed. New York, NY, USA: Free Press, 2003.

[10]. Reserve Bank of India, "Report on trend and progress of banking in India 2023–24," RBI, Mumbai, India, 2024.

[11]. Bank for International Settlements, "Artificial intelligence and machine learning in financial services: Market developments and financial stability implications," FSI Insights on Policy Implementation, no. 33, Nov. 2021.

[12]. R. Prasanth, "Demographic profile and adoption of artificial intelligence enabled banking services: An empirical study of bank customers in Coimbatore district," Int. J. Latest Technol. Eng. Manag. Appl. Sci., vol. 14, no. 12, pp. 106–112, Dec. 2025.

[13]. A. Raj and A. Puri, "Adoption of artificial intelligence in Indian banking industry," Innovative Res. Thoughts, vol. 10, no. 4, pp. 42–53, Nov. 2024.
[14]. S. Pareek, K. Khatri, D. K. Adhana, Rashmi, D. Sharma, and M. Srivastava, "An empirical study on the factors influencing the use of artificial intelligence in Indian financial services," Acad. Mark. Stud. J., vol. 29, no. S1, pp. 1–7, 2025.
[15]. V. K. Sharma and S. Singh, "Determinants of AI adoption in banking services in India: An empirical investigation," J. Inform. Educ. Res., vol. 4, no. 2, 2024.
[16]. M. Moharrak and E. Mogaji, "Generative AI in banking: Empirical insights on integration, challenges and opportunities in a regulated industry," Int. J. Bank Mark., vol. 43, no. 4, pp. 871–896, 2025.
[17]. S. Dharmadhikari, R. Kaur, and S. Khurjekar, "Assessing the customer adoption and perceptions for AI-driven sustainable initiatives in Indian banks," Environ. Soc. Psychol., vol. 9, no. 2, 2024.
[18]. J. F. Hair, W. C. Black, B. J. Babin, and R. E. Anderson, Multivariate Data Analysis, 8th ed. Boston, MA, USA: Cengage, 2019.
[19]. L. J. Cronbach, "Coefficient alpha and the internal structure of tests," Psychometrika, vol. 16, no. 3, pp. 297–334, Sep. 1951.
[20]. J. Cohen, Statistical Power Analysis for the Behavioral Sciences, 2nd ed. Hillsdale, NJ, USA: Lawrence Erlbaum, 1988.
[21]. M. I. Jordan and T. M. Mitchell, "Machine learning: Trends, perspectives, and prospects," Science, vol. 349, no. 6245, pp. 255–260, Jul. 2015.
[22]. I. Goodfellow, Y. Bengio, and A. Courville, Deep Learning. Cambridge, MA, USA: MIT Press, 2016.
[23]. A. Agrawal, J. Gans, and A. Goldfarb, Prediction Machines: The Simple Economics of Artificial Intelligence. Boston, MA, USA: Harvard Business Review Press, 2018.
[24]. C. B. Frey and M. A. Osborne, "The future of employment: How susceptible are jobs to computerisation?," Technol. Forecast. Soc. Change, vol. 114, pp. 254–280, Jan. 2017.
[25]. E. Brynjolfsson and A. McAfee, The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies. New York, NY, USA: W. W. Norton, 2014.
[26]. V. Venkatesh, J. Y. L. Thong, and X. Xu, "Consumer acceptance and use of information technology: Extending the unified theory of acceptance and use of technology," MIS Quart., vol. 36, no. 1, pp. 157–178, Mar. 2012.
[27]. NITI Aayog, "Responsible AI for all: Principles for responsible AI," Government of India, New Delhi, India, Tech. Rep., Feb. 2021.
[28]. International Monetary Fund, "Gen-AI: Artificial intelligence and the future of work," IMF Staff Discussion Note SDN/24/01, Washington, DC, USA, Jan. 2024.
[29]. C. Rudin, "Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead," Nature Mach. Intell., vol. 1, no. 5, pp. 206–215, May 2019.
[30]. A. Jobin, M. Ienca, and E. Vayena, "The global landscape of AI ethics guidelines," Nature Mach. Intell., vol. 1, no. 9, pp. 389–399, Sep. 2019.

### *Cite this Article:*

*Andani, S. G. (2026). Adoption of artificial intelligence in banking and financial services: An empirical study of customer and organizational perspectives. International Journal of Humanities, Commerce and Education, 2(10), 1–10.*

***Journal URL:*** *https://ijhce.com/*

***DOI****: https://doi.org/10.59828/ijhce.v2i10.141*

***Disclaimer/Publisher's Note****: The author(s) are solely responsible for the content and claims of this article. The publisher and editors assume no legal liability for any errors, copyright violations, or damages arising from its use.*