© International Academic Research Journal of Business and Management ISSN:2227-1287
Volume 14, Issue No:1, April 2026, Page no. 27- 36 DOI: https://doi.org/10.5281/zenodo.19911681

Consumer Perceptions of Fairness and Trust in AI-Driven Personalized Pricing

Dev Aadhisesh S K
MBA Student, Faculty of Management Studies,
CMS Business School,
JAIN (Deemed-to-be University), Bengaluru

Prof. Amita Gupta
Professor, Faculty of Management Studies,
CMS Business School,
JAIN (Deemed-to-be University), Bengaluru

Abstract

The rapid proliferation of artificial intelligence (AI) in digital commerce has given rise to personalized pricing, wherein algorithms dynamically adjust prices for individual consumers based on behavioral data, browsing history, geographic location, and purchase patterns. While this practice offers commercial efficiency for firms, it raises fundamental questions about consumer fairness and trust. This research investigates how consumers perceive the fairness of AI-driven personalized pricing and how such perceptions shape their trust in both AI systems and the companies that deploy them. Employing a mixed-methods primary research design, the study collected data from 155 respondents through a structured online survey comprising quantitative Likert-scale instruments alongside open-ended qualitative responses. The theoretical framework integrates Equity Theory, Information Asymmetry Theory, and the Technology Acceptance Model to interpret findings within established scholarly traditions. Key findings reveal that consumer acceptance of personalized pricing is not binary but conditional: the vast majority of respondents expressed willingness to accept such practices only when accompanied by transparency, data privacy guarantees, and opt-out mechanisms. The mean fairness score of 2.98 out of 5 and the mean AI trust score of 2.97 out of 5 both indicate low-to-moderate perceptions, while 86% of respondents agreed that algorithmic disclosure would significantly improve their trust. Regulatory demand was strong, with 4.22 out of 5 mean agreement on the need for government oversight. These findings carry significant implications for digital marketing strategy, AI ethics, and public policy formulation.

Keywords: Personalized Pricing, Algorithmic Fairness, Consumer Trust, AI Ethics, Digital Commerce, India

INTRODUCTION

The contemporary digital marketplace is undergoing a profound transformation driven by the integration of artificial intelligence into core commercial functions. Among the most commercially consequential and ethically contested applications of AI is personalized pricing — a practice whereby firms leverage algorithmic systems to set individualized prices for consumers based on an extensive array of personal data points. These include browsing behavior, purchase history, geographic location, device type, and even the time of day at which a consumer accesses a platform, resulting in a pricing environment that is neither uniform nor transparent.

Personalized pricing is not a new phenomenon in economic theory; price discrimination has long been studied as a mechanism by which firms extract consumer surplus. However, the advent of big data analytics and machine learning has elevated this practice from a theoretical possibility to an operational reality at unprecedented scale. Companies operating in e-commerce, ride-hailing, hospitality, travel, food delivery, and streaming services now routinely deploy AI-powered dynamic pricing engines that continuously recalibrate prices in real time.

For consumers, this creates an environment of informational asymmetry in which companies possess granular knowledge of individual willingness to pay, while consumers remain largely unaware that such mechanisms are at work. At its core, personalized pricing challenges the foundational market principle of price transparency. Research across consumer psychology and behavioral economics consistently shows that perceived price unfairness triggers negative emotional responses including anger, betrayal, and reduced willingness to engage with the offending firm.

The Indian digital commerce market provides a particularly rich context for this inquiry. India is among the world's fastest-growing digital economies, with over 850 million internet users and a rapidly expanding base of online shoppers. Platforms such as Amazon India, Flipkart, Swiggy, Zomato, MakeMyTrip, and Uber are deeply embedded in daily consumer life. Yet consumer awareness of AI-driven pricing mechanisms remains incomplete, and the regulatory landscape continues to evolve under the Digital Personal Data Protection Act 2023.

This study addresses the central research question: How do consumers perceive the fairness and trustworthiness of AI-driven personalized pricing, and what conditions mediate their acceptance of this practice? The study makes three principal contributions: first, it provides primary empirical evidence on consumer perceptions in the AI-driven personalized pricing domain; second, it identifies the conditions under which consumers willingly accept such practices; and third, it develops a conditional acceptance model applicable to AI decision-making contexts broadly.

REVIEW OF LITERATURE

The academic literature relevant to this study spans three intersecting bodies of knowledge: price fairness research, AI trust literature, and the specific emerging field of algorithmic pricing.

Price Fairness Perceptions

The foundational theoretical work on price fairness draws on Equity Theory (Adams, 1963), which holds that consumers evaluate fairness by comparing their own outcomes to reference standards. Campbell (1999) established that consumer perceptions of price unfairness are triggered not only by the price level itself but by inferences about the seller's motive, distinguishing between prices perceived as reflecting cost necessity and those attributed to opportunistic profit maximization. Haws and Bearden (2006) demonstrated that dynamic pricing depresses overall fairness perceptions even when consumers benefit from it, attributing this to the asymmetric information conditions that dynamic pricing creates. Xia, Monroe, and Cox (2004) developed a comprehensive conceptual framework for price fairness perceptions, emphasizing that procedural fairness — the fairness of the process by which prices are set — is at least as important as distributive fairness in determining consumer responses.

Trust in AI Systems

The literature on trust in AI systems establishes that algorithmic opacity is a significant driver of distrust. Consumers tend to find it harder to attribute intention and accountability to AI agents than to human agents, generating a specific form of distrust rooted in opacity rather than demonstrated malfeasance. McKnight, Choudhury, and Kacmar (2002) identified benevolence, integrity, and competence as the three core dimensions of technology trust, all of which are compromised when AI systems operate without transparency. Bart et al. (2005) confirmed that online trust is multidimensional and context-sensitive, with privacy and security perceptions playing especially strong roles in commercial digital contexts. Raghavan (2019) specifically examined algorithmic accountability, arguing that the technical complexity of AI systems creates structural barriers to the transparency that trust requires.

Algorithmic Personalization and Privacy

The privacy dimension of personalized pricing has been examined through the lens of the personalization-privacy paradox (Awad & Krishnan, 2006; Xu et al., 2011), wherein consumers desire the benefits of personalization while simultaneously being uncomfortable with the data sharing it requires. Chellappa and Sin (2005) found that the degree of personalization offered significantly moderates privacy concerns, with highly customized offers increasing discomfort proportionally. Martin, Borah, and Palmatier (2017) demonstrated empirically that consumer data privacy perceptions have measurable effects on both customer and firm performance, linking privacy trust to commercial loyalty. Zuiderveen Borgesius and Poort (2017) examined the regulatory dimensions of online price discrimination under EU data privacy law, establishing that algorithmic personalization implicates not merely commercial but fundamental rights concerns.

THEORETICAL FRAMEWORK

This study is grounded in three complementary theoretical frameworks that together provide a comprehensive lens for interpreting consumer perceptions of AI-driven personalized pricing.

Equity Theory

Originally proposed by Adams (1963) and extended to consumer contexts by Oliver and DeSarbo (1988), Equity Theory holds that individuals evaluate the fairness of exchanges by comparing their own input-outcome ratios to reference standards. In the context of personalized pricing, the consumer's relevant comparison is with other consumers paying different prices for identical goods. The theory predicts that unfavorable comparisons will generate perceptions of injustice and negative behavioral responses, while favorable comparisons may generate acceptance but not necessarily satisfaction.

Information Asymmetry Theory

Rooted in the work of Akerlof (1970) and Stiglitz (2001), Information Asymmetry Theory holds that market failures arise when one party to a transaction possesses significantly more information than the other. AI-driven personalized pricing creates a pronounced informational asymmetry: firms have access to detailed consumer data and sophisticated analytical tools, while consumers typically lack knowledge of how their prices are determined. This asymmetry undermines the conditions for voluntary, informed consent that underpin ethical commercial practice and creates conditions favorable to exploitation.

Technology Acceptance Model (TAM)

The Technology Acceptance Model provides a framework for understanding under what conditions consumers are willing to accept and use AI-driven systems. Perceived usefulness, perceived ease of use, and, in extended versions of the model, perceived trust and privacy risk all mediate technology acceptance decisions. The TAM framework predicts that consumers who perceive personal benefit from personalized pricing will be more accepting, while those who perceive privacy risk or opacity will resist adoption. Together, these three frameworks form an integrated theoretical basis for the study's hypotheses and analysis.

RESEARCH METHODOLOGY

Scope and Objectives

This study examines consumer perceptions of fairness and trust in AI-driven personalized pricing within the Indian digital commerce context. The scope is defined along four dimensions: thematic, geographical, temporal, and methodological. Thematically, the study focuses on three interrelated constructs: perceived price fairness, trust in AI systems, and conditional acceptance of personalized pricing. Geographically, the study is situated in the Indian digital market, capturing data from the Bengaluru metropolitan area. Temporally, the study captures consumer perceptions during 2024-2025.

The study pursues five research objectives: (1) to assess consumer awareness of AI-driven personalized pricing across major Indian digital platforms; (2) to measure consumer perceptions of fairness in personalized pricing; (3) to evaluate trust in AI pricing systems and identify trust-influencing factors; (4) to identify conditions under which consumers would accept personalized pricing; and (5) to examine how awareness of personalized pricing has affected consumer behavior.

Research Hypotheses

Based on the theoretical framework and review of literature, the following five hypotheses were advanced:

- H1 (Awareness and Fairness): Higher levels of consumer awareness of AI-driven personalized pricing are associated with lower perceptions of fairness.
- H2 (Fairness and Trust): Lower perceptions of pricing fairness are significantly associated with lower levels of trust in both AI pricing systems and the companies that deploy them.
- H3 (Transparency and Acceptance): Consumer acceptance of personalized pricing is significantly higher when transparency about pricing mechanisms is provided.
- H4 (Demographics and Attitude): Younger and more highly educated consumers demonstrate significantly more positive attitudes toward AI-driven personalized pricing than older or less educated consumers.
- H5 (Experience and Behavior): Consumers who have personally experienced price disparities are significantly more likely to have adopted defensive behaviors such as price comparison, use of incognito mode, or platform switching.

Research Design

This study employs a cross-sectional, mixed-methods primary research design. The quantitative component is built around a structured online survey comprising 15 questions organized across five thematically coherent sections: demographic profile, awareness and experience, perceived fairness, trust in AI systems and companies, and overall attitude and behavioral response. The survey instrument drew on validated scales from the academic literature: the Price Fairness Scale (Xia, Monroe, & Cox, 2004) and the Trust in Technology scale adapted from McKnight, Choudhury, and Kacmar (2002).

All Likert-scale items are scored on a five-point scale.

The qualitative component is embedded within the survey as an open-ended question asking respondents to describe conditions under which they would feel more comfortable with personalized pricing. Sampling was conducted through a combination of convenience and snowball approaches via the researcher's professional and academic networks and digital platforms. A screening question confirmed that all respondents were adults aged 18 or above who had engaged in online shopping at least occasionally.

Sample Profile

The study collected data from 155 complete respondents. The sample was predominantly drawn from younger age cohorts: the 25-34 age group was the most represented at 42.6%, followed by the 18-24 group at 23.2%, respondents aged 35-44 at 18.1%, the 45-54 group at 11.0%, and those aged 55 and above at 5.2%. In terms of educational attainment, the sample was highly educated, with postgraduate degree holders constituting the largest group at 32.3%, followed by undergraduates at 29.7%.

DATA ANALYSIS AND FINDINGS

Consumer Awareness and Experience

A significant majority of respondents demonstrated familiarity with personalized pricing. Sixty-seven respondents (43.2%) reported being fully aware, while 68 (43.9%) had heard of it but lacked detailed knowledge. Only 20 respondents (12.9%) reported no prior awareness. This distribution indicates that personalized pricing has entered mainstream consumer consciousness, though comprehension of its mechanics remains incomplete for a large share of the population.

Experience with price disparities was similarly mixed. Fifty respondents (32.3%) reported definitively having experienced paying a different price than others for identical products. A larger group of 66 respondents (42.6%) suspected this had occurred without being certain, reflecting the opacity of AI pricing systems. The combined total of those with direct or suspected experience amounts to 74.8% of the sample. Platform attribution was broadly distributed: ride-hailing applications were most frequently identified (71.0%), followed by food delivery (69.0%), e-commerce platforms (68.4%), travel and hotel booking (67.7%), and streaming services (65.8%).

Perceived Fairness Analysis

The overall fairness rating yielded a mean score of 2.98 out of 5 — marginally below the midpoint, indicating that respondents perceive AI-driven personalized pricing as slightly unfair. The distribution reveals 46 respondents (29.7%) rating fairness at 1 or 2, 68 (43.9%) rating it at 3 (neutral), and 41 (26.5%) rating it at 4 or 5. The modal response of neutrality reflects the conditional acceptance stance characterizing the sample.

The four-item fairness matrix revealed a more differentiated picture. The statement that companies should be transparent about their pricing mechanisms received the highest fairness-section mean of 4.26, indicating near-consensus on the importance of transparency. The belief that personalized pricing benefits companies more than consumers scored 3.84, indicating broad agreement with an asymmetric outcome narrative. The view that personalized pricing constitutes price discrimination scored 3.67, while conditional acceptance tied to personal benefit scored 3.43.

Trust in AI Pricing Systems

The overall trust score in AI pricing systems yielded a mean of 2.97 out of 5, virtually identical to the fairness mean, suggesting that the two constructs are closely aligned. Trust in companies to use consumer data responsibly was similarly low at 2.94, indicating that distrust extends from AI systems to the firms that deploy them.

However, the trust matrix reveals important nuances regarding conditions under which trust can be rebuilt. The statement that the respondent would trust a company more if it disclosed its pricing algorithm received the highest score in the entire survey at 4.30, indicating near-consensus on the trust-building value of transparency. The statement that government regulation is needed scored 4.22, reflecting strong demand for external oversight. The view that AI-driven pricing is too opaque to be trusted scored 3.86, indicating widespread concern about algorithmic opacity.

Conditions for Acceptance and Behavioral Responses

The conditions for acceptance identified by respondents were striking in their consistency. Transparency (clearly informed beforehand) was cited by 100 respondents (64.5%), loyalty discounts or personalized offers by 97 (62.6%), data privacy protection by 96 (61.9%), and opt-out options by 89 (57.4%). Only 21 respondents (13.5%) stated that personalized pricing was unacceptable under any circumstances. These four conditions form a coherent cluster describable as an informed consent framework.

Regarding behavioral responses, the most common reaction to discovering an overcharge was to feel annoyed but continue using the platform, selected by 61 respondents (39.4%). This finding is central to the conditional acceptance thesis: consumers are tolerating a practice they find unfair because switching costs are high. Thirty-three respondents (21.3%) indicated they would use technical countermeasures such as VPN or incognito mode, 25 (16.1%) would leave the platform, 22 (14.2%) would accept it as normal, and 14 (9.0%) would report or complain. Overall attitude distribution showed the modal response was neutrality (34.2%), followed by somewhat positive (26.5%) and somewhat negative (23.2%).

Hypothesis Testing Results

Each hypothesis was evaluated against the empirical evidence. H1 received partial support: respondents with no prior awareness showed the lowest fairness score (2.70), suggesting that initial awareness heightens unfairness perceptions, though fully aware respondents (3.00) showed similar scores to partially aware respondents (3.04). H2 was supported by the near-identical mean scores for overall fairness (2.98) and overall AI trust (2.97). H3 was strongly supported, with transparency being the most frequently cited acceptance condition (64.5%) and the highest single-item score for algorithmic disclosure (4.30). H4 was not supported: minimal variation in fairness perceptions across age groups (range 2.96 to 3.12) indicated that demographic differences are not reliable predictors of attitude. H5 was supported, with respondents who reported direct or suspected price disparity experience more likely to adopt defensive behaviors.

DISCUSSION

Theoretical Implications

The findings extend the Equity Theory framework to the algorithmic pricing context, demonstrating that consumers apply equity norms to AI-mediated transactions and that violations generate similar negative responses to those documented in conventional pricing contexts. The specific dynamics of algorithmic equity evaluation appear more sensitive to procedural fairness concerns than to outcome fairness alone, a refinement of established theory.

The findings contribute to Information Asymmetry Theory by demonstrating empirically that consumers are acutely aware of and sensitive to the informational advantages held by firms. The overwhelming demand for transparency and algorithmic disclosure suggests that reducing informational asymmetry is not merely a regulatory compliance issue but a fundamental condition for maintaining consumer trust. This finding has implications for the broader literature on information economics and market design in digital contexts.

The findings provide empirical support for and refinement of the Technology Acceptance Model in the AI pricing context. The TAM prediction that perceived usefulness mediates technology acceptance is partially confirmed, but the data also demonstrate that trust and perceived privacy risk are equally powerful determinants, suggesting that extended TAM models incorporating these constructs are better suited to the AI pricing context. The conditional acceptance model that emerges from this study offers a theoretical contribution in its own right, conceptualizing consumer acceptance as a contingent outcome determined by specific enabling conditions.

Managerial Implications

The primary managerial recommendation emerging from this study is that ethical personalization is not merely a compliance obligation but a source of competitive advantage. Firms should implement algorithmic transparency disclosure mechanisms providing plain-language explanations of how prices are determined. The finding that 4.30 out of 5 respondents agreed that algorithmic disclosure would increase their trust demonstrates the high return on investment such disclosures can generate.

Firms should also redesign their personalization value proposition to ensure that consumers subject to data-based pricing receive a clear and tangible benefit in return. Loyalty programs, personalized discounts, early access to sales, and free service upgrades are mechanisms through which firms can convert collected data into consumer value. The finding that loyalty discounts were cited as a condition for acceptance by 62.6% of respondents indicates strong consumer receptiveness to this approach. Finally, firms should invest in robust data privacy infrastructure and communicate this investment clearly, as 61.9% of respondents cited data privacy protection as a condition for acceptance.

LIMITATIONS AND SCOPE FOR FUTURE RESEARCH

This study is subject to several limitations. First, the sampling approach — relying on convenience and snowball sampling via digital networks — introduces selection bias toward digitally active and educated consumers, limiting generalizability to less digitally engaged segments. Second, the cross-sectional design captures consumer perceptions at a single point in time and cannot establish causal relationships. Third, self-reported data on awareness and experience may not accurately reflect actual consumer knowledge and behavior due to social desirability bias.

This study opens several productive avenues for future research. Longitudinal studies tracking changes in consumer awareness, fairness perceptions, and trust over time as AI pricing becomes more prevalent would provide valuable evidence on how attitudes evolve with experience and regulatory change. Experimental designs exposing participants to different transparency disclosure formats would enable causal assessment of the relationship between disclosure and trust. Cross-national comparative studies would illuminate how cultural factors mediate consumer responses across different regulatory and commercial contexts. Finally, sector-specific studies examining healthcare, insurance, and financial services pricing would extend the conditional acceptance model to high-stakes domains where the consequences of price discrimination are particularly significant.

CONCLUSION

This study provides empirical evidence that consumer acceptance of AI-driven personalized pricing is fundamentally conditional. The data reveal a landscape in which most consumers are aware of personalized pricing, find it somewhat unfair, distrust the AI systems that implement it, and yet continue to engage with the platforms that deploy it. This paradox of informed tolerance reflects the practical realities of a digital marketplace in which the perceived costs of switching platforms are high and the benefits of engagement are substantial.

The central conclusion is that the path to sustainable personalized pricing lies through transparency, reciprocal value creation, and consumer control. Firms that disclose their pricing mechanisms, offer consumers meaningful benefits in return for their data, and provide genuine opt-out rights will be better positioned to maintain the consumer trust that is ultimately the foundation of long-term commercial viability. Policymakers should develop regulatory frameworks that codify these conditions as minimum standards rather than voluntary aspirations. The future of AI-driven personalized pricing depends not on whether it is permitted but on whether it is practiced in a manner that respects the rights and interests of the consumers whose data makes it possible.

REFERENCES

Acquisti, A., & Varian, H. R. (2005). Conditioning prices on purchase history. Marketing Science, 24(3), 367–381. https://doi.org/10.1287/mksc.1040.0103

Awad, N. F., & Krishnan, M. S. (2006). The personalization privacy paradox: An empirical evaluation of information transparency and the willingness to be profiled online for personalization. MIS Quarterly, 30(1), 13–28.

Bart, Y., Shankar, V., Sultan, F., & Urban, G. L. (2005). Are the drivers and role of online trust the same for all web sites and consumers? Journal of Marketing, 69(4), 133–152.

Campbell, M. C. (1999). Perceptions of price unfairness: Antecedents and consequences. Journal of Marketing Research, 36(2), 187–199.

Chellappa, R. K., & Sin, R. G. (2005). Personalization versus privacy: An empirical examination of the online consumer's dilemma. Information Technology and Management, 6(2), 181–202.

Haws, K. L., & Bearden, W. O. (2006). Dynamic pricing and consumer fairness perceptions. Journal of Consumer Research, 33(3), 304–311.

Macha, M., Bhargave, R., & Chen, H. (2022). Robots set prices: Consumer perceptions and acceptance of AI-generated pricing recommendations. Journal of the Academy of Marketing Science, 50(4), 752–769.

Martin, K. D., Borah, A., & Palmatier, R. W. (2017). Data privacy: Effects on customer and firm performance. Journal of Marketing, 81(1), 36–58.

McKnight, D. H., Choudhury, V., & Kacmar, C. (2002). The impact of initial consumer trust on intentions to transact with a web site. Journal of Strategic Information Systems, 11(3–4), 297–323.

Morey, T., Forbath, T., & Schoop, A. (2015). Customer data: Designing for transparency and trust. Harvard Business Review, 93(5), 96–105.

Raghavan, M. (2019). Algorithmic accountability and the limits of transparency. Journal of Information Policy, 9, 303–340.

Shiller, B. R. (2020). Approximating purchase propensities and reservation prices from broad consumer tracking. International Economic Review, 61(2), 847–870.

Stiglitz, J. E. (2001). Information and the change in the paradigm in economics. American Economic Review, 92(3), 460–501.

Voigt, P., & Von dem Bussche, A. (2017). The EU general data protection regulation (GDPR): A practical guide. Springer.

Xia, L., Monroe, K. B., & Cox, J. L. (2004). The price is unfair! A conceptual framework of price fairness perceptions. Journal of Marketing, 68(4), 1–15.

Xu, H., Luo, X., Carroll, J. M., & Rosson, M. B. (2011). The personalization privacy paradox: An exploratory study of decision making process for location-aware marketing. Decision Support Systems, 51(1), 42–52.

Zuiderveen Borgesius, F. J., & Poort, J. (2017). Online price discrimination and EU data privacy law. Journal of Consumer Policy, 40(3), 347–366.

Table 1: Demographic Profile of Respondents (n=155)

Age Group	25–34 years	66 (42.6%)
	18–24 years	36 (23.2%)
	35–44 years	28 (18.1%)
	45–54 years	17 (11.0%)
	55+ years	8 (5.2%)
Education	Postgraduate	50 (32.3%)
	Undergraduate	46 (29.7%)
	Professional Degree	27 (17.4%)
	Doctoral	17 (11.0%)
Shopping Frequency	Several times/week	50 (32.3%)
	Once/week	49 (31.6%)
	Few times/month	36 (23.2%)

Table 2: Mean Scores for Perceived Fairness Statements (n=155)

Statement	Mean Score (1–5)
Companies should be transparent about pricing mechanisms	4.26
Pricing benefits companies more than consumers	3.84
Personalized pricing is a form of price discrimination	3.67
Acceptable if I receive a lower price	3.43
Overall fairness of personalized pricing	2.98

Table 3: Mean Scores for Trust in AI Pricing Statements (n=155)

Statement	Mean Score (1–5)
Would trust more if algorithm disclosed	4.30
Government regulation needed	4.22
AI pricing too opaque to trust	3.86
Overall trust in AI pricing systems	2.97
Trust companies to use data responsibly	2.94

Table 4: Summary of Hypothesis Testing Results

Hypothesis	Result
H1: Awareness → Lower Fairness Perception	Partially Supported
H2: Lower Fairness → Lower Trust	Supported
H3: Transparency → Higher Acceptance	Supported
H4: Younger/Educated → More Positive Attitude	Not Supported
H5: Direct Experience → Defensive Behavior	Supported