

## AI-Driven Personalization and Its Effect on Customer Trust, Satisfaction and Loyalty: a Conceptual and Evidence-Based Review with Reference to India

Dr. Kunal Soni HOD Accountancy<sup>1</sup>, Dr. Srinivasan Iyer<sup>2\*</sup>

<sup>1</sup>Maharashi Dayanand College of Arts, Science and Commerce, Maharashtra, India

<sup>2</sup>Associate Professor, CDOE, Manipal University Jaipur, Rajasthan, India

**\*Corresponding Author**

Dr. Srinivasan Iyer

Associate Professor, CDOE, Manipal University Jaipur, Rajasthan, India

**Cite This Paper as:** Dr. Kunal Soni HOD Accountancy, Dr. Srinivasan Iyer (2026) AI-Driven Personalization and Its Effect on Customer Trust, Satisfaction and Loyalty: a Conceptual and Evidence-Based Review with Reference to India The Journal Of African Development L, Vol.7, No.1, 158-168

### KEYWORDS

AI-driven personalization; customer trust; customer satisfaction; customer loyalty; privacy; personalization paradox; data protection; India

### ABSTRACT

Firms are turning more and more to artificial intelligence (AI) to adapt offers, content and service exchanges to each individual customer. Shoppers now expect this level of tailoring, but recent surveys point to weakening confidence in how companies handle data and AI. The present paper asks how AI-driven personalization shapes customer trust, satisfaction and loyalty. Drawing on a structured review of peer-reviewed research together with recent industry and regulatory material (McKinsey, Salesforce, Cisco and India's Digital Personal Data Protection framework), it builds a conceptual model containing seven propositions. The model holds that personalization reaches trust by two opposing paths: perceived relevance, which builds trust, and perceived privacy vulnerability, which erodes it. Reported statistics indicate that over 70% of consumers anticipate personalized interactions and roughly three in four are frustrated when they do not get them. Meanwhile, across Salesforce's successive studies, the proportion trusting companies to apply AI ethically dropped from 58% to 42%, while the proportion feeling they are treated as individuals climbed from 39% to 73%. Experimental and field findings suggest that the impact of personalization hinges on existing trust and on how openly data are gathered. The paper concludes that personalization lifts satisfaction and loyalty only if it is delivered through a design that protects trust, and it supplies a managerial framework, measurement scales and an agenda for primary research in India

## 1. INTRODUCTION

### 1.1 Background

Personalization is no longer an optional marketing enhancement; it is what people assume when they use digital services. Streaming services suggest the next programme, online marketplaces rearrange their landing pages shopper by shopper, and banking and food-delivery apps vary their offers by time of day, place and earlier activity. Most of these experiences are driven by machine-learning models that predict what each customer is likely to want and choose, frequently in a few milliseconds, what to present. Industry trend reviews for 2026 still rank AI-driven personalization among the core priorities of marketing, next to generative content, privacy-first data approaches and short-form video (GWI, 2025).

The business argument for personalization is well established. McKinsey (2021, 2023) finds that most consumers expect tailored interactions and that firms executing it well achieve markedly better marketing returns. But the technologies that enable tailoring rely on vast quantities of personal data and on automated inferences that customers cannot observe. Survey results from Salesforce (2025) and Cisco (2024) suggest customers are guarding their data more closely and feel less certain that firms will deploy AI responsibly. Personalization thus lies where two pressures meet: growing demand for relevance and growing unease over data use.

This tension is particularly timely for Indian marketers. The Government of India notified the Digital Personal Data Protection Rules, 2025 in November 2025, putting the Digital Personal Data Protection Act, 2023 into operational form (Hogan Lovells, 2025). When the core provisions become effective, firms will need a clear, purpose-specific basis for any personal data they use for personalization. How customers react to personalization under this regime is a practical concern for businesses and a worthwhile subject for researchers

## 1.2 Problem statement

Personalizing more does not automatically earn a firm more trust. Scholars have largely studied personalization, privacy concern, trust, satisfaction and loyalty in separate strands. Moreover, much of the best-known industry evidence appears as headline percentages with scant methodological information. What is needed is a single model that connects these strands and a judgement on what the available evidence can genuinely support.

## 1.3 Research objectives

Five objectives guide the study:

To define AI-driven personalization and link it with the established constructs of trust, satisfaction and loyalty.

To build a conceptual model showing how personalization can both increase and reduce customer trust.

To consolidate recent published evidence on consumer expectations, attitudes to privacy and trust in AI.

To evaluate how strongly the evidence supports each proposed relationship.

To draw out managerial implications and a research agenda, referring to India's data-protection regime.

## 1.4 Research questions

**RQ1:** How does AI-driven personalization affect customer trust?

**RQ2:** Via trust, how does personalization shape customer satisfaction and loyalty?

**RQ3:** Which conditions (transparency, control, prior trust and regulation) decide whether personalization strengthens or weakens these outcomes?

## 1.5 Significance of the study

For practitioners, the paper frames personalization as a decision about trust and not merely about targeting, and supplies measurable practices to go with it (Table 9). For students and scholars, it brings verifiable data from recent industry surveys together with classic theory, shows which links in the chain are well evidenced and which are not, and proposes scales for a primary study. For readers interested in policy, it illustrates how the principles of India's data-protection framework can be turned into design choices for personalization systems.

## 1.6 Structure of the paper

Section 2 covers the relevant concepts and earlier studies, and Section 3 sets out the conceptual framework with seven propositions. Section 4 explains the method and its boundaries. Section 5 analyses the evidence on expectations, attitudes to privacy, trust trends, behavioural findings and the Indian regulatory setting. Section 6 discusses findings and implications, Section 7 notes limitations and future research, and Section 8 concludes.

## 2. LITERATURE REVIEW

### 2.1 AI-driven personalization

Here, personalization refers to a firm tailoring content, offers, timing or service to a single customer on the basis of data about that customer. This is distinct from segmentation, where similar customers are treated as a group, and from customization, where the customer selects the settings. Describing it as AI-driven indicates that predictive or generative models, which learn from behaviour and refresh almost in real time, produce the tailoring instead of fixed rules. Bleier and Eisenbeiss (2015b) characterise the effectiveness of personalized advertising as the interaction of what is displayed, when, and where, a useful reminder that personalization offers several design levers.

Industry estimates put its value high. McKinsey (2023) pulls together evidence that personalization may cut customer acquisition costs by up to half, raise revenues by 5 to 15 percent and improve marketing return on investment by 10 to 30 percent. These should be treated as ranges seen across company cases and not as assured results, an issue revisited in Section 7.

### 2.2 Customer trust

Trust is commonly regarded as the basis of durable exchange relationships. Morgan and Hunt (1994) describe it as confidence in a partner's reliability and integrity, and show that together with commitment it supports relationship marketing. Gefen, Karahanna and Straub (2003) demonstrate that, online, trust reduces perceived risk and encourages purchase intention. Three facets matter most for personalization: competence (the firm can anticipate what I need), benevolence (the firm uses my data in my interest) and integrity (the firm keeps its data promises). Survey questions asking whether a company uses data in ways that benefit the customer correspond closely to the benevolence facet.

### 2.3 Customer satisfaction

Satisfaction is mostly explained through the expectation-disconfirmation model (Oliver, 1980). Customers hold expectations before an encounter, compare them with the performance they perceive, and are satisfied when performance equals or exceeds expectations. Personalization works through this mechanism in two ways. Relevant recommendations and well-timed service can raise perceived performance. At the same time, the spread of personalization lifts what customers take as the norm. The consumer research from McKinsey discussed in Section 5.1 indicates that not personalizing now frustrates most consumers. It is plausible, then, that personalization behaves partly as a hygiene factor,

penalised when missing but rewarded less than firms hope when present. This reading requires direct testing.

## 2.4 Customer loyalty

Oliver (1999) views loyalty as a strong commitment to buy a favoured product or service again in the future, even when circumstances and competing marketing might prompt switching. Researchers separate attitudinal loyalty (preference and advocacy) from behavioural loyalty (repeat buying). Satisfaction has long been tied to loyalty and profitability (Anderson, Fornell and Lehmann, 1994), and likelihood to recommend has become a popular practitioner metric through the Net Promoter Score (Reichheld, 2003). For personalization, loyalty is the result that ultimately justifies spending on data and technology, so its antecedents, trust and satisfaction, merit careful attention.

## 2.5 The personalization-privacy paradox

The main tension in the literature is that the data enabling relevance also produce vulnerability. Awad and Krishnan (2006) observed that customers who care more about information transparency are less willing to be profiled online for personalization. Aguirre, Mahr, Grewal, de Ruyter and Wetzels (2015) showed the paradox experimentally: personalized ads increased click-through intentions when information was gathered openly but not when it was gathered covertly, since covert gathering heightened customers' sense of vulnerability. They also found that cues building trust could cancel out the adverse effect. Bleier and Eisenbeiss (2015a) found that how well personalized banners worked depended on trust in the retailer, and that displaying a selection of previously viewed products outperformed displaying all of them, a result they attributed to unease about intrusiveness. Lambrecht and Tucker (2013) reported that very specific retargeted ads did not consistently beat generic ones and did best once customers' preferences were well formed. Table 2 summarises these studies.

**Table 1. Key constructs used in this paper**

| Construct                          | Working definition                                                                                                                                    | Theoretical anchor                            |
|------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------|
| AI-driven personalization          | Tailoring by the firm of content, offers, timing or service to one person, generated by predictive or generative models drawing on that person's data | Bleier & Eisenbeiss (2015b); McKinsey (2023)  |
| Perceived relevance and usefulness | Extent to which the customer sees the tailored output as matching their needs and saving them effort                                                  | Aguirre et al. (2015)                         |
| Perceived privacy vulnerability    | Unease and feeling of exposure when personalization reveals how much the firm knows or how it found out                                               | Aguirre et al. (2015); Awad & Krishnan (2006) |
| Customer trust                     | Belief in the firm's competence, benevolence and integrity, including in how it manages data and AI                                                   | Morgan & Hunt (1994); Gefen et al. (2003)     |
| Customer satisfaction              | Assessment that performance matches or surpasses expectations                                                                                         | Oliver (1980)                                 |
| Customer loyalty                   | Dedication to repurchasing and readiness to recommend even when alternatives exist                                                                    | Oliver (1999); Reichheld (2003)               |

**Table 2. Selected empirical studies on personalization, privacy and trust**

| Study                     | Method                                | Main finding relevant to this paper                                                                                                                                       |
|---------------------------|---------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Awad & Krishnan (2006)    | Empirical survey evaluation           | Customers who attach more value to information transparency are less ready to be profiled online for personalization.                                                     |
| Lambrecht & Tucker (2013) | Field data from retargeted ads        | Highly specific retargeted ads do not always beat generic ones; they perform better after preferences have formed.                                                        |
| Aguirre et al. (2015)     | One field study and three experiments | Personalization improves response when collection is overt; covert collection heightens vulnerability and reduces click-through; trust-building cues can counteract this. |



| Study                       | Method                                | Main finding relevant to this paper                                                                                                                                                                    |
|-----------------------------|---------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Bleier & Eisenbeiss (2015a) | Field data and laboratory experiments | Acceptance of personalized banners rests on trust in the retailer; one previously viewed item outperformed all of them; in the field study click-through was 27% higher for the more trusted retailer. |

## 2.6 Research gap

This review points to three gaps. First, the classic studies address personalized advertising and retargeting, and little work examines AI-driven personalization over the full customer journey, including generative content and conversational agents. Second, trust, satisfaction and loyalty are seldom modelled alongside the two opposing perceptions, relevance and vulnerability, that personalization generates. Third, evidence from India, where regulation is in flux, is still scarce. The model in Section 3 answers the first two gaps, while Section 5.5 and Section 7 take up the third.

## 3. CONCEPTUAL FRAMEWORK AND PROPOSITIONS

Figure 1 presents the proposed model. AI-driven personalization is expected to affect trust along two parallel paths: a positive path via perceived relevance and usefulness, and a negative path via perceived privacy vulnerability. Trust then passes the combined effect on to satisfaction and loyalty, and also links directly to loyalty.




**Figure 1. Conceptual model linking AI-driven personalization, trust, satisfaction and loyalty. Moderators (not shown): transparency of data collection, customer control and consent, prior brand trust, awareness of privacy law, and age.**

Seven propositions follow from the model:

- P1.** Better AI-driven personalization raises perceived relevance and usefulness.
- P2.** Greater personalization intensity, particularly with covert data collection, raises perceived privacy vulnerability.
- P3.** Perceived privacy vulnerability lowers customer trust.
- P4.** Perceived relevance and usefulness raise customer trust.
- P5.** Customer trust raises customer satisfaction.
- P6.** Customer satisfaction raises customer loyalty.
- P7.** Customer trust also raises customer loyalty directly.

Transparency, control, prior brand trust, awareness of privacy law and age are expected to moderate P2 and P3. For instance, open collection and simple controls should weaken the path from personalization to vulnerability, and customers who already trust the brand should accept more personalization before they feel exposed.

## 4. RESEARCH METHODOLOGY

### 4.1 Design

This is a structured secondary-data review combined with conceptual model building; it presents no new survey or experiment. It brings together published empirical studies, large industry surveys and legal texts, and applies simple

calculations (percentage-point and relative changes) to compare figures between survey editions. The design fits the aim of consolidating existing knowledge and identifying what should be tested next.

#### 4.2 Sources and selection

Three criteria governed source selection: relevance to personalization, privacy, trust in AI or customer outcomes; recency (industry and legal sources dated 2021 to 2025, academic sources by original publication date); and traceability to a named publisher or institution. Peer-reviewed work was favoured for causal claims, while industry surveys served for descriptive trends. The searches took place in October 2026. Table 3 lists the main evidence sources.

**Table 3. Principal evidence sources**

| Source                                             | Type                                           | Used for                                                                   |
|----------------------------------------------------|------------------------------------------------|----------------------------------------------------------------------------|
| McKinsey & Company (2021, 2023)                    | Consultancy research                           | Consumer expectations about personalization; business results reported     |
| Cisco (2024)                                       | Yearly worldwide consumer survey               | Privacy awareness, use of generative AI, behaviour linked to data trust    |
| Salesforce (2025)                                  | Large customer survey across several countries | Trends in trust in AI, in data use and in feeling treated as an individual |
| Aguirre et al. (2015); Bleier & Eisenbeiss (2015a) | Peer-reviewed, Journal of Retailing            | Experimental and field findings on trust and personalization               |
| Government of India (2023); MeitY (2025)           | Statute and rules                              | Regulatory background for personalization in India                         |
| Hogan Lovells (2025); Social Samosa (2025)         | Legal and trade commentary                     | Reading of the Rules; context of Indian advertising                        |

#### 4.3 Limitations of the method

Four limitations should be borne in mind. First, vendors and consultancies with commercial interests produce the industry surveys, and the web pages and press releases summarising them frequently leave out sample sizes and question wording. Second, editions differ in year, sample and country mix, so comparisons show the direction of change and not its exact size. Third, stated attitudes and actual behaviour do not always align; Cisco's data, discussed below, reveal a gap between concern and conduct. Fourth, the model is conceptual and its propositions are not statistically tested here. Figures cited in this paper are paraphrased from the sources and should be verified against the originals before reuse.

### 5. EVIDENCE AND ANALYSIS

#### 5.1 Consumer expectations and reported business outcomes

McKinsey's consumer research finds that 71 percent of consumers expect companies to provide personalized interactions and 76 percent become frustrated when that does not happen (McKinsey, 2021). The same research reports that fast-growing firms generate 40 percent more of their revenue from personalization than slower-growing peers, and, as summarised by the CMO Council (2021), that 76 percent of consumers are more inclined to consider purchasing from brands that personalize and 78 percent are more inclined to recommend them to friends and family. Table 4 gathers these indicators.

**Table 4. Expectations of personalization and reported outcomes**

| Indicator                                                              | Reported figure | Source                                          |
|------------------------------------------------------------------------|-----------------|-------------------------------------------------|
| Share of consumers expecting personalized interactions                 | 71%             | McKinsey (2021)                                 |
| Share of consumers frustrated when personalization is absent           | 76%             | McKinsey (2021)                                 |
| Revenue from personalization, fast-growing versus slower-growing firms | 40% more        | McKinsey (2021)                                 |
| Consumers more inclined to consider brands that personalize            | 76%             | McKinsey (2021), as cited by CMO Council (2021) |



| Indicator                                                  | Reported figure | Source                                          |
|------------------------------------------------------------|-----------------|-------------------------------------------------|
| Consumers more inclined to recommend to friends and family | 78%             | McKinsey (2021), as cited by CMO Council (2021) |
| Possible fall in customer acquisition cost                 | Up to 50%       | McKinsey (2023)                                 |
| Possible revenue uplift                                    | 5% to 15%       | McKinsey (2023)                                 |
| Possible gain in marketing return on investment            | 10% to 30%      | McKinsey (2023)                                 |

Note: Figures appear as given by the cited sources; outcomes are ranges across company cases and not guaranteed.

These data can be read in two ways. Optimistically, personalization is both expected and rewarded, which supports P1 and the later links to loyalty, referral intent especially. Cautiously, the 76 percent frustration figure indicates that personalization has become a minimum standard, so firms cannot use it to stand out unless they also get the trust conditions right. Because the evidence is largely self-reported, P1 receives only a moderate strength rating in Table 8.

## 5.2 Privacy awareness, generative AI and data-trust behaviour

Cisco's 2024 Consumer Privacy Survey shows the other side of the paradox. It reports that 53 percent of consumers know about privacy laws, 63 percent think AI can improve their lives, and 59 percent say strong privacy laws make them more comfortable sharing information in AI applications (Cisco, 2024). Regular use of generative AI almost doubled in a year, from 12 to 23 percent. Even so, 30 percent of generative AI users type personal or confidential details into these tools, although 84 percent worry about such data becoming public, a 54-percentage-point difference between concern and behaviour.

The survey also connects trust with buying decisions. Over three-quarters of consumers say they would not buy from an organisation whose handling of their data they do not trust, and 49 percent of 25-to-34-year-olds say they have changed companies or providers because of data policies or sharing practices, against 18 percent of those aged 75 and above. The younger group's rate is roughly 2.7 times that of the older one. Since younger consumers also use digital services most heavily, this age pattern is significant for firms aiming personalization at them. Table 5 summarises these indicators.

**Table 5. Selected indicators from the Cisco 2024 Consumer Privacy Survey**

| Indicator                                                                                      | Reported figure |
|------------------------------------------------------------------------------------------------|-----------------|
| Consumers who know about privacy laws                                                          | 53%             |
| Consumers who think AI can improve their lives                                                 | 63%             |
| Consumers for whom strong privacy laws increase comfort sharing information in AI applications | 59%             |
| Regular generative AI use (prior year in brackets)                                             | 23% (12%)       |
| Generative AI users who input personal or confidential data                                    | 30%             |
| Consumers worried that data entered into generative AI could become public                     | 84%             |
| Consumers who hold businesses responsible for ethical AI use                                   | 78%             |
| Consumers unwilling to buy from an organisation whose data practices they distrust             | More than 75%   |
| Changed provider because of data practices: ages 25-34 versus 75+                              | 49% vs 18%      |

Source: Cisco (2024). The difference between concern (84%) and behaviour (30%) is the author's calculation.

## 5.3 Trends in trust: the personalization-trust divergence

Salesforce's State of the AI Connected Customer study (seventh edition, with over 16,000 consumers and business buyers worldwide) gives the clearest evidence of a gap between how personal service feels and how far customers trust the system behind it. The proportion of customers who say companies treat them as individuals and not as numbers climbed from 39 percent in 2023 to 73 percent. Over the same time, those trusting businesses to use AI ethically fell from 58 percent in 2023 to 42 percent, and those believing companies use their data in ways that benefit them dropped from 60 percent in 2022 to 49 percent (Salesforce, 2025). Table 6 and Figure 2 show the comparison.

**Table 6. Change in selected customer perceptions (Salesforce)**

| Indicator                                                | Earlier edition | 2025 edition | Change (pp) | Relative change |
|----------------------------------------------------------|-----------------|--------------|-------------|-----------------|
| Trust in businesses to use AI ethically                  | 58% (2023)      | 42%          | -16         | -28%            |
| Belief that companies use data to the customer's benefit | 60% (2022)      | 49%          | -11         | -18%            |
| Feel treated as an individual rather than a number       | 39% (2023)      | 73%          | +34         | +87%            |

Source: Salesforce (2025); changes computed by the author. Because samples differ between editions, the comparison indicates direction and not exact magnitude.




**Figure 2. Customer views on AI, data use and being treated as an individual: earlier edition compared with 2025 (Salesforce, 2025).**

The report further notes that 71 percent of customers are growing more protective of their personal information, 64 percent think companies are careless with customer data, 61 percent believe advances in AI make corporate trustworthiness even more important, and 72 percent consider it important to know whether they are dealing with an AI agent (Salesforce, 2025). Comfort also varies by task: 38 percent of customers accept agents creating personalized content, but just 17 percent accept an agent making financial decisions on their behalf.

Taken together, these figures depict a market where personalization is increasingly felt to be present and effective while confidence in the motives behind it declines. This is consistent with the model's two-path logic: relevance is increasing (the positive path), but so are perceived vulnerability and doubts about benevolence (the negative path). The pattern does not prove the mechanism, as the surveys were not built to test it, yet it fits better than a simple view that more personalization is always better.

#### 5.4 Behavioural and experimental evidence

Experimental and field studies support stronger causal claims than opinion surveys. Aguirre et al. (2015) show that personalization raises response only when information is collected overtly; with covert collection, customers feel more vulnerable and click-through declines. Their field study on Facebook and secondary campaign data also showed steep falls in click-through once customers learned their data had been gathered without consent. These findings support P2 and P3 and show that the moderating role of transparency proposed in Section 3 has empirical grounding.

Bleier and Eisenbeiss (2015a) add that an identical personalized ad performs differently depending on its sender. In their

field study, click-through was 27 percent higher for the more trusted retailer, and privacy concerns and perceived intrusiveness prompted negative reactions, particularly toward retailers that had not built trust through earlier positive experiences. Their result that a single previously viewed product beat a display of every viewed product points to a practical design rule: calibrate how much the system reveals it knows. Lambrecht and Tucker (2013) point the same way, showing that specificity pays off only when the customer is ready for it.

### 5.5 The Indian regulatory context

India's Digital Personal Data Protection Act was passed on 11 August 2023, and the implementing Rules were notified in November 2025 (Hogan Lovells, 2025). The Act covers digital personal data and is built on seven principles: consent and transparency, purpose limitation, data minimisation, accuracy, storage limitation, security safeguards and accountability; it also demands verifiable consent before children's data is processed, subject to limited exemptions (Indian Infrastructure, 2025). Under the Act, consent has to be free, specific, informed, unconditional and unambiguous, given through a clear affirmative action.

Trade commentary observes that the main provisions on notices, consent, processing and fiduciary duties are due to take effect 18 months after notification, and that purpose-specific, informed consent will make mass collection of personal data for unspecified later targeting difficult to defend (Social Samosa, 2025). The same commentary notes that programmatic buying made up 42 percent (about Rs 20,686 crore) of India's digital media industry at the end of 2024, citing Dentsu's digital report, which shows how much ad spending relies on data flows that the new regime will reshape. Table 7 turns each principle into an implication for AI-driven personalization. The table reflects the author's interpretation and is not legal advice.

**Table 7. DPDP principles and implications for AI-driven personalization**

| Principle                       | What it requires (summary)                                                        | Implication for personalization                                                                          |
|---------------------------------|-----------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------|
| Consent and transparency        | Informed, specific consent given by clear affirmative action, with a clear notice | Say in plain language which data drive which personalized experience; choose opt-in over buried defaults |
| Purpose limitation              | Data may be used only for the consented purpose                                   | Do not reuse browsing or purchase data for unrelated profiling without fresh consent                     |
| Data minimisation               | Gather only what the purpose requires                                             | Prefer a few high-value signals, including data customers offer voluntarily, to exhaustive profiles      |
| Accuracy and storage limitation | Keep data accurate and do not hold it longer than necessary                       | Define retention periods for model inputs and update outdated preferences                                |
| Security safeguards             | Reasonable safeguards and notification of breaches                                | Regard personalization datasets as prime targets; a breach harms trust directly                          |
| Accountability                  | The data fiduciary is answerable, and a Data Protection Board handles complaints  | Chart data flows, record the legal basis for every marketing use and review vendor contracts             |
| Children's data                 | Verifiable consent is required before processing                                  | Restrict or switch off behavioural personalization for minors unless consent has been verified           |

Source: Compiled from Indian Infrastructure (2025), Hogan Lovells (2025) and Social Samosa (2025); the implications column reflects the author's interpretation.

## 6. DISCUSSION

### 6.1 Assessing the propositions

Table 8 rates the evidence behind each proposition. The negative path (P2 and P3) and the loyalty links (P6 and P7) are best supported, resting on experimental studies and long-standing theory. P1 is supported, but mainly by self-reported survey data. P4 and P5 have the weakest direct evidence in an AI setting: existing studies show that trust conditions responses to personalization, but few test whether perceived relevance itself builds trust or whether trust raises satisfaction with AI-personalized services. These links are therefore priorities for primary research.

**Table 8. Evidence assessment for the propositions**

| Proposition                                                                 | Principal evidence                                                                                | Strength                        |
|-----------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------|---------------------------------|
| P1: Better personalization raises perceived relevance                       | McKinsey (2021) data on expectation and frustration; Bleier & Eisenbeiss (2015a) on ad design     | Moderate (mostly self-reported) |
| P2: Intensive personalization, particularly covert, increases vulnerability | Aguirre et al. (2015); Awad & Krishnan (2006)                                                     | Strong (experimental)           |
| P3: Vulnerability lowers trust                                              | Aguirre et al. (2015); Cisco (2024) on refusal to purchase and switching; Salesforce (2025) trend | Moderate to strong              |
| P4: Relevance raises trust                                                  | Indirect: Bleier & Eisenbeiss (2015a); Salesforce (2025)                                          | Limited; needs direct testing   |
| P5: Trust raises satisfaction                                               | Theory (Morgan & Hunt, 1994; Gefen et al., 2003); AI-specific testing is scarce                   | Limited; needs direct testing   |
| P6: Satisfaction raises loyalty                                             | Anderson et al. (1994); Oliver (1999); McKinsey (2021) on referral intent                         | Strong (established)            |
| P7: Trust directly increases loyalty                                        | Morgan & Hunt (1994); Cisco (2024) on switching                                                   | Strong                          |

## 6.2 Why more personalization does not automatically mean more trust

The synthesis indicates that personalization creates customer value only when the customer understands and accepts the exchange behind it. When the exchange is open, relevance is perceived as service; when it is covert or excessive, the same relevance is perceived as surveillance. This accounts for how McKinsey can report that most consumers expect personalization while Salesforce reports declining trust in data and AI use: both can hold together. It also explains the weight of prior trust in Bleier and Eisenbeiss (2015a). A trusted firm has more latitude to personalize since customers presume benevolent intent, while a firm lacking that credit must first earn it.

Measurement is a further implication. If personalization is partly a hygiene factor (Section 2.3), firms tracking only click-through or conversion will overlook the slower decline in trust that shows up in satisfaction, switching and referral measures. Trust and loyalty indicators such as CSAT, NPS and the share of customers opting in to data use should therefore be reported next to short-term campaign metrics.

## 6.3 Managerial implications

Table 9 converts the discussion into seven practices, each paired with an indicator a team can monitor. Several correspond directly to the findings reviewed earlier: open collection (Aguirre et al., 2015), calibrated display of past behaviour (Bleier and Eisenbeiss, 2015a), disclosure of AI use with clear escalation to humans (Salesforce, 2025), and consent and minimisation (Indian data-protection principles).

**Table 9. A trust-preserving framework for AI-driven personalization**

| Practice          | What to do                                                                         | Indicator to track                                   |
|-------------------|------------------------------------------------------------------------------------|------------------------------------------------------|
| 1. Collect openly | Tell customers what is collected and why when collecting it; avoid hidden tracking | Opt-in rate; privacy complaints                      |
| 2. Give control   | Offer a preference centre and simple opt-out; apply changes promptly               | Preference-centre use; opt-out rate                  |
| 3. Minimise data  | Favour volunteered data and contextual signals over exhaustive profiling           | Proportion of personalization based on declared data |
| 4. Calibrate      | Reveal only part of what you know; do not display every past action                | Click-through; a survey item on feeling watched      |
| 5. Explain        | Show a visible reason for every recommendation or offer                            | CSAT on recommendations                              |

| Practice                        | What to do                                                                                              | Indicator to track                              |
|---------------------------------|---------------------------------------------------------------------------------------------------------|-------------------------------------------------|
| 6. Disclose AI and offer humans | Mark AI interactions clearly and give an easy route to a human                                          | Escalation resolution CSAT; complaint rate      |
| 7. Measure what matters         | Compare personalized and non-personalized experiences through holdout groups on CSAT, NPS and retention | CSAT, NPS and repeat purchase by exposure group |

#### 6.4 Theoretical contribution

The paper's contribution lies in joining three bodies of work that are normally kept apart: the expectation-disconfirmation theory of satisfaction, the trust-commitment account of loyalty and the personalization-privacy paradox. Its central conceptual claim is that trust is where they converge. Personalization sends trust two opposing signals, and the net outcome depends on conditions firms control: transparency, control and the degree of data intrusion. Researchers thus receive a testable structure, and the assumption that personalization effects are uniformly positive is avoided.

#### 7. LIMITATIONS AND FUTURE RESEARCH

The limitations listed in Section 4.3 apply here. Beyond them, the propositions are not statistically tested, evidence on generative AI and conversational agents is recent and still evolving, and India's regulatory regime is in transition, so its behavioural effects cannot yet be seen. Future research should therefore run a primary study with Indian online shoppers. A cross-sectional survey analysed through structural equation modelling could test P1 to P7 and the moderators. An experiment varying whether data collection is disclosed and whether a control option is offered would test the causal claims in P2 and P3. Longitudinal data could show whether lost trust recovers. Table 10 proposes indicative measures, worded for this study, that can be adapted for such work.

**Table 10. Indicative measurement scales for a primary study**

| Construct                          | Indicative item (author's wording)                                            | Response format                                             |
|------------------------------------|-------------------------------------------------------------------------------|-------------------------------------------------------------|
| Perceived personalization quality  | The recommendations this app gives me fit my interests.                       | 5-point Likert                                              |
| Perceived relevance and usefulness | The offers shown to me here are useful for my needs.                          | 5-point Likert                                              |
| Perceived privacy vulnerability    | It makes me uncomfortable how much this company appears to know about me.     | 5-point Likert (adapt from Aguirre et al., 2015)            |
| Customer trust                     | I trust this company to handle my data responsibly.                           | 5-point Likert                                              |
| Customer satisfaction              | On the whole, I am satisfied with my experience of this company.              | 5-point scale (CSAT style)                                  |
| Customer loyalty                   | I plan to continue buying from this company and would recommend it to others. | 5-point Likert; 0-10 recommendation scale (Reichheld, 2003) |

Note: Sample size should follow a power analysis; the model is compact enough to be estimated with a few hundred respondents.

#### 8. CONCLUSION

AI-driven personalization has become an expectation rather than a bonus: around seven in ten consumers anticipate it and roughly three-quarters are frustrated when it is absent. The evidence reviewed here shows, though, that providing it is insufficient. Customers increasingly feel treated as individuals, yet fewer trust firms to use AI ethically or to use data for their benefit, and many say they will leave organisations whose data handling they distrust. Experimental studies explain why: personalization boosts response when data collection is open and trusted, and undermines it when collection seems covert or intrusive.

The proposed model puts trust at the centre of the chain from personalization to satisfaction and loyalty, and names the practices that protect it: open collection, customer control, data minimisation, calibrated display, explanation, disclosure of AI with access to humans, and measurement beyond clicks. For Indian firms, the Digital Personal Data Protection framework converts several of these from good practice into legal expectation. The paper's main contribution is a clear,

testable structure; the next step is to test the weaker links, especially P4 and P5, with primary data.

## REFERENCE

1. Aguirre, E., Mahr, D., Grewal, D., de Ruyter, K., & Wetzels, M. (2015). Unraveling the personalization paradox: The effect of information collection and trust-building strategies on online advertisement effectiveness. *Journal of Retailing*, 91(1), 34-49. <https://doi.org/10.1016/j.jretai.2014.09.005>
2. Anderson, E. W., Fornell, C., & Lehmann, D. R. (1994). Customer satisfaction, market share, and profitability: Findings from Sweden. *Journal of Marketing*, 58(3), 53-66.
3. 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.
4. Bleier, A., & Eisenbeiss, M. (2015a). The importance of trust for personalized online advertising. *Journal of Retailing*, 91(3), 390-409.
5. Bleier, A., & Eisenbeiss, M. (2015b). Personalized online advertising effectiveness: The interplay of what, when, and where. *Marketing Science*, 34(5), 669-688.
6. Cisco. (2024, October 30). New Cisco survey shows strong relationship between privacy awareness and trust in AI [Press release; Cisco 2024 Consumer Privacy Survey].
7. <https://newsroom.cisco.com/c/r/newsroom/en/us/a/y2024/m10/cisco-survey-shows-strong-relationship-between-privacy-awareness-and-trust-in-ai.html>
8. CMO Council. (2021). Business is personal [CMO blog]. <https://redemo.cmocouncil.org/expert-views/cmo-blog/business-is-personal>
9. Gefen, D., Karahanna, E., & Straub, D. W. (2003). Trust and TAM in online shopping: An integrated model. *MIS Quarterly*, 27(1), 51-90.
10. Government of India. (2023). The Digital Personal Data Protection Act, 2023 (Act No. 22 of 2023).
11. GWI. (2025, December). Connecting the dots: 5 major marketing trends for 2026 you can't afford to overlook. <https://www.thinkdigital.travel/research-directory/connecting-the-dots-5-major-marketing-trends-for-2026-you-can-t-afford-to-overlook>
12. Hogan Lovells. (2025). India's Digital Personal Data Protection Act 2023 brought into force.
13. <https://www.hoganlovells.com/en/publications/indias-digital-personal-data-protection-act-2023-brought-into-force->
14. Indian Infrastructure. (2025, November 18). Government notifies Digital Personal Data Protection Rules 2025. <https://indianinfrastructure.com/2025/11/18/government-notifies-digital-personal-data-protection-rules-2025/>
15. Lambrecht, A., & Tucker, C. (2013). When does retargeting work? Information specificity in online advertising. *Journal of Marketing Research*, 50(5), 561-576.
16. McKinsey & Company. (2021, November). The value of getting personalization right, or wrong, is multiplying. McKinsey & Company.
17. McKinsey & Company. (2023, May). What is personalization? McKinsey Explainers.
18. [https://mckinsey.com/~media/mckinsey/featured\\_insights/mckinsey\\_explainers/what\\_is\\_personalization/what-is-personalization.pdf](https://mckinsey.com/~media/mckinsey/featured_insights/mckinsey_explainers/what_is_personalization/what-is-personalization.pdf)
19. Ministry of Electronics and Information Technology. (2025). Digital Personal Data Protection Rules, 2025 [Notified November 2025]. Government of India.
20. Morgan, R. M., & Hunt, S. D. (1994). The commitment-trust theory of relationship marketing. *Journal of Marketing*, 58(3), 20-38.
21. Oliver, R. L. (1980). A cognitive model of the antecedents and consequences of satisfaction decisions. *Journal of Marketing Research*, 17(4), 460-469.
22. Oliver, R. L. (1999). Whence consumer loyalty? *Journal of Marketing*, 63(Special Issue), 33-44.
23. Reichheld, F. F. (2003). The one number you need to grow. *Harvard Business Review*, 81(12), 46-54.
24. Salesforce. (2025). State of the AI connected customer (7th ed.). <https://salesforce.com/uk/resources/research-reports/state-of-the-connected-customer>
25. Social Samosa. (2025, November). DPDP to push a major reset for ad targeting and personalisation. <https://www.socialsamosa.com/experts-speak/dpdp-to-push-a-major-reset-for-ad-targeting-and-personalisation-10785467>