

REFRAMING AI-DRIVEN PERSONALIZATION IN RETAIL: A CONCEPTUAL SYNTHESIS OF TRUST AND DATA PRIVACY AS MEDIATING CONSTRUCTS**Djazabel Petrov**

ORCID ID - 0009-0004-7455-0360

Student, University of East London

<https://orcid.org/0009-0004-7455-0360>**ABSTRACT**

This study conceptually examines the relationship between AI-driven personalization and customer loyalty in retail, with a particular focus on the mediating roles of trust and data privacy perceptions. Although personalization is widely presented as a driver of customer experience and retention, the literature does not support a consistently direct relationship between personalized retail interactions and loyalty outcomes. Instead, prior work suggests that customer responses depend on how personalization is interpreted, especially in relation to trust and the perceived handling of personal data (Awad and Krishnan, 2006; Martin and Murphy, 2017; Jarrar, 2026).

This paper adopts a structured literature synthesis approach to integrate foundational and recent research on personalization, privacy, trust, and customer response in retail and adjacent digital environments. The analysis suggests that AI-driven personalization does not directly create loyalty. Rather, its effectiveness depends on whether customers perceive the personalization process as acceptable, transparent, and aligned with responsible data practices. Data privacy perception appears to function as an early evaluative filter, while trust acts as the relational mechanism through which accepted personalization can translate into repeat engagement and loyalty.

The study contributes by reframing personalization as a conditional rather than direct driver of loyalty. It also proposes a process-based conceptual framework that distinguishes privacy evaluation from trust formation and explains how positive and negative customer responses can emerge from the same personalization practice. From a managerial perspective, the paper highlights that the competitive value of personalization no longer lies in technical sophistication alone, but in the ability of retailers to balance relevance with transparency, restraint, and trustworthiness.

Keywords:

AI-driven personalization, customer loyalty, trust, data privacy, retail, conceptual framework

1. INTRODUCTION

Artificial intelligence has become a central part of contemporary retail strategy. Retailers increasingly use AI to personalize recommendations, offers, interfaces, and communications based on customer data, browsing behavior, purchase history, and predictive analytics. This shift is commonly justified on the basis that more relevant interactions should improve customer experience and eventually strengthen loyalty (Grewal, Roggeveen and Nordfält, 2017; Huang and Rust, 2018). Yet the real picture is messier.

Customers do not respond to personalization in a uniformly positive way. Some perceive it as useful, convenient, and relevant. Others perceive it as intrusive, manipulative, or vaguely unsettling. This tension has long been recognized in the literature through what Awad and Krishnan (2006) describe as the personalization–privacy paradox: customers value relevance, but they also become uncomfortable when relevance implies extensive profiling or opaque data usage.

This problem has become more important, not less, in AI-enabled environments. AI makes personalization more scalable, faster, and more predictive. It also makes it easier for firms to infer preferences, automate segmentation, and intervene in customer decision processes in ways that are not always visible to the customer. As a result, the issue is no longer simply whether personalization works, but **under what conditions** it works and **through which mechanisms** it influences loyalty.

Much of the literature still treats personalization as though it has a mostly direct positive effect on customer outcomes. However, both foundational and recent work suggest that the relationship is mediated by psychological and evaluative processes, especially trust and privacy-related perceptions (Martin and Murphy, 2017; Teepapal, 2025; Markou et al., 2025). Your own recent article also supports this direction by showing the mediating relevance of trust and data privacy perceptions in the Saudi retail context (Jarrar, 2026).

This study therefore does not ask whether personalization matters. That part is already obvious. It asks a more useful question: **how does AI-driven personalization get translated into loyalty, and why does it sometimes fail or even backfire?**

The paper addresses this question through a conceptual synthesis of literature. Its main argument is that AI-driven personalization should not be framed as a direct driver of customer loyalty. Instead, it should be understood as a conditional stimulus that first passes through a privacy evaluation stage and then, if acceptable, contributes to trust formation. Loyalty emerges only after that process succeeds.

This reframing contributes to the literature in three ways. First, it separates **data privacy perception** from **trust**, rather than collapsing them into one general attitude. Second, it proposes a **sequential mechanism** in which privacy evaluation precedes trust formation. Third, it explains both positive and negative personalization outcomes within the same framework, which helps make sense of inconsistent findings in prior studies.

2. LITERATURE REVIEW

2.1 AI-Driven Personalization in Retail

AI-driven personalization refers to the use of algorithms, predictive models, and customer data to tailor retail experiences at the individual level. In practice, this includes product recommendations, dynamic promotions, personalized pricing logic, curated content, chat-based assistance, and automated interaction design. In the broader retail literature, technology-enabled personalization is often presented as part of the future of retailing because it improves relevance and facilitates better matching between customer needs and firm offerings (Grewal, Roggeveen and Nordfält, 2017).

In service and marketing settings, AI also changes the architecture of customer interaction. Huang and Rust (2018) argue that AI affects not only firm efficiency but the way service is delivered and perceived. This matters because personalization is not experienced merely as a technical output. It is interpreted by customers as part of the relationship between themselves and the firm.

The positive case for personalization is straightforward: more relevant interactions should reduce search effort, improve convenience, and increase perceived usefulness. However, the literature also shows that personalization can become counterproductive when it is too aggressive, too visible, poorly timed, or perceived as invasive. Bleier and Eisenbeiss (2015), for example, show that the effectiveness of personalized online advertising depends on contextual factors such as timing and placement, rather than on personalization alone.

Recent work reinforces the same point in AI-enabled settings. Teepapal (2025) finds that AI-driven personalization can increase trust, privacy concerns, and perceived usefulness at the same time. That result alone is enough to kill the lazy assumption that personalization is simply “good.” It is more accurate to say it triggers a bundle of reactions, some positive and some defensive.

2.2 Customer Loyalty in Digital Retail

Customer loyalty is commonly understood as a combination of favorable attitudes and repeat behavioral intentions toward a firm or brand. In digital retail, loyalty depends not only on product satisfaction but on the broader customer journey, including pre-purchase, purchase, and post-purchase interactions (Lemon and Verhoef, 2016). This makes personalization potentially relevant because it can shape how customers experience multiple touchpoints.

Still, loyalty is not produced by relevance alone. A personalized recommendation may improve convenience, but that does not guarantee a durable relational outcome. If customers feel monitored, exploited, or manipulated, the same personalization practice may weaken loyalty rather than strengthen it. This helps explain why empirical findings are often mixed and why mediation-based interpretations are more convincing than direct-effect assumptions.

Jarrar (2026) adds useful context here by showing that trust and data privacy perceptions are central to understanding how AI-driven personalization relates to customer loyalty in Saudi Arabia’s retail sector. That study is important not just because it supports mediation, but because it reflects a regional retail context rather than recycling the usual Western platform examples.

2.3 Trust as a Mediating Construct

Trust is one of the most established constructs in relationship marketing. Morgan and Hunt (1994) position trust, alongside commitment, as a key mediating variable in successful relational exchange. In digital and data-intensive environments, trust becomes even more important because uncertainty is higher and the mechanisms shaping customer treatment are often less visible.

In the context of AI-driven personalization, trust can be understood as the customer’s belief that the retailer is reliable, credible, and not acting in a harmful or exploitative way. Trust reduces perceived risk and supports

willingness to continue engaging with the brand. It is therefore not just a soft emotional variable. It functions as a conversion mechanism between evaluated firm behavior and future customer action.

Recent evidence supports this interpretation. Teepapal (2025) reports that trust positively affects consumer engagement in AI-personalized social media marketing, while Markou et al. (2025) emphasize that trust and ethical perceptions shape acceptance of AI-based personalized advertising. In plain English: customers do not just ask whether the recommendation is useful. They also ask, consciously or not, whether the firm behind it is behaving properly.

2.4 Data Privacy Perceptions as a Mediating Construct

Data privacy perception refers to how customers evaluate the collection, processing, visibility, and use of their personal information. Martin and Murphy (2017) show that data privacy is not merely a legal matter. It is a marketing issue because it shapes how customers interpret firm behavior and how willing they are to engage.

This matters heavily in personalization. Personalization depends on data. The more tailored the interaction feels, the more likely customers are to infer that the firm knows something about them. Sometimes that feels useful. Sometimes it feels like someone reading your diary and calling it customer experience.

Awad and Krishnan (2006) demonstrated that willingness to be profiled for personalization is linked to perceptions of information transparency. More recently, Alhitmi (2024) highlights that AI-driven marketing raises persistent concerns around data security, privacy, and transparency, reinforcing the need for governance and clear communication. Canhoto (2024) also shows that AI-enabled personalization in physical retail settings produces uncertainty because customers respond not only to the value of the offer but to the context in which personalization occurs and the privacy implications that surround it.

Together, these studies show that privacy perception is not a side issue. It is one of the first interpretive filters through which personalization is judged.

3. CONCEPTUAL SYNTHESIS AND FRAMEWORK DEVELOPMENT

3.1 Reframing the Personalization–Loyalty Relationship

The literature reviewed above suggests that the relationship between AI-driven personalization and customer loyalty is best understood as **conditional rather than linear**. Personalization introduces value in the form of relevance, efficiency, and convenience, but it also introduces risk in the form of surveillance concerns, uncertainty, and perceived loss of control (Awad and Krishnan, 2006; Martin and Murphy, 2017; Alhitmi, 2024).

This means customers do not respond to personalization as a pure output. They respond to what personalization **signals** about the firm's data practices and intentions. That distinction is important because it explains why technically effective personalization can still produce weak or negative customer outcomes.

A central problem in some prior discussions is that **trust** and **privacy** are treated as if they are interchangeable. They are not.

- **Data privacy perception** is primarily a **cognitive evaluation**: *Is my data being handled properly? Is this acceptable?*
- **Trust** is a **relational judgment**: *Do I believe this company is credible, reliable, and acting appropriately toward me?*

That distinction strengthens the framework because it gives each construct a specific job rather than letting both float around as generic "customer feelings."

3.2 Mechanism of Influence

The proposed mechanism consists of four stages.

Stage 1: Exposure to AI-Driven Personalization

The customer encounters personalized recommendations, targeted communication, adaptive interface content, or other AI-enabled retail stimuli. At this stage, the firm's intention is usually to improve relevance and increase engagement (Grewal, Roggeveen and Nordfält, 2017).

Stage 2: Cognitive Evaluation Through Data Privacy Perception

The customer evaluates whether the personalization practice appears acceptable. This evaluation includes perceived transparency, control, appropriateness, sensitivity of data involved, and the general sense of whether the firm crossed a line. This stage aligns closely with the logic of the personalization–privacy paradox and broader privacy calculus thinking, where benefits are weighed against risks (Awad and Krishnan, 2006; Martin and Murphy, 2017; Canhoto, 2024).

This stage is crucial because **negative evaluation can stop the process early**. If customers perceive the personalization as intrusive or opaque, trust may not develop at all.

Stage 3: Relational Judgment Through Trust Formation

If privacy concerns remain within acceptable limits, the customer proceeds to a more relational interpretation of the firm. At this point, trust can develop based on perceived competence, fairness, credibility, and responsible conduct. Morgan and Hunt (1994) provide the foundational logic here, while recent research shows that AI-based personalization acceptance is strongly shaped by trust and ethical interpretation (Teepapal, 2025; Markou et al., 2025).

Stage 4: Behavioral Outcome in the Form of Customer Loyalty

Only after the previous stages succeed does personalization have a realistic path toward loyalty-related outcomes such as repeat purchase intention, customer retention, and ongoing engagement. This is consistent with broader customer journey literature showing that loyalty is the result of accumulated, interpreted experience rather than isolated touchpoint efficiency (Lemon and Verhoef, 2016).

3.3 Sequential Logic: Privacy Before Trust

One of the strongest contributions of this paper is the argument that **privacy perception precedes trust formation**.

That sequence matters. It means privacy is not simply another variable sitting next to trust. It acts as an earlier evaluative screen. If the personalization practice fails at the privacy stage, the customer may never move into a trust-building process at all. This helps explain why some personalization efforts fail despite being accurate or useful in technical terms.

In short:

- **Privacy asks:** *Is this acceptable?*
- **Trust asks:** *Is this company worthy of reliance?*

That ordering makes the framework more defensible and better able to explain contradictory findings in prior research.

3.4 Conceptual Propositions

Based on the synthesis, the following propositions are developed:

P1. AI-driven personalization influences customer loyalty indirectly rather than through a simple direct effect (Jarrar, 2026; Teepapal, 2025).

P2. Data privacy perception functions as a primary cognitive filter through which customers evaluate the acceptability of AI-driven personalization (Awad and Krishnan, 2006; Martin and Murphy, 2017; Canhoto, 2024).

P3. Trust functions as a relational mediator that converts acceptable personalization into positive loyalty-related outcomes (Morgan and Hunt, 1994; Markou et al., 2025; Jarrar, 2026).

P4. Negative privacy perceptions can weaken or reverse the effect of AI-driven personalization by increasing resistance, undermining trust, and reducing loyalty (Bleier and Eisenbeiss, 2015; Alhitmi, 2024).

P5. The same personalization practice may generate both value and concern simultaneously, meaning that personalization effectiveness depends on the balance between perceived usefulness and perceived ethical acceptability (Teepapal, 2025; Canhoto, 2024).

4. DISCUSSION

4.1 Personalization Is Not a Direct Loyalty Engine

The main implication of this synthesis is that AI-driven personalization should not be treated as a direct loyalty engine. That assumption is too neat for reality. The literature supports a more contingent explanation: personalization provides a stimulus, but customer outcomes depend on how that stimulus is interpreted in terms of privacy, fairness, and trustworthiness (Awad and Krishnan, 2006; Martin and Murphy, 2017; Teepapal, 2025). This matters because many managerial and academic discussions still talk about personalization as if more precision automatically means better customer relationships. It does not. Better targeting can improve relevance, but relevance is only one part of the psychological equation.

4.2 Explaining Inconsistent Prior Findings

The framework also helps explain why prior studies often report inconsistent effects.

- Studies showing **positive effects** of personalization likely reflect contexts in which customers perceived the practice as acceptable and the firm as trustworthy.
- Studies showing **weak, mixed, or negative effects** likely reflect contexts in which customers interpreted the personalization as invasive, poorly timed, or insufficiently transparent.

Bleier and Eisenbeiss (2015) already show that personalization effectiveness depends on contextual conditions. Recent AI-focused studies make the picture even clearer. Teepapal (2025) finds that AI-enabled personalization can increase both trust and privacy concerns at once, and that personalization itself may not significantly affect

engagement directly. That is a strong sign that mediation and interpretation matter more than brute personalization intensity.

4.3 Why Privacy and Trust Must Be Separated

Separating privacy perception from trust is not just a semantic cleanup. It improves explanatory power.

If both constructs are treated as one broad “customer comfort” variable, the model becomes mushy. By separating them, the framework can show that a customer may acknowledge value in personalization while still objecting to the data practice behind it. Likewise, a customer may not object strongly to the data use but still withhold trust because the firm appears manipulative or self-serving.

Recent research supports this distinction. Markou et al. (2025) connect acceptance of AI-based personalization not only to utility but to trust, ethics, and identity-related concerns. Canhoto (2024) similarly shows that AI-enabled retail personalization creates contextual uncertainty that cannot be explained by usefulness alone.

4.4 When Personalization Backfires

A particularly important implication is that personalization can do more than merely fail. It can backfire.

When personalization is perceived as excessive, intrusive, or poorly explained, customers may feel discomfort, loss of control, or suspicion. This can produce distrust and avoidance rather than engagement. Alhitmi (2024) emphasizes that privacy and security concerns remain central barriers in AI-driven marketing, while Canhoto (2024) demonstrates that in-store AI-enabled personalization can trigger uncertainty because customers react not only to the content of the offer but to the circumstances and implications of its delivery.

That means firms should stop measuring personalization success as if it were just a conversion game. A campaign that lifts clicks while eroding trust is not a clean win. It is a delayed problem wearing a KPI costume.

4.5 Theoretical Contribution

This paper contributes conceptually by integrating relationship marketing logic with privacy-oriented evaluation logic into a process-based model.

From relationship marketing, the framework draws on the idea that trust is a key mediating mechanism in relational exchange (Morgan and Hunt, 1994). From privacy scholarship, it draws on the idea that customer acceptance of profiling depends on information transparency, risk perception, and the broader handling of personal data (Awad and Krishnan, 2006; Martin and Murphy, 2017). More recent AI-focused work extends this by showing that personalization in algorithmic settings intensifies the salience of ethics, transparency, and privacy governance (Alhitmi, 2024; Markou et al., 2025).

The paper’s specific contribution is to place these pieces in sequence: **AI-driven personalization** → **data privacy perception** → **trust** → **customer loyalty**, with a negative path for intrusive personalization leading toward distrust and reduced loyalty.

4.6 Practical Relevance in Current Retail Reality

In current retail practice, AI-driven personalization is no longer unusual. It is increasingly standard. That shifts the competitive question. The differentiator is no longer simply whether a retailer can personalize, but whether it can do so in a way that customers perceive as legitimate, proportionate, and trustworthy.

Recent scholarship supports that shift. Alhitmi (2024) stresses the growing importance of transparency, regulation, and security in AI-driven marketing. Teepapal (2025) shows that trust and perceived usefulness matter more than personalization intensity alone. Markou et al. (2025) further tie acceptance to ethical perception and identity-related concerns.

So the practical takeaway is blunt: **personalization is not the real moat anymore; trust-preserving personalization is.**

5. MANAGERIAL IMPLICATIONS

The findings have direct implications for retail managers, CX leaders, marketers, and digital teams.

First, firms should stop equating personalization quantity with personalization quality. More data use and more aggressive targeting do not automatically improve customer outcomes. Personalization should be designed around **appropriateness**, not just precision.

Second, data privacy should be treated as a core part of the customer experience, not pushed into a legal or IT corner. Customers do not experience privacy as an internal compliance category. They experience it as part of whether the brand feels safe, respectful, and in control.

Third, trust should be viewed as a design objective. Retailers often invest in recommendation quality, automation, and optimization, but underinvest in transparency and explainability. A simple “why you’re seeing this” explanation, clear preference controls, and visible restraint in personalization can matter more than another layer of predictive cleverness.

Fourth, performance measurement needs to improve. If personalization is evaluated only through click-through rates, conversions, or short-term sales lift, firms may miss the reputational and relational damage caused by intrusive tactics. Metrics should also include privacy complaints, opt-out behavior, trust indicators, and signals of customer discomfort.

Fifth, cross-functional governance matters. Personalization sits at the intersection of marketing, technology, customer experience, and compliance. If these functions are not aligned, firms tend to optimize the algorithm while neglecting the customer's interpretation of what the algorithm is doing.

In short, the managerial lesson is simple: **the future is not “use more AI.” It is “use AI in a way that customers can live with.”**

6. CONCLUSION, LIMITATIONS, AND FUTURE RESEARCH

This study set out to reframe the relationship between AI-driven personalization and customer loyalty in retail. Based on a structured synthesis of foundational and recent literature, the paper argues that personalization should not be treated as a direct driver of loyalty. Instead, its effects are mediated by customer interpretations, especially data privacy perceptions and trust.

The proposed framework suggests that customers first evaluate whether the personalization practice is acceptable in terms of data use, transparency, and control. Only if that evaluation is sufficiently positive can trust develop. Loyalty-related outcomes then emerge through that relational pathway. This sequential interpretation helps explain why personalization can produce positive, neutral, or negative outcomes across different contexts.

The paper contributes by distinguishing privacy perception from trust, placing them in sequence, and incorporating both positive and negative response paths into one conceptual model. That gives the framework more explanatory precision than generic “personalization improves loyalty” claims.

This study is conceptual and therefore has limitations. It does not empirically test the proposed sequence or compare its strength across retail categories, platforms, or customer segments. Future research should validate the framework using quantitative and qualitative methods in different contexts, including physical retail, e-commerce, and omnichannel settings. It would also be useful to test boundary conditions such as digital literacy, regulatory awareness, culture, product sensitivity, and the type of data used for personalization.

Finally, because AI capability is evolving quickly, future work should continue examining not only whether customers accept personalization, but what kinds of governance, transparency, and restraint make that acceptance sustainable.

REFERENCES

- 1) Alhitmi, H.K. (2024) ‘Data security and privacy concerns of AI-driven marketing in the context of economics and business field: an exploration into possible solutions’, *Cogent Business & Management*, 11(1).
- 2) Awad, N.F. and 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), pp. 13–28.
- 3) Bleier, A. and Eisenbeiss, M. (2015) ‘Personalized online advertising effectiveness: The interplay of what, when, and where’, *Marketing Science*, 34(5), pp. 669–688.
- 4) Canhoto, A.I. (2024) ‘Unpacking the personalisation-privacy paradox in the context of AI-enabled personalisation in the physical retail environment’, *Information Systems Frontiers*, 26(3).
- 5) Grewal, D., Roggeveen, A.L. and Nordfält, J. (2017) ‘The future of retailing’, *Journal of Retailing*, 93(1), pp. 1–6.
- 6) Huang, M.-H. and Rust, R.T. (2018) ‘Artificial intelligence in service’, *Journal of Service Research*, 21(2).
- 7) Jarrar, S.M. (2026) ‘The role of trust and data privacy perceptions in mediating the relationship between AI-driven personalization and customer loyalty in Saudi Arabia’s retail sector’, *International Journal of Engineering Technology Research & Management (IJETRM)*, 10(03), pp. 381–388. DOI: 10.5281/zenodo.19325425
- 8) Lemon, K.N. and Verhoef, P.C. (2016) ‘Understanding customer experience throughout the customer journey’, *Journal of Marketing*, 80(6), pp. 69–96.
- 9) Markou, V., Serdaris, P., Antoniadis, I. and Spinthiropoulos, K. (2025) ‘Personalization, trust, and identity in AI-based marketing: An empirical study of consumer acceptance in Greece’, *Administrative Sciences*, 15(11), 440.

IJETRM

International Journal of Engineering Technology Research & Management (IJETRM)

Journal Article

<https://ijetrm.com/issue/>

- 10) Martin, K.D. and Murphy, P.E. (2017) 'The role of data privacy in marketing', *Journal of the Academy of Marketing Science*, 45(2), pp. 135–155.
- 11) Morgan, R.M. and Hunt, S.D. (1994) 'The commitment-trust theory of relationship marketing', *Journal of Marketing*, 58(3), pp. 20–38.
- 12) Teepapal, T. (2025) 'AI-driven personalization: Unraveling consumer perceptions in social media engagement', *Computers in Human Behavior*, 165.