From Artificial Intelligence to Human-Centric Marketing: Examining the Impact of AI-Powered Personalization on Consumer Trust and Purchase Intention

Ms. Shrishti Singh[1], Ms. Anjali Raghuvanshi[2], Ms. Mansi Bhati[3]

[1]Assistant Professor; KCC Institute of Legal and Higher Education, Greater Noida. Email: Shrishti734@gmail.com
[2]Research Scholar; Amity College of Commerce and Finance, Amity University, Noida. Email: Anjali.raghuvanshi@s.amity.edu
[3]Research Scholar; Amity College of Commerce and Finance, Amity University, Noida. Email: Mansi.bhati@s.amity.edu

Abstract
Artificial intelligence (AI) is reshaping marketing practice by enabling firms to personalize recommendations, advertisements, and digital interactions at an individual level. Yet the growing technical sophistication of AI-powered personalization does not automatically translate into favorable consumer responses. Prior research has separately examined relevance, privacy concerns, trust, and purchase intention as outcomes of AI-enabled personalization, but the literature remains theoretically fragmented regarding the psychological sequence through which algorithmic personalization is converted into a human-centric consumer experience. This conceptual paper addresses that gap by developing a proposition-based framework, grounded in the Stimulus–Organism–Response (S-O-R) paradigm, that explains how AI-powered personalization (stimulus) shapes perceived relevance and consumer trust (organism) before influencing purchase intention (response). The paper argues that personalization functions as a technological capability, whereas relevance and trust represent the consumer's subjective translation of that capability into meaningful experience. Four propositions are developed: personalization positively influences perceived relevance; relevance positively influences trust; trust positively influences purchase intention; and relevance and trust sequentially mediate the personalization–purchase intention relationship. The paper's principal theoretical contribution is the reconceptualization of perceived relevance as the translation mechanism connecting algorithmic intelligence to subjective consumer value, and the argument that human-centricity in AI marketing is an emergent property of consumer interpretation rather than of technological sophistication alone. The framework is positioned within the Industry 5.0 discourse on human-centric, responsible technological innovation and offers managers a basis for evaluating AI systems by relevance and trustworthiness rather than by data intensity alone. The paper concludes with boundary conditions and a research agenda for future empirical testing.

Keywords: artificial intelligence; AI-powered personalization; human-centric marketing; perceived relevance; consumer trust; purchase intention; Industry 5.0

1. Introduction

Artificial intelligence has moved from a peripheral technical capability to a structural component of contemporary marketing strategy. Firms increasingly rely on AI to analyze behavioral data, infer preferences, automate interactions, and generate individualized recommendations across the customer journey. Huang and Rust's (2021) strategic framework for AI in marketing captures this shift by distinguishing mechanical, thinking, and feeling AI and by showing how each expands firms' capacity to personalize marketing activity at a scale that earlier segmentation-based approaches could not achieve.

The rise of AI-enabled personalization has changed what consumers encounter in digital markets. Rather than standardized mass communication, consumers routinely receive product recommendations, advertisements, search results, and promotional content that have been algorithmically tailored to their individual behavior. AI-enabled personalization now extends across the customer journey, from initial profiling and navigation through nudging and retention (Gao & Liu, 2023).

Yet the technical capacity to personalize does not by itself guarantee that consumers will respond favorably. Personalization has been shown to operate as a paradox: the same underlying data practices that increase the relevance and usefulness of marketing content can simultaneously increase consumers' sense of vulnerability and resistance (Aguirre et al., 2015). Greater algorithmic sophistication compounds this tension, because AI systems process a wider range of behavioral signals than earlier rule-based personalization, which can heighten consumer awareness of surveillance and algorithmic influence

even as it improves recommendation accuracy (Cai & Mardani, 2023; Walker, 2024).

The emergence of Industry 5.0 provides a relevant frame for revisiting this tension. Where earlier industrial paradigms emphasized efficiency and automation, Industry 5.0 foregrounds human-centricity, sustainability, and resilience, positioning technology as a means of supporting human needs and stakeholder value rather than an end in itself (Breque et al., 2021). Extended to marketing, this logic implies that the strategic question for AI-enabled personalization is not only how accurately a system can predict consumer preferences, but whether consumers experience the resulting interactions as meaningful, relevant, and trustworthy. Consistent with this view, recent consumer research argues that AI technologies carry experiential, social, and psychological consequences for consumers that cannot be reduced to measures of predictive accuracy (Puntoni et al., 2021; Hermann & Puntoni, 2024).

Existing research on AI and marketing has expanded rapidly but remains fragmented across distinct outcome variables. Some studies examine the technical capabilities and strategic applications of AI (Huang & Rust, 2021); others examine consumer trust in AI-enabled shopping and service contexts (Chakraborty et al., 2024), privacy and purchase resistance (Cai & Mardani, 2023; Walker, 2024), authenticity perceptions of AI-generated communication (Kirk & Givi, 2024), or relational breakdowns with AI-powered assistants (Saenger et al., 2024). Recent empirical work on AI-powered personalized advertising has begun to connect several of these strands, showing that perceived relevance, usefulness, and trust jointly help explain purchase intention (An & Ngo, 2025). However, the literature has not yet articulated, at a theoretical level, the sequential psychological process through which algorithmic personalization is translated into a human-centric consumer experience and subsequently into behavioral intention. Studies tend to treat relevance, trust, and purchase intention as parallel or loosely connected outcomes of personalization rather than as stages in a coherent organismic process.

This paper addresses that gap by asking: **How does AI-powered personalization become translated into a human-centric consumer experience that fosters trust and purchase intention?**

To answer this question, the paper develops a conceptual framework grounded in the Stimulus–Organism–Response (S-O-R) paradigm, proposing that AI-powered personalization operates as an external technological stimulus that first shapes consumers' cognitive evaluation of relevance, which in turn shapes their relational evaluation of trust, which finally shapes their behavioral purchase intention:

AI-Powered Personalization → Perceived Relevance → Consumer Trust → Purchase Intention

The paper makes four contributions. First, it reframes AI-powered personalization as a technological capability that is analytically distinct from the human-centric experience it may or may not produce, shifting the evaluative focus from algorithmic sophistication to consumer meaning-making. Second, it develops perceived relevance as the conceptual mechanism through which algorithmic intelligence is translated into subjective consumer value, since consumers do not experience algorithms directly but only their outputs. Third, it establishes a sequential organismic mechanism—relevance preceding trust—that offers a more theoretically precise account of AI-enabled consumer response than models that treat relevance and trust as independent, parallel mediators. Fourth, it extends the emerging discourse on human-centric marketing into the Industry 5.0 era by operationalizing human-centricity through the constructs of relevance, trust, and consumer value.

The remainder of the paper is organized as follows. Section 2 reviews the conceptual foundations of AI in marketing, the personalization paradox, and human-centric marketing under Industry 5.0. Section 3 develops the S-O-R theoretical framework underlying the model. Section 4 develops the paper's four propositions. Section 5 presents the conceptual model. Sections 6 through 9 discuss the framework's theoretical contributions, managerial implications, and boundary conditions, and Section 10 concludes.

2. Conceptual Foundations and Literature Background

2.1 Artificial Intelligence in Marketing: From Automation to Consumer Experience

AI in marketing has historically been framed around what algorithmic systems can accomplish: processing large volumes of behavioral data, automating decisions, and improving the accuracy of targeting and recommendation. Huang and Rust's (2021) framework remains the most influential articulation of this capability-oriented view, distinguishing mechanical AI (automating repetitive tasks), thinking AI (processing data to generate insight), and feeling AI (recognizing and responding to emotion), and mapping each onto marketing research, strategy, and action.

A parallel and increasingly influential stream of research has shifted attention from what AI can do to how consumers experience what AI does. Puntoni et al. (2021) argue that AI should not be treated as a neutral instrument of efficiency; rather, AI reshapes consumers' sense of agency, their social relationships with firms and other actors, and even aspects of their identity. Hermann and Puntoni (2024) extend this

experiential perspective to generative AI, noting that as AI systems move from predictive to generative capabilities, their psychological and social consequences for consumers become correspondingly more consequential. Related work has examined trust formation in AI-enabled shopping assistance (Chakraborty et al., 2024), authenticity judgments toward AI-generated marketing content (Kirk & Givi, 2024), and the effects of AI-enabled service failures on vulnerable consumers (Hermann et al., 2024).

Taken together, this literature suggests that the marketing discipline's understanding of AI has progressively moved from a technology-centric question—what can AI do?—toward a consumer-centric question—how do consumers interpret and evaluate what AI does? This distinction is foundational to the present paper: AI-powered personalization is treated here as a technological stimulus whose marketing value depends entirely on how it is subsequently interpreted by consumers, not on its underlying computational sophistication.

2.2 AI-Powered Personalization and the Personalization Paradox

Personalization refers to the adaptation of marketing content, recommendations, or experiences to the characteristics of individual consumers. AI substantially extends earlier segmentation-based personalization by enabling firms to process heterogeneous behavioral signals—browsing history, purchase patterns, search activity, and contextual data—at a scale and granularity that earlier systems could not achieve (Gao & Liu, 2023). AI-enabled personalization now spans profiling, navigation support, real-time nudging, and retention-oriented communication across the customer journey.

The promise of this capability is intuitive: the more precisely a system can infer an individual's preferences, the more closely marketing communication can, in principle, align with that individual's needs. This promise, however, is not self-fulfilling. Aguirre et al.'s (2015) foundational work on the personalization paradox demonstrates that personalized marketing simultaneously increases perceived relevance and adoption while generating feelings of vulnerability tied to how consumer information is collected and disclosed. The same underlying mechanism—data-driven inference—thus produces both the benefit (relevance) and the risk (intrusiveness) of personalization.

This paradox is amplified in AI-enabled contexts. Because AI systems draw on a wider and less transparent range of behavioral inputs than earlier personalization methods, they can simultaneously improve recommendation quality and heighten consumer awareness of data collection and algorithmic influence (Cai & Mardani, 2023). Consumers may therefore hold a positive appraisal of what personalization delivers while holding a more ambivalent appraisal of how it was produced. Walker (2024) similarly shows that privacy concerns and perceptions of platform responsibility remain central to whether AI-driven personalization is experienced as beneficial or intrusive.

The theoretical implication is that the effectiveness of AI-powered personalization cannot be inferred from its technical properties alone. It is contingent on consumer interpretation—on whether the personalized output is judged to be relevant to the individual's needs and produced by a system the individual is willing to trust. This is the conceptual space the present paper occupies.

2.3 Human-Centric Marketing in the Industry 5.0 Era

Industry 5.0 represents a deliberate reorientation of technological strategy away from efficiency-maximization alone and toward human-centricity, sustainability, and resilience (Breque et al., 2021). Its central premise is that technological systems should be designed to serve human needs, agency, and well-being rather than to optimize organizational metrics at the expense of the individuals those systems affect.

Extended to marketing, this principle motivates the concept of human-centric marketing, which this paper defines as the design and deployment of marketing technologies in ways that prioritize consumer needs, agency, meaningful value, and trustworthy experience, rather than data intensity or targeting efficiency alone. Human-centric marketing is distinguished from three adjacent orientations. It differs from technology-centric marketing, which evaluates AI systems primarily by computational or predictive sophistication. It differs from data-centric marketing, which treats the volume and granularity of consumer data as the primary driver of marketing value. And it differs from automation-centric marketing, which prioritizes the efficiency and scalability of marketing execution over the quality of the consumer's experience of that execution.

Under a human-centric framing, AI-powered personalization does not become valuable simply because it processes more consumer data or automates more interactions. It becomes valuable when consumers experience its outputs as personally relevant and when that experience is sufficiently trustworthy to support their decision-making. This reframing motivates the paper's central constructs—perceived relevance and consumer trust—as the mechanisms through which algorithmic personalization is, or is not, converted into a human-centric marketing outcome.

3. Theoretical Framework

3.1 The Stimulus–Organism–Response Framework

The proposed model is grounded in the Stimulus–

Organism–Response (S-O-R) framework, which holds that external environmental stimuli shape individuals' internal cognitive and affective states (the organism), which in turn shape behavioral responses. The S-O-R paradigm is well suited to AI-enabled marketing contexts because consumers are increasingly exposed to algorithmically generated stimuli—recommendations, advertisements, and personalized content—whose behavioral effects are mediated by internal psychological evaluation rather than occurring automatically.

The present paper maps the S-O-R components as follows:

S-O-R Component	Construct
Stimulus	AI-Powered Personalization
Organism – Cognitive Evaluation	Perceived Relevance
Organism – Relational Evaluation	Consumer Trust
Response	Purchase Intention

This mapping should not be applied mechanically. The theoretical contribution of this paper lies specifically in how the organism component is conceptualized. Rather than treating the organism as a single undifferentiated internal state, the paper argues that the organism must be modeled as a sequential internal process comprising two analytically distinct evaluations: a cognitive evaluation of relevance, followed by a relational evaluation of trust.

This sequencing is theoretically necessary for the following reasons. First, consumers encounter AI-powered personalized communication as a discrete stimulus—a recommendation, an advertisement, a search result. Second, before any relational judgment is possible, consumers must cognitively evaluate whether that stimulus is relevant to their needs; relevance is a precondition for further engagement, not a byproduct of it. Third, once relevance has been established, it provides consumers with experiential evidence of the system's competence and responsiveness—evidence that supports a subsequent, qualitatively different evaluation: whether the system can be relied upon, i.e., trusted. Fourth, trust reduces the uncertainty that is characteristic of AI-enabled recommendations and digital purchasing, and it is this reduction in uncertainty—rather than relevance alone—that translates into a willingness to act.

In short, relevance answers the question "is this useful to me?" while trust answers the qualitatively different question "can I rely on this?" Because the second question is only meaningfully posed once the first has been answered affirmatively, relevance must theoretically precede trust in the proposed mechanism. This ordering is one of the paper's central theoretical contributions and is developed further in Section 4.

4. Theory Development and Research Propositions

4.1 AI-Powered Personalization and Perceived Relevance

Perceived relevance refers to the extent to which a consumer judges marketing content, recommendations, or advertisements to be meaningful and applicable to their individual needs, interests, and circumstances. Relevance is a subjective consumer evaluation, not an objective property of the technology that produced it.

AI-powered personalization draws on browsing behavior, purchase history, search activity, and contextual information to generate individualized content. The underlying premise of personalization is that better inference about an individual should improve the alignment between marketing content and that individual's needs. Recent empirical work on AI-powered personalized advertising supports this premise, showing that perceived personalization is positively associated with perceived relevance (An & Ngo, 2025).

This relationship, however, should not be treated as automatic or self-evident. The present paper emphasizes a distinction that is often collapsed in the literature: personalization is what the system does; relevance is what the consumer experiences. A system may be technically personalized—drawing on extensive behavioral data and sophisticated inference—without being experienced as relevant, if the resulting output does not correspond to the consumer's actual needs, timing, or context. Relevance is therefore the first point at which algorithmic capability is filtered through consumer judgment, and it is this filtering, rather than the technical sophistication of personalization itself, that determines whether personalization generates any downstream marketing value.

Proposition 1: AI-powered personalization positively influences consumers' perceived relevance of marketing communication and recommendations.

4.2 Perceived Relevance and Consumer Trust

Consumer trust is understood here as a willingness to rely on a firm, platform, or recommendation

system based on perceptions of its reliability, competence, and benevolent intent. Digital consumption environments are characterized by substantial uncertainty regarding product quality, recommendation accuracy, data usage, and algorithmic decision-making; trust functions as a mechanism for reducing this uncertainty.

The present paper argues that perceived relevance is a critical antecedent of trust in AI-enabled marketing environments, for the following reasons. First, relevant recommendations provide direct experiential evidence that the system has correctly interpreted the consumer's preferences—evidence of competence. Second, when relevance is experienced repeatedly rather than as an isolated instance, it demonstrates a pattern of responsiveness that goes beyond a single accurate prediction. Third, this demonstrated responsiveness reduces the consumer's uncertainty about whether the system's future outputs can be relied upon. Fourth, this reduction in uncertainty is the proximate condition for trust formation.

It is important to maintain a clear conceptual boundary between relevance and trust rather than treating them as interchangeable. Relevance is a cognitive value evaluation—a judgment about whether a specific output is useful. Trust is a relational evaluation—a broader judgment about whether the system, as an ongoing relational partner, can be relied upon. Relevance concerns a discrete output; trust concerns an ongoing relationship. Empirical work on AI-powered personalized advertising is consistent with this ordering, indicating that relevance plays an important explanatory role in trust formation and, subsequently, purchase intention (An & Ngo, 2025), while broader research on AI-enabled service systems shows that trust is shaped by accumulated experience with system performance rather than by any single interaction (Chakraborty et al., 2024).

Proposition 2: Perceived relevance positively influences consumer trust in AI-enabled marketing environments.

4.3 Consumer Trust and Purchase Intention

Purchase intention refers to a consumer's conscious willingness or likelihood to purchase a product or service. Purchasing decisions in AI-enabled digital environments are typically made under conditions of incomplete information: consumers must evaluate not only the product being recommended but also the credibility of the recommendation system itself. This introduces a layer of uncertainty—concerning information asymmetry, recommendation credibility, and algorithmic decision-making—that is additional to ordinary purchase-related risk.

Trust reduces this uncertainty by providing consumers with a basis for relying on the recommendation system's outputs without independently verifying each one. Recent studies confirm that trust is a significant predictor of purchase intention in AI-powered personalized advertising contexts (An & Ngo, 2025) and in AI-enabled commerce and service settings more broadly (Chakraborty et al., 2024). The inverse relationship is also instructive: research on relational breakdowns with AI-powered virtual assistants shows that when consumers perceive betrayal by an AI system, their purchase intentions for the recommended products decline sharply (Saenger et al., 2024), underscoring that trust—rather than relevance or personalization alone—is the proximate psychological condition for behavioral commitment.

Proposition 3: Consumer trust positively influences purchase intention in AI-enabled marketing environments.

4.4 The Sequential Human-Centric Personalization Mechanism

The preceding propositions, considered individually, establish plausible pairwise relationships. The central theoretical contribution of this paper, however, lies in integrating them into a single sequential mechanism rather than treating relevance and trust as parallel, independent mediators of the personalization–purchase intention relationship.

The theoretical case for sequential rather than parallel mediation follows directly from the logic developed in Sections 3.1 and 4.1–4.3. Personalization without relevance risks being experienced as noise, or worse, as surveillance: data-intensive targeting that does not correspond to genuine consumer need offers no experiential evidence on which trust could be built, and may instead heighten the vulnerability associated with the personalization paradox (Aguirre et al., 2015). Relevance without trust, meanwhile, may generate momentary interest in a specific recommendation without producing the relational confidence required for behavioral commitment; a single relevant output does not, by itself, resolve the broader uncertainty surrounding an AI system's reliability. Trust is therefore the mechanism that converts a positive but isolated evaluation (relevance) into a durable willingness to act (purchase intention). This reasoning yields a specific process structure:

Stage 1 (Technological Stimulus): AI-powered personalization generates individualized marketing content.

Stage 2 (Consumer Meaning-Making): The consumer cognitively evaluates whether this content is relevant to their needs.

Stage 3 (Relational Evaluation): Relevant experiences accumulate into evidence of system competence and reliability, supporting trust.

Stage 4 (Behavioral Response): Trust reduces uncertainty and increases willingness to purchase.

Recent empirical work on AI-powered personalized advertising lends support to this sequential logic, finding that personalization does not translate into trust directly but operates through the value consumers derive from personalized communication—consistent with relevance functioning as a necessary intermediate evaluation rather than an optional or parallel one (An & Ngo, 2025).

The central argument advanced here is that human-centricity emerges through consumer interpretation rather than through technological sophistication alone. A system does not become human-centric by virtue of processing more consumer data or generating more precisely targeted content; it becomes human-centric only insofar as its outputs are interpreted by consumers as relevant, and only insofar as that relevance accumulates into trust sufficient to support decision-making.

Proposition 4: Perceived relevance and consumer trust sequentially mediate the relationship between AI-powered personalization and purchase intention.

5. Conceptual Framework

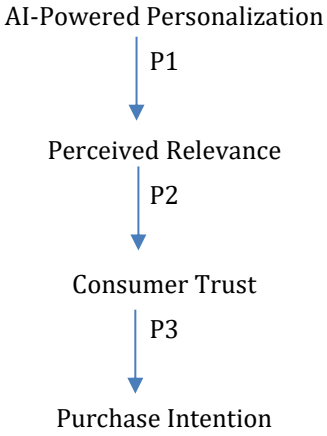

Figure 1. A Human-Centric AI Personalization Framework

Sequential mediation (P4):

AI-Powered Personalization → Perceived Relevance → Consumer Trust → Purchase Intention

The model depicts a single directional pathway rather than a set of independent bivariate relationships. AI-powered personalization is positioned as the exogenous stimulus; perceived relevance and consumer trust are positioned as sequential organismic evaluations, ordered to reflect the theoretical argument developed in Sections 3 and 4; and purchase intention is positioned as the behavioral response. Propositions 1 through 3 correspond to the three direct paths in the model, and Proposition 4 corresponds to the sequential indirect pathway that constitutes the model's central mechanism.

6. Discussion

The framework developed in this paper reorients the evaluation of AI-powered personalization away from a purely technological question and toward a consumer-experiential one. The conventional managerial question—how accurately can AI predict consumer preferences?—is theoretically insufficient, because predictive accuracy and consumer acceptance are not identical constructs. A system can generate highly accurate predictions of consumer behavior while nonetheless producing negative consumer reactions if its outputs are experienced as intrusive, poorly timed, or untrustworthy. The personalization paradox illustrates this precisely: the same data practices that increase relevance can simultaneously increase perceived vulnerability, depending on how information collection is perceived (Aguirre et al., 2015).

Two contrasting logics can be distinguished. A technology-centric logic holds that better algorithms produce better targeting, which produces higher conversion—a direct, largely mechanical chain in which consumer psychology is treated as a pass-through rather than a determining factor. A human-centric logic, by contrast, holds that AI personalization influences consumer relevance judgments, which shape trust, which shapes behavioral intention—a chain in which consumer interpretation is the decisive variable at every stage. These are not merely stylistic differences in description; they imply different diagnostic questions when personalization underperforms. Under a technology-centric logic, underperformance suggests a need for better algorithms or more data. Under the human-centric logic developed here,

underperformance more plausibly indicates a breakdown in the relevance or trust stages of the process—a diagnosis that data intensity alone cannot correct.

This distinction is directly relevant to the Industry 5.0 context introduced in Section 2.3. Industry 5.0's emphasis on human-centricity implies that AI-enabled marketing should be evaluated not by what it enables firms to know about consumers, but by what it enables consumers to experience: relevance, understanding, and trustworthy support for their own decision-making. The framework developed here operationalizes that principle within a testable psychological mechanism.

7. Theoretical Contributions

7.1 Reconceptualizing AI Personalization

The paper contributes to AI marketing literature by decoupling technological capability from consumer outcome. Where much of the existing literature has emphasized what AI enables firms to do—automate, predict, target, and personalize at scale (Huang & Rust, 2021)—the present framework repositions the evaluative criterion around what consumers experience as a result. This reframes human-centricity as an emergent, interpretive outcome rather than a designed technical property.

7.2 Perceived Relevance as a Translation Mechanism

The paper's central conceptual contribution is the development of perceived relevance as the mechanism that translates algorithmic intelligence into subjective consumer value. Consumers do not experience algorithms directly; they experience recommendations, advertisements, and interactions. Algorithmic capability therefore influences consumer behavior only through the way consumers interpret its outputs. Positioning relevance explicitly as this translation point provides a more precise account of why some technically sophisticated personalization efforts fail to generate favorable consumer response: the failure occurs not necessarily at the algorithmic level but at the interpretive level, where the translation from data to meaning breaks down.

7.3 The Sequential Organism Mechanism

The paper's third contribution is the theoretical argument that relevance must precede trust, rather than the two constructs functioning as independent, parallel mediators. This ordering provides a more theoretically precise account of AI-enabled consumer response than models that treat relevance and trust as interchangeable outcomes of personalization, because it specifies not only that both constructs matter but why they matter in a particular sequence—relevance as evidence of competence, trust as the relational judgment that evidence supports.

7.4 Extending Human-Centric Marketing to AI Personalization

Finally, the paper extends the human-centric marketing concept, developed in relation to Industry 5.0 principles, into the specific domain of AI-powered personalization. It offers a way to operationalize human-centricity conceptually through the constructs of relevance, trust, and consumer value, without overstating the novelty of the underlying psychological constructs themselves, each of which has an established history in consumer research. The contribution lies in their sequential integration within an AI-specific stimulus–organism–response mechanism, not in the invention of the constructs themselves.

8. Managerial Implications

The proposed framework carries several implications for marketing managers, digital platform managers, e-commerce firms, and AI technology developers.

1. Optimize for relevance, not personalization intensity. Firms frequently equate more data and more granular targeting with better personalization. The framework suggests that the relevant managerial objective is the consumer's experience of relevance, not the volume or precision of the underlying data processing. Investment in additional data inputs that do not demonstrably improve perceived relevance is unlikely to generate the downstream trust and purchase-intention benefits the framework identifies.

2. Measure perceived relevance directly, not only behavioral proxies. Digital marketing systems typically track click-through and conversion rates as proxies for personalization quality. The framework implies that firms should also directly measure consumers' perceived relevance—whether recommendations are seen as matching genuine needs—since relevance, not engagement alone, is the theorized precondition for trust.

3. Design trust deliberately into AI systems, rather than assuming it follows from accuracy. Because trust is a relational rather than purely cognitive evaluation, it is unlikely to emerge automatically from technically accurate recommendations. Firms should support trust formation through transparency about how recommendations are generated, consistency in recommendation quality, meaningful consumer control over personalization settings, and responsible data practices (Walker, 2024).

4. Avoid personalization that outpaces perceived relevance. Given the personalization paradox, increasing the intensity or frequency of personalized contact without a corresponding increase in perceived relevance risks shifting consumer interpretation from "helpful" toward "surveilling," undermining rather than building the trust the framework identifies as necessary for purchase intention (Aguirre et al., 2015; Cai & Mardani, 2023).

5. Design AI for consumer decision support rather than behavioral manipulation. The framework's emphasis on relevance and trust as sequential preconditions for purchase intention implies that AI systems designed primarily to maximize short-term conversion, at the expense of genuine alignment with consumer need, are likely to be self-limiting: they may generate short-term response without the trust required to sustain longer-term consumer relationships.

9. Boundary Conditions and Future Research Agenda

The framework developed here is deliberately parsimonious, and its applicability is likely to be conditioned by several consumer-level, technology-level, and contextual factors that constitute a natural agenda for future research.

Consumer-level factors. AI literacy may condition how consumers interpret personalized outputs, since less AI-literate consumers may evaluate relevance and system competence using different cues than more literate consumers. Privacy concern is a plausible moderator of the personalization–relevance relationship, consistent with the personalization–privacy paradox (Aguirre et al., 2015; Cai & Mardani, 2023). Technology readiness, consumer involvement, and individual need for personalization are additional candidate moderators of the strength of the proposed relationships.

Technology-level factors. AI transparency and explainability may influence whether consumers interpret personalized outputs as helpful or manipulative, and may therefore moderate the relevance–trust relationship. Recommendation accuracy, the degree of personalization, and the presence or absence of anthropomorphic design cues are further candidates for future investigation.

Contextual factors. The proposed relationships may vary by product type (hedonic versus utilitarian, high- versus low-involvement), by the level of perceived purchase risk, by platform type, and by cultural context, given that trust formation and privacy sensitivity are known to vary across cultural settings.

Potential future relationships. Future research could examine moderated relationships such as AI Personalization × Privacy Concern → Perceived Relevance, and Perceived Relevance × AI Transparency → Consumer Trust. Comparative research contrasting human-generated and AI-generated personalization would help isolate which elements of the proposed mechanism are specific to algorithmic sources. Longitudinal research is particularly warranted to examine how trust in AI-enabled personalization develops, and how it may erode following relational breakdowns such as those documented by Saenger et al. (2024).

While the present paper is conceptual, the model is designed for empirical testing. A cross-sectional survey of consumers who have encountered AI-powered recommendations, advertisements, or digital suggestions, using a combination of purposive and convenience sampling, would provide an appropriate initial empirical design. A sample of approximately 250–400 respondents would support structural equation modeling, for example using a partial least squares approach, with assessment of indicator reliability, composite reliability, average variance extracted, and discriminant validity at the measurement-model stage, and path coefficients, bootstrapped indirect effects, R^2, f^2, and Q^2 at the structural-model stage. The sequential indirect effect central to Proposition 4—AI Personalization → Relevance → Trust → Purchase Intention—could be examined directly through bootstrapped mediation analysis.

10. Conclusion

This paper set out to answer a specific theoretical question: when does AI-powered personalization become human-centric? Drawing on the Stimulus–Organism–Response framework, the paper's answer is that AI-powered personalization becomes human-centric when algorithmic capability is translated, through consumer interpretation, into an experience that is perceived as relevant and sufficiently trustworthy to support decision-making. This translation is not automatic and cannot be assumed from technological sophistication alone; it depends on a sequential psychological process in which relevance functions as evidence of system competence, trust functions as the relational judgment that evidence supports, and purchase intention functions as the behavioral outcome of that judgment.

The paper's contribution is therefore conceptual rather than technical: it does not propose a better algorithm, but a better account of what makes an algorithm's outputs matter to the people who receive them. As AI systems continue to expand their capacity to infer and act upon consumer preferences, the more consequential determinant of their marketing effectiveness may not be how much a system knows about a consumer, but how effectively that consumer perceives the system to understand and serve them. Future empirical work testing the propositions developed here would help establish whether, and under what boundary conditions, this sequential mechanism holds.

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. An, G. K., & Ngo, T. T. A. (2025). AI-powered personalized advertising and purchase intention in Vietnam's digital landscape: The role of trust, relevance, and usefulness. Journal of Open Innovation: Technology, Market, and Complexity, 11(3), 100580. https://doi.org/10.1016/j.joitmc.2025.100580
3. Breque, M., De Nul, L., & Petridis, A. (2021). Industry 5.0: Towards a sustainable, human-centric and resilient European industry. European Commission.
4. Cai, H., & Mardani, A. (2023). Research on the impact of consumer privacy and intelligent personalization technology on purchase resistance. Journal of Business Research, 113811.
5. Chakraborty, D., Kar, A. K., Patre, S., & Gupta, S. (2024). Enhancing trust in online grocery shopping through generative AI chatbots. Journal of Business Research, 180, 114737. https://doi.org/10.1016/j.jbusres.2024.114737
6. Gao, Y., & Liu, H. (2023). Artificial intelligence-enabled personalization in interactive marketing: A customer journey perspective. Journal of Research in Interactive Marketing, 17(5), 663–680. https://doi.org/10.1108/JRIM-01-2022-0023
7. Hermann, E., & Puntoni, S. (2024). Artificial intelligence and consumer behavior: From predictive to generative AI. Journal of Business Research, 180, 114720.
8. Hermann, E., Williams, G. Y., & Puntoni, S. (2024). Deploying artificial intelligence in services to aid vulnerable consumers. Journal of the Academy of Marketing Science, 52, 1431–1451.
9. Huang, M.-H., & Rust, R. T. (2021). A strategic framework for artificial intelligence in marketing. Journal of the Academy of Marketing Science, 49, 30–50. https://doi.org/10.1007/s11747-020-00749-9
10. Kirk, C. P., & Givi, J. (2024). The AI-authorship effect: Understanding authenticity, moral disgust, and consumer responses to AI-generated marketing communications. Journal of Business Research, 114984.
11. Puntoni, S., Reczek, R. W., Giesler, M., & Botti, S. (2021). Consumers and artificial intelligence: An experiential perspective. Journal of Marketing, 85(1), 131–151. https://doi.org/10.1177/0022242920953847
12. Saenger, C., Kuchmaner, C. A., & Bateman, P. J. (2024). Betrayed by AI: How perceived betrayal by a virtual assistant affects consumers' purchase intentions for recommended products. Journal of Business Research, 114940.
13. Walker, K. L. (2024). AI-driven technology and privacy: The value of social media responsibility. Journal of Research in Interactive Marketing, 18(5), 815–835.

Table A: Research Gap and Conceptual Contribution

Existing Knowledge	Limitation	Proposed Advancement
AI in marketing has been extensively studied as a strategic and technical capability (Huang & Rust, 2021).	Capability-focused accounts do not explain when or why AI personalization succeeds or fails from the consumer's perspective.	Reframes AI personalization as a stimulus whose value depends on consumer interpretation, not technical sophistication.
The personalization paradox shows that personalization can increase both relevance and vulnerability (Aguirre et al., 2015).	The paradox identifies competing outcomes but does not specify the psychological sequence linking personalization to behavioral response.	Positions relevance and trust as sequential, ordered evaluations rather than competing or parallel outcomes.
Recent studies show relevance, usefulness, and trust jointly predict purchase intention in AI-powered advertising (An & Ngo, 2025).	These outcomes are typically modeled as parallel predictors rather than as a theoretically ordered mechanism.	Establishes relevance as theoretically prior to trust, based on the S-O-R logic of cognitive evaluation preceding relational evaluation.
Consumer-AI research emphasizes experiential and psychological consequences of AI (Puntoni et al., 2021; Hermann & Puntoni, 2024).	This experiential perspective has not been integrated with the human-centric marketing and Industry 5.0 discourse in relation to personalization specifically.	Extends human-centric marketing into AI-powered personalization by operationalizing human-centricity through relevance and trust.

Table B: Summary of Research Propositions

Proposition	Relationship	Theoretical Logic
P1	AI-powered personalization → Perceived relevance (positive)	Personalization is a technological capability; relevance is the consumer's subjective evaluation of whether that capability's output meets their needs.
P2	Perceived relevance → Consumer trust (positive)	Relevant experiences provide experiential evidence of system competence and responsiveness, which reduces uncertainty and supports trust formation.
P3	Consumer trust → Purchase intention (positive)	Trust reduces the uncertainty inherent in AI-mediated recommendations and digital purchasing, increasing willingness to act.
P4	AI-powered personalization → Perceived relevance → Consumer trust → Purchase intention (sequential mediation)	Human-centricity emerges through a sequential, ordered process of consumer interpretation rather than as a direct effect of technological sophistication.