

Paper Type: Original Article

When Personalization Feels Like Surveillance: A Marketing Framework for Artificial Intelligence-Personalized Advertising

Amirhossein Nafei* 

Department of Industrial Engineering and Management, National Taipei University of Technology, Taipei, Taiwan; amir.nafei@gmail.ntut.edu.tw.

Citation:

Received: 02 January 2026

Revised: 14 April 2026

Accepted: 19 May 2026

Nafei, A. (2026). When personalization feels like surveillance: A marketing framework for artificial intelligence-personalized advertising. *Systemic Analytics*, 4(2), 102-117.

Abstract

Artificial Intelligence (AI)-personalized advertising can increase message relevance, consumer engagement, and persuasive effectiveness, yet the same targeting practices can expose the extent of consumer tracking and trigger concerns about surveillance, manipulation, and loss of autonomy. Existing research documents both the benefits and the risks of personalization, but it does not fully explain how personalization intensity, data-use transparency, and meaningful consumer control jointly shape advertising evaluations and brand-level outcomes. Through an integrative synthesis of personalized advertising, privacy-calculus theory, psychological reactance, procedural fairness, and brand-trust research, this conceptual article develops a dual-pathway framework. The value-creation pathway predicts that high-quality personalization strengthens perceived relevance and advertising value, which subsequently enhance purchase intention, brand trust, and continued engagement. The surveillance-threat pathway predicts that unexpected data use, sensitive inferences, excessive specificity, and opaque targeting practices increase perceived surveillance and intrusiveness, thereby stimulating reactance, advertising avoidance, and trust erosion. The framework distinguishes transparency from control and argues that these mechanisms become most effective when combined: transparency clarifies how and why personal data are used, whereas control allows consumers to modify, limit, or refuse personalization. Their joint presence creates procedural legitimacy and can preserve the benefits of personalization while reducing perceived vulnerability. Twelve propositions specify the principal mediating mechanisms, boundary conditions, and an inverted-U relationship between personalization intensity and net marketing response. The article also proposes a three-study empirical validation agenda and practical design principles for developing personalized advertising systems that improve relevance and effectiveness without undermining consumer autonomy, privacy expectations, or long-term brand trust.

Keywords: Artificial intelligence, Personalized advertising, Consumer privacy, Perceived intrusiveness, Brand trust, Transparency, Consumer control.

1 | Introduction

Artificial Intelligence (AI) is changing digital marketing from segment-level targeting toward continuous, individual-level prediction and message generation. Contemporary systems can infer consumer preferences from browsing, purchase, location, platform, and conversational data; select a persuasive appeal; generate

 Corresponding Author: amir.nafei@gmail.ntut.edu.tw

 <https://doi.org/10.31181/sa42202679>



Licensee System Analytics. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<http://creativecommons.org/licenses/by/4.0>).

copy and imagery; and optimize delivery in real time. Marketing research has therefore moved beyond treating AI as a back-office analytics tool and increasingly examines it as an active participant in targeting, customer interaction, and persuasion [1–4]. Within this transformation, AI-personalized advertising is especially consequential because it combines two powerful capabilities: the capacity to make messages more relevant and the capacity to make previously invisible data practices perceptible to consumers.

Online Behavioral Advertising (OBA) already monitors digital behavior to tailor advertising, but AI substantially increases the scale, speed, granularity, and inferential depth of personalization [5]. A conventional retargeting banner may display a product previously viewed by a consumer. An AI-personalized advertisement can go further by predicting latent needs, adapting emotional tone, selecting identity-relevant language, and producing individualized creative content. These capabilities create obvious commercial opportunities. Personalization can improve message fit, reduce search costs, and increase advertising effectiveness when the right content is delivered at the right stage and in the right context [6–9]. At the same time, the same advertisement can reveal that a firm knows more about the consumer than the consumer expected it to know.

This tension motivates the present article. The managerial problem is not simply whether personalization should be increased or reduced. The more difficult question is how to design AI-personalized advertising that creates genuine relevance without producing the feeling of being watched, inferred, or manipulated. Recent advertising research confirms that personalization can strengthen perceived relevance and brand response while intrusiveness limits or reverses these benefits [10–12]. Highly personalized messages can also activate reactance when consumers cannot understand why a particular personal attribute was used or when the value of the offer does not justify the depth of personalization [13], [14]. These mixed findings indicate that personalization is not a unitary treatment. Its effects depend on data provenance, message specificity, data sensitivity, timing, platform trust, and the consumer's ability to exercise meaningful control.

The practical importance of the problem is increasing as AI-generated advertising moves from product recommendations to adaptive persuasion. A consumer may reasonably expect a retailer to recommend shoes similar to a recently viewed pair. The same consumer may react negatively when an advertisement appears to infer financial insecurity, health concerns, political identity, or emotional vulnerability from unrelated digital behavior. In both cases, the message can be statistically relevant. Yet the second case violates expectations concerning appropriate information flows and threatens perceived autonomy. The distinction between useful prediction and unacceptable surveillance is therefore not captured by algorithmic accuracy alone.

The literature offers several partial explanations. Advertising-value research emphasizes informativeness and relevance [15]. Intrusiveness research explains how interruption and perceived interference generate avoidance [16], [17]. Psychological reactance theory explains resistance to threats against freedom [18], while the persuasion knowledge model explains how consumers recognize and cope with marketers' influence tactics [19]. Privacy research conceptualizes concerns about data collection, control, and awareness [20–25], and contextual integrity explains why information use becomes objectionable when it violates context-specific norms [26]. These perspectives are individually valuable, but they have not been fully integrated into a marketing account of AI-personalized advertising in which message relevance, surveillance perception, transparency, control, and brand trust operate simultaneously.

The central research question is therefore: Under what conditions does AI-personalized advertising create value through relevance, and under what conditions does it trigger a surveillance response that damages advertising effectiveness and brand trust? To answer this question, the article develops a dual-pathway framework. The value-creation pathway begins when personalization is experienced as useful, appropriate, and preference-consistent. Perceived relevance increases advertising value and supports favorable market responses. The surveillance-threat pathway begins when consumers infer excessive tracking, unexpected data combination, sensitive profiling, or opaque prediction. These appraisals increase perceived intrusiveness and psychological reactance, which encourage avoidance and weaken trust.

The main innovation is to treat data-use transparency and meaningful consumer control as distinct governance mechanisms rather than interchangeable privacy assurances. Transparency informs consumers about data sources, purposes, and inferences, but information alone is not always reassuring. Disclosure can backfire when it reveals information flows that consumers regard as unacceptable [27]. Control is different: it provides the capacity to inspect, change, refuse, or reset personalization. Evidence from personalized advertising and data-privacy research indicates that control can improve responses even when the underlying targeting technology does not change [8], [28–30]. The framework argues that transparency without control may create awareness without remedy, whereas transparency combined with meaningful control can establish procedural legitimacy.

The article makes four contributions. First, it integrates the positive relevance pathway and the negative surveillance pathway into one marketing framework, thereby explaining why similar personalization practices can produce opposite responses. Second, it distinguishes personalization intensity from personalization quality. More data and greater specificity do not necessarily create greater value; the relevant issue is whether the message improves decision usefulness while respecting expected information boundaries. Third, it theorizes transparency and control as complementary but conditionally effective mechanisms. Fourth, it proposes a non-linear relationship between AI-personalization intensity and net marketing response: performance should improve as relevance increases, peak in a responsible-personalization zone, and decline when surveillance, intrusiveness, and reactance become dominant.

The article is organized into six sections. Section 2 reviews AI-personalized advertising, advertising value, privacy calculus, intrusiveness, reactance, transparency, control, and trust. Section 3 develops the conceptual framework and propositions. Section 4 derives managerial implications for digital advertising strategy and platform design. Section 5 presents a cumulative empirical validation agenda. Section 6 concludes with the theoretical contribution, limitations, and future research directions.

2 | Literature Review and Theoretical Foundations

2.1 | Conceptual Synthesis Approach

This article uses an integrative theory-synthesis approach. It does not claim to be a systematic review and does not report empirical findings. Conceptual research contributes by connecting established theories, specifying mechanisms, resolving tensions, and identifying boundary conditions that are not visible within isolated streams [31], [32]. Consistent with guidance for rigorous literature synthesis, the review prioritizes peer-reviewed studies that directly define the focal constructs or establish the relationships needed to explain AI-personalized advertising [33]. Five streams are integrated: AI and personalized advertising, relevance and advertising value, privacy and surveillance appraisal, intrusiveness and reactance, and transparency, control, and brand trust.

2.2 | Artificial Intelligence and Personalized Advertising

AI in marketing supports data collection, prediction, segmentation, targeting, creative generation, service interaction, and decision automation [1–4]. In advertising, personalization can vary in depth, breadth, timing, and context. Depth refers to how specifically a message reflects a consumer's behavior or characteristics, whereas breadth refers to the range of data or categories incorporated into the message [6], [34]. Timing determines whether the advertisement appears when the consumer's need is active or after the personalization has become stale. Placement determines whether the advertisement appears in a context that supports or conflicts with the consumer's current goal [6].

The literature consistently rejects a simple “more personalization is better” assumption. Personalization can improve effectiveness, but high specificity may decay quickly, become intrusive, or produce reactance when it is not justified [6], [13]. Targeting and visual obtrusiveness can each improve performance in isolation yet undermine one another when combined [9]. Social-media evidence likewise shows that relevance can mediate

favorable brand responses while perceived intrusiveness weakens the relevance effect [10]. These findings imply that AI should be evaluated not only by predictive accuracy but also by the consumer's interpretation of how the prediction was produced and why the message was delivered.

AI adds a further layer because the system may infer attributes that the consumer never explicitly disclosed. Consumers can appreciate algorithms in objective tasks but resist them in subjective or identity-sensitive tasks [35], [36]. Disclosure of an AI agent can also change purchase behavior even when the underlying service remains competent [37]. In personalized advertising, AI awareness may therefore influence both perceived usefulness and perceived manipulation. The technology's role is not neutral: It can signal sophisticated assistance, but it can also make the persuasive infrastructure feel more autonomous and less accountable.

2.3 | The Value-Creation Pathway: Relevance and Advertising Value

Advertising value refers to consumers' evaluation of the relative utility or worth of advertising, with informativeness and entertainment traditionally operating as positive antecedents [15]. Personalized advertising can increase value by reducing irrelevant exposure, matching the message to current goals, and helping consumers identify suitable products. When consumers interpret personalization as a service rather than an extraction practice, message specificity can improve perceived relevance, attitude toward the advertisement, click intention, and purchase response [6], [10–12].

The relevance mechanism is fundamentally relational. A message is relevant when the consumer can connect its content to an active need, preference, or decision. Relevance is therefore not equivalent to the amount of personal data embedded in an advertisement. A moderately personalized advertisement based on a recent on-site search may be more relevant than a deeply personalized message derived from distant or sensitive data. The framework accordingly distinguishes personalization intensity, the extent of individualized data and message adaptation, from personalization quality, the degree to which adaptation is useful, timely, accurate, and norm-consistent.

The distinction helps explain overpersonalization. High-depth advertisements can initially attract attention but become less effective as time passes or as the consumer's decision state changes [6]. Similarly, perceived personalization can strengthen brand responses through relevance, yet the same message can produce weaker responses when it appears intrusive [10–12]. The positive pathway therefore requires more than technical matching. It requires consumer-recognizable benefit.

2.4 | Privacy Calculus, Information Boundaries, and Surveillance Appraisal

Information-privacy research portrays consumer response as a context-dependent evaluation of benefits, risks, trust, and control. Internet users' information privacy concerns include collection, control, and awareness [20]. Consumers differ in the types of information they are willing to disclose, and greater control over collection and dissemination can reduce resistance [21]. Privacy-calculus models further propose that consumers may accept information disclosure when perceived benefits and trust outweigh expected risks [22]. However, the calculus is imperfect because preferences are unstable, consequences are uncertain, and privacy judgments are strongly shaped by context [24], [25].

The personalization–privacy paradox emerges because the same data that improve service relevance also create vulnerability. Transparency can increase understanding, but privacy-sensitive consumers may remain unwilling to be profiled even when information practices are disclosed [23]. Privacy-safe personalization can improve engagement by retaining consumer information locally or reducing boundary intrusion [28]. These findings show that privacy is not merely a cost to be traded for relevance; it is also a judgment about whether information movement is appropriate.

Contextual integrity provides a useful explanation. Information use is experienced as legitimate when it conforms to expectations about the actors, attributes, purposes, and transmission principles governing a context [26]. In advertising, consumers are generally more accepting of information they directly provided on

the same platform than of information inferred or obtained across contexts [27]. AI personalization increases the risk of violating contextual integrity because models can combine weak signals, infer sensitive attributes, and produce messages whose relevance reveals an unexpected information flow.

Perceived surveillance is the consumer's appraisal that the firm or platform has continuously monitored, combined, or inferred personal information beyond what was expected. This perception is related to, but distinct from, abstract privacy concern. Privacy concern is a relatively general orientation; surveillance perception is a proximal interpretation of a specific advertisement. A transparency message can therefore increase both perceived transparency and perceived surveillance when it clarifies how extensive the targeting process has been.

2.5 | Intrusiveness, Psychological Reactance, and Advertising Avoidance

Perceived intrusiveness captures the extent to which an advertisement interrupts, interferes with, or invades a consumer's cognitive or personal space [16]. Intrusiveness increases when the advertisement disrupts an ongoing task, is overly intense, or appears to use personal information in an unwelcome manner [17]. Advertising avoidance can then take behavioral, cognitive, or affective forms, including ignoring the message, closing the advertisement, restricting data access, leaving the platform, or developing skepticism toward the advertiser [14].

Psychological reactance explains why these responses can occur even when the advertised product is relevant. Reactance is a motivational state triggered by a perceived threat to freedom [18]. Highly personalized advertising can threaten freedom by making consumers feel categorized, manipulated, or unable to escape tracking. The persuasion knowledge model adds that consumers draw on knowledge about marketing motives and tactics to interpret persuasion attempts [19]. AI-personalized advertising can make the tactic unusually salient: The advertisement itself becomes evidence that the marketer has profiled the consumer.

Intrusiveness and reactance are related but not identical. Intrusiveness is an appraisal of the advertisement or data practice, whereas reactance is the motivational response to perceived interference or control. A consumer may find an advertisement intrusive without actively resisting it when the message provides substantial value. Conversely, a consumer may resist a relatively unobtrusive message if the underlying data use is judged manipulative. The framework therefore models surveillance and intrusiveness as proximal appraisals that activate reactance and avoidance when personalization is not legitimized by relevance, transparency, and control.

2.6 | Transparency, Meaningful Control, Procedural Fairness, and Brand Trust

Transparency is often proposed as a universal remedy for data-driven marketing, but its effects are conditional. Ad-transparency disclosures can improve effectiveness when they reveal acceptable information flows on trusted platforms; they can backfire when they expose cross-context tracking or inferred attributes that violate consumer norms [18]. The visibility and framing of privacy information can also influence perceived intrusiveness and ethicality [39], while privacy policies can increase both transparency and surveillance perceptions. Transparency therefore does not automatically create trust. It changes what consumers know, and that knowledge can be reassuring or alarming.

Meaningful control concerns whether consumers can influence the collection, use, and consequences of their data. Field evidence shows that increasing perceived privacy control can nearly double responsiveness to personalized advertising even when advertisers' actual targeting capabilities remain unchanged [8]. Data-privacy research similarly finds that transparency and control can suppress vulnerability and violation effects [30]. Control is also relevant to algorithm use more broadly: People become more willing to rely on imperfect algorithms when they can modify outcomes, even slightly [40].

The combination of transparency and control is theoretically important. Transparency without control may produce informed powerlessness. Control without transparency may produce superficial choice because

consumers do not know what they are controlling. When both are present, consumers can understand the information flow and exercise agency over it. This combination should strengthen procedural fairness, the belief that the process governing personalization is understandable, respectful, and correctable. Recent targeting research further demonstrates that perceived relevance and consumer controllability are central to whether targeting practices are judged fair and worthy of brand support [41].

Brand trust is the willingness to rely on a brand under conditions of vulnerability. Trust supports purchase loyalty and attitudinal loyalty and can affect brand performance [42]. In personalized advertising, trust operates both as an antecedent and an outcome. Trusted retailers can use personalization more effectively [34], while opaque or norm-violating data practices can erode trust [27], [30]. AI-specific trust and literacy may also affect whether personalization is interpreted as competent assistance or as invasive automation [43], [44].

Table 1. Literature synthesis and focal research gap.

Stream	Established Insight	Unresolved Issue in AI-Personalized Advertising	Framework Contribution
Personalized ads	Targeted messages can improve relevance and effectiveness, but effects depend on depth, timing, and context [6–12].	Predictive accuracy does not explain why a technically relevant ad can feel unacceptable.	Separates targeting intensity from targeting quality and consumer-perceived legitimacy.
Privacy boundaries	Consumers evaluate collection, control, awareness, benefits, risk, and context [20–26], [28].	General privacy concern does not capture the advertisement-specific feeling of being watched.	Introduces perceived surveillance as a proximal appraisal of data-driven targeting.
Reactance	Intrusiveness and threats to freedom can produce resistance and avoidance [13-19].	Negative reactions are often modeled without the competing relevance pathway.	Integrates intrusiveness and reactance with relevance and advertising value.
Disclosure/control	Transparency can help or backfire; control can reduce vulnerability and improve response [8], [27], [30], [40].	Transparency and control are frequently treated as equivalent privacy assurances.	Models transparency and meaningful control as distinct, interacting governance mechanisms.
Brand trust	Trust shapes responses to targeting and technology-mediated persuasion [34-37], [40], [42-45].	It is unclear when trust is protected versus depleted by AI-personalized advertising.	Positions brand trust as both a boundary condition and a relational outcome.

3 | Conceptual Framework and Propositions

3.1 | The Dual-Pathway Structure

The proposed framework treats AI-personalized advertising as a consumer appraisal process rather than a direct technological effect. The same system configuration can generate different responses because consumers interpret both the benefit and the legitimacy of personalization. Three advertising-design inputs are central: personalization intensity, data-use transparency, and meaningful consumer control. These inputs shape a value-creation pathway and a surveillance-threat pathway, which converge on advertising outcomes and brand trust.

Personalization intensity refers to the extent to which an advertisement is adapted to an individual through personal data, inferred attributes, message specificity, dynamic creative generation, and delivery optimization. Data-use transparency refers to the clarity and specificity with which the firm explains the source, purpose, inference, and use of personal information. Meaningful consumer control refers to the ability to inspect,

modify, refuse, delete, or reset the data and personalization settings that affect advertising. *Fig. 1* presents the complete framework.




Fig. 1. Dual-pathway marketing framework of AI-personalized advertising.

3.2 | The Value-Creation Pathway

Personalization can create value when it helps consumers recognize that an advertisement addresses a current need. AI can increase usefulness by combining product information with observed preferences, adapting the message to the consumer's decision stage, and reducing exposure to irrelevant alternatives. The first psychological translation of these capabilities is perceived relevance, not purchase intention. Consumers must first conclude that the message is connected to their goals in a useful and appropriate manner.

Proposition 1. Higher personalization quality increases perceived advertising relevance, whereas personalization intensity has a weaker positive effect when quality is held constant.

Perceived relevance should increase advertising value because consumers receive useful information with lower search and evaluation costs. Advertising value then supports favorable attitudes and behavioral intentions. The effect should extend beyond immediate clicks because a relevant advertisement can signal that the firm understands the consumer's needs. However, this relational inference depends on whether personalization is experienced as service rather than surveillance.

Proposition 2. Perceived relevance positively influences advertising value, which in turn increases advertising attitude, click intention, and purchase intention.

Personalization can also strengthen brand trust when relevance is delivered consistently and without unexpected data use. A useful message can communicate competence and benevolence: the brand appears capable of understanding needs and willing to reduce consumer effort. This pathway is most plausible when the source of the personalization is obvious, the message is accurate, and the consumer can connect the advertisement to an intentional action such as a search, saved preference, or prior purchase.

Proposition 3. Advertising value positively influences brand trust when consumers perceive the personalization as accurate, timely, and preference-consistent.

3.3 | The Surveillance-Threat Pathway

Personalization can simultaneously reveal the existence of tracking and inference. Consumers may ask how the advertiser knew a preference, whether information moved across platforms, and whether the message exposes a sensitive or identity-relevant attribute. These questions generate perceived surveillance when the

data practice exceeds expected information boundaries. The effect should depend less on the objective number of data points than on the psychological distance between the consumer's action and the advertisement.

Proposition 4. Personalization intensity increases perceived surveillance when the advertisement relies on inferred, cross-context, or sensitive information rather than consumer-stated, same-context information.

Perceived surveillance should increase intrusiveness because the advertisement is no longer evaluated only as a message; it is evaluated as evidence of monitoring. Intrusiveness is expected to be especially strong when personalization is surprising, when the source of the data is unclear, or when the message appears immediately after a private or identity-sensitive activity. In these conditions, even a relevant advertisement can feel invasive.

Proposition 5. Perceived surveillance positively influences perceived advertising intrusiveness, with the effect strengthened by data sensitivity and information-source distance.

Intrusiveness should activate psychological reactance when consumers interpret the targeting practice as a threat to autonomy. Reactance shifts attention from the advertised benefit to the persuasive tactic. Consumers may then avoid the advertisement, reject the recommendation, change privacy settings, or punish the brand. This pathway explains why improving message accuracy can reduce performance if accuracy reveals an unacceptable degree of tracking.

Proposition 6. Perceived intrusiveness increases psychological reactance, which reduces advertising attitude, click intention, purchase intention, and willingness to accept future personalization.

The surveillance-threat pathway also has relational consequences. Data practices create vulnerability because consumers cannot fully observe how information is combined or reused. When an advertisement reveals an unexpected inference, the consumer may perceive the firm as opportunistic or manipulative. Trust should decline even when the advertised product is suitable, because the consumer evaluates the integrity of the process rather than only the quality of the recommendation.

Proposition 7. Perceived surveillance and psychological reactance reduce brand trust, and brand-trust erosion partially mediates their negative effects on purchase response and data-sharing willingness.

3.4 | Transparency and Control as Legitimacy Mechanisms

Transparency can reduce uncertainty by explaining why the consumer is seeing an advertisement and which information was used. Yet transparency is not uniformly positive. A disclosure that confirms an acceptable flow, such as the use of a recent search on the same retailer's website, can increase comprehension and relevance. A disclosure that reveals inferred financial status or cross-platform tracking can make the surveillance pathway more salient. Transparency, therefore, amplifies the consumer's evaluation of the underlying practice rather than automatically improving it.

Proposition 8. Data-use transparency increases perceived procedural fairness when it reveals expected and bounded information flows, but increases perceived surveillance when it reveals norm-violating or unexpectedly sensitive information flows.

Meaningful control provides a different form of assurance. Consumers should feel less vulnerable when they can view why an advertisement was selected, remove an inferred interest, restrict sensitive categories, pause personalization, or reset the advertising profile. These features convert personalization from an imposed process into a manageable service. Token controls, however, should have little effect when they are difficult to locate, do not alter targeting, or require excessive effort.

Proposition 9. Meaningful consumer control reduces perceived intrusiveness and psychological reactance while increasing procedural fairness and brand trust.

Transparency and control should interact. Transparency without control can create informed powerlessness: the consumer understands the surveillance practice but cannot change it. Control without transparency can

create uninformed choice: the consumer can toggle settings but does not know what the settings govern. Their joint presence creates procedural legitimacy because consumers can understand, evaluate, and correct the personalization process.

Proposition 10. Transparency and meaningful control have a positive interaction such that transparency improves brand trust and advertising response more strongly when consumers possess effective control over personalization.

3.5 | The Responsible-Personalization Zone and Boundary Conditions

The combined framework predicts a non-linear relationship between personalization intensity and net marketing response. At very low personalization, advertisements provide limited relevance and little differentiated value. As personalization increases, relevance and advertising value improve. Beyond a boundary-sensitive point, however, additional data depth and message specificity yield diminishing relevance while increasing surveillance and intrusiveness. The resulting curve should peak in a responsible-personalization zone where the message is useful, the information flow is expected, and the consumer retains agency. *Fig. 2* illustrates this prediction.




Fig. 2. Predicted non-linear relationship between AI personalization intensity and net marketing response.

Proposition 11. The relationship between personalization intensity and net marketing response is inverted-U shaped, with the highest response occurring when relevance gains are substantial but perceived surveillance remains low.

Several boundary conditions determine the location and steepness of the curve. Data sensitivity should move the backfire point to lower levels of personalization because consumers attach greater risk to health, finance, location, identity, and intimate-life information. Cross-context data use should have a similar effect. Brand and platform trust should delay the backfire point by reducing opportunistic inferences, although trust cannot legitimize every information flow. Individual privacy concern should strengthen the negative pathway. AI literacy may have mixed effects: knowledgeable consumers may understand legitimate system functions, but they may also recognize inference and tracking more accurately. Product involvement should strengthen the relevance pathway when consumers actively seek information, while identity-sensitive products may increase surveillance concerns when personalization appears to infer private traits.

Proposition 12. The negative effects of personalization intensity are stronger for sensitive data, cross-context information flows, high privacy concern, and low brand or platform trust, but weaker when personalization quality, consumer involvement, and meaningful control are high.

Table 2. Construct definitions and operational guidance.

Construct	Conceptual Definition	Illustrative Manipulation or Measure	Expected Role
Intensity	Extent of individual adaptation through data depth, inference, creative specificity, and delivery timing.	Manipulate generic, moderate, and high-depth AI-personalized advertisements.	Exogenous advertising-design variable.
Quality	Usefulness, accuracy, timeliness, and preference fit of the adaptation.	Measure relevance, accuracy, usefulness, and decision support.	Positive-pathway input.
Surveillance	Ad-specific inference that the firm monitored or inferred more than expected.	Items assessing tracking awareness, unexpected knowledge, and feeling watched.	First-stage threat appraisal.
Intrusiveness	Extent to which an ad invades, interferes with, or disrupts personal and cognitive space.	Adapt the validated advertising-intrusiveness scale [16].	Threat-pathway mediator.
Reactance	Motivational resistance generated by a perceived threat to freedom or autonomy.	Measure irritation, resistance, counterarguing, and desire to restore freedom.	Behavioral-resistance mediator.
Transparency	Clarity concerning data source, purpose, inference, and use in ad delivery.	Manipulate, no explanation, vague explanation, and specific information-flow disclosure.	Conditional legitimacy mechanism.
Consumer control	Capacity to inspect, modify, refuse, delete, or reset ad settings.	Manipulate editable interests, category exclusions, opt-out, and profile reset.	Agency and legitimacy mechanism.
Brand trust	Willingness to rely on a brand when data use creates vulnerability	Adapt established brand-trust measures [42].	Relational mediator and outcome.
Market response	Attitudinal and behavioral outcomes of the ad encounter.	Ad attitude, click, product choice, purchase intention, willingness to pay, data sharing, and future personalization acceptance.	Dependent outcomes.

4 | Managerial Implications for Digital Marketing and Advertising

4.1 | Optimize Personalization Quality, Not Data Volume

Advertising managers should separate the question “How much can we personalize?” from “How much personalization creates recognizable consumer value?” AI systems make it inexpensive to incorporate more attributes into a message, but data volume is not a consumer benefit. Managers should prioritize variables that improve choice quality, timing, and usefulness. A recent on-site search, a stated preference, or a saved product may justify personalization more clearly than an inferred psychological profile. The goal should be high explanatory relevance with minimum necessary data.

Performance dashboards should therefore compare personalization intensity and personalization quality. A campaign can be technically sophisticated yet low quality if the message is inaccurate, stale, or disconnected from the consumer’s current goal. Managers should monitor not only Click-Through Rate (CTR) and conversion, but also perceived relevance, intrusiveness, opt-out behavior, profile corrections, brand trust, and post-exposure avoidance. A short-term lift that produces long-term trust erosion is not an efficient outcome.

4.2 | Design Transparency around Information Flows

Transparency should answer four practical questions: What information was used? Where did it come from? What was inferred? Why did that information affect this advertisement? Generic statements such as “shown based on your activity” provide little procedural value. More specific statements can be useful, but only when the underlying practice is defensible. Firms should not expect disclosure language to repair an unacceptable data flow. Transparency is most effective when it makes a bounded and consumer-recognizable process legible. A contribution-style explanation can show the direct link between consumer action and advertisement delivery: “You viewed running shoes on this site, so we are showing related products.” For inferred attributes, the platform should identify the inference and allow correction: “We estimated that you may be interested in home fitness. Change this interest.” This design transforms transparency from a passive notice into an actionable explanation.

4.3 | Provide Control at the Moment of Advertising

Privacy controls are often buried in account menus, separated from the advertisement that activates concern. The framework recommends in-flow control: consumers should be able to inspect and modify personalization at the point of exposure. Useful controls include removing an interest, excluding sensitive categories, choosing contextual rather than behavioral advertising, pausing off-platform tracking, and resetting the advertising profile. The control must produce a visible consequence; otherwise, it becomes symbolic rather than meaningful. *Fig. 3* organizes transparency and control into four managerial states. Opaque targeting combines low understanding with low agency and creates high suspicion risk. Explained-but-powerless targeting can intensify surveillance salience because consumers learn about data use without receiving a remedy. Unexplained user management provides some autonomy but leaves uncertainty about what is being managed. Responsible personalization combines clear information flows with granular, effective control.




Fig. 3. Transparency-control matrix for responsible AI-personalized advertising.

4.4 | Segment Governance by Data and Product Sensitivity

A uniform personalization policy is unlikely to be appropriate across all campaigns. Managers should classify data according to sensitivity, source distance, and inferential depth. Low-sensitivity, same-context data can support deeper personalization with relatively light governance. Sensitive, cross-context, or inferred data require stronger justification, explicit control, and potentially a prohibition on individual targeting. Contextual

advertising may be preferable when the relevance benefit of profiling is modest relative to the expected privacy cost.

Product category also matters. Consumers may welcome strong personalization for utilitarian search tasks when it reduces effort. The same targeting depth may feel manipulative for products connected to health, finance, identity, family, or vulnerability. Governance rules should therefore be campaign-specific and should be tested with consumers before deployment.

4.5 | Treat Trust as an Advertising Performance Asset

Brand trust should be incorporated into media and personalization decisions rather than treated as a general corporate-reputation variable. Trusted brands may receive greater latitude, but using that latitude aggressively can deplete the very trust that supports personalization. Managers should estimate the long-term relational effect of data practices, including opt-outs, complaint rates, profile deletions, and changes in willingness to share information. Responsible personalization is not only an ethical safeguard; it is a strategy for preserving the data relationship on which future advertising effectiveness depends.

5 | Empirical Validation Agenda

The propositions are designed for empirical testing and are not presented as established effects. A cumulative research program should separate causal mechanisms, boundary conditions, and behavioral consequences. The studies should be preregistered, use realistic advertising interfaces, record manipulation checks, and include behavioral outcomes where feasible. *Table 3* summarizes the proposed program.

5.1 | Study 1: Personalization Quality and the Dual Pathway

Study 1 should test whether personalization intensity and quality exert distinct effects. A factorial experiment can manipulate personalization intensity at three levels and personalization quality at two levels. Participants could view a social-media or retail-media advertisement for a familiar product category. High-quality personalization would connect the message to a current, explicitly stated need; low-quality personalization would use equally specific but weakly diagnostic information. Primary mediators would be perceived relevance, advertising value, perceived surveillance, and intrusiveness. Outcomes would include advertising attitude, click intention, product choice, and purchase intention.

A sample of approximately 600–750 participants would permit planned contrasts, parallel mediation, and tests of the predicted non-linear pattern while allowing exclusions for failed manipulation checks. The analysis should compare linear and quadratic personalization terms and should report whether relevance gains remain after surveillance and intrusiveness are controlled. Independent judges can assess message accuracy and creative quality to distinguish consumer appraisal from objective execution.

5.2 | Study 2: Transparency, Control, and Information-Flow Norms

Study 2 should manipulate data-use transparency, meaningful control, and information-flow acceptability. The transparency condition would explain the source and purpose of the targeting. The control condition would allow participants to remove the inferred interest or select contextual advertising. Information-flow acceptability would vary between same-site, consumer-stated data and cross-site, algorithmically inferred data. This design directly tests whether transparency helps acceptable practices but backfires for unacceptable practices, and whether control reduces the backfire.

The primary outcomes should include procedural fairness, surveillance perception, reactance, brand trust, willingness to accept future personalization, and actual use of the control feature. The study should distinguish the availability of control from control usage because the mere presence of an effective option may provide reassurance even when consumers do not exercise it. A moderated-mediation model can test whether transparency and control influence market response through fairness, surveillance, reactance, and trust.

5.3 | Study 3: Field Validation and Long-Term Relationship Effects

Study 3 should evaluate the responsible-personalization zone in a field or high-fidelity marketplace. A retailer or simulated platform could randomly vary personalization intensity and provide either standard privacy settings or advertisement-level transparency and control. Behavioral outcomes should include CTR, browsing depth, product choice, conversion, opt-out, profile correction, and return visits. A delayed follow-up should measure brand trust and willingness to share data after the immediate advertising encounter.

The field design should also examine heterogeneity. Privacy concern, AI literacy, product involvement, brand familiarity, and data sensitivity can be measured or manipulated. Latent-class or finite-mixture analysis may reveal segments that differ in their responsible-personalization threshold. Such evidence would help managers avoid a single global personalization level and instead choose governance settings that reflect consumer and campaign characteristics.

5.4 | Measurement and Analytical Safeguards

Measures should be adapted conservatively from established scales. Perceived intrusiveness can draw on the validated scale developed by Li et al. [16], privacy concern on the collection-control-awareness tradition [20], and brand trust on established marketing measures [42]. Perceived surveillance requires clear differentiation from general privacy concern, creepiness, and intrusiveness. Items should capture advertisement-specific inferences that the firm monitored, combined, or inferred more information than expected.

The empirical program should separate manipulations, mediators, and outcomes temporally where possible. It should include attention checks without making the study purpose obvious, preregister exclusions, and report alternative model specifications. Behavioral data should complement intentions. Researchers should archive advertisement stimuli, targeting explanations, platform interfaces, model versions, and prompt templates because AI-generated creative content can change across system updates. Replications should span cultures, platforms, and product categories.

Table 3. Proposed cumulative empirical validation program.

Study	Core Design	Primary Mechanisms	Behavioral Outcomes
Study 1: quality (P1–P7; P11)	Three personalization levels crossed with high versus low consumer-value fit.	Relevance, advertising value, surveillance, and intrusiveness.	Click, product choice, purchase intention, and ad avoidance.
Study 2: control design (P8–P10; P12)	Specific disclosure and actionable control crossed with acceptable versus norm- violating data use.	Procedural fairness, surveillance, reactance, and brand trust.	Control use, personalization acceptance, data sharing, and product choice.
Study 3: field (P11–P12; full model)	Live or high-fidelity campaign varying intensity and governance over time.	Responsible- personalization threshold and trust dynamics.	CTR, conversion, opt-out, profile correction, return visits, and delayed trust.

6 | Conclusion

AI-personalized advertising creates a distinctive marketing challenge because the same analytical capability can generate both relevance and surveillance. The conventional question whether personalized advertising is more effective than generic advertising is therefore incomplete. The more important question is how consumers interpret the information flow, the persuasive process, and their own agency when an advertisement appears to know them.

The dual-pathway framework explains the competing mechanisms. In the value-creation pathway, personalization quality increases perceived relevance and advertising value, supporting favorable attitudes,

clicks, purchase responses, and brand trust. In the surveillance-threat pathway, unexpected or sensitive data use increases perceived surveillance and intrusiveness, which activate psychological reactance, avoidance, and trust erosion. The framework further shows why transparency and control should not be treated as substitutes. Transparency can clarify or expose; control can restore agency. Their joint presence creates the strongest basis for procedural legitimacy.

The article contributes to digital marketing and advertising theory by distinguishing personalization intensity from personalization quality, introducing perceived surveillance as an advertisement-specific appraisal, and predicting an inverted-U relationship between personalization intensity and net marketing response. It also provides a managerial principle: the most advanced advertising system is not necessarily the system that can infer the most. It is the system that can create recognizable value with the least unnecessary data while keeping information flows understandable, bounded, and correctable.

This conceptual framework has limitations. Its propositions require empirical testing, and the responsible-personalization threshold may vary across cultures, platforms, product categories, regulatory environments, and consumer segments. Perceived surveillance also requires dedicated scale development and discriminant validation against privacy concern, intrusiveness, vulnerability, and creepiness. The framework focuses on consumer response and does not fully model advertiser competition, publisher incentives, legal compliance, or algorithmic discrimination. Future research should implement the proposed experiments and field studies, compare behavioral and contextual targeting, investigate conversational and generative advertising formats, examine long-term trust after repeated exposures, and determine how transparency and control can protect consumer autonomy without eliminating the relevance benefits of responsible personalization.

Acknowledgments

The author acknowledges no external assistance requiring formal recognition.

Author Contribution

Conceptualization, A.N.; methodology, A.N.; investigation, A.N.; writing original draft preparation, A.N.; writing review and editing, A.N.; visualization, A.N. The author has read and agreed to the submitted version of the manuscript.

Funding

This research received no external funding.

Data Availability

No new data were created or analyzed in this conceptual study. Data sharing is therefore not applicable.

Conflicts of Interest

The author declares no conflict of interest.

References

- [1] Davenport, T., Guha, A., Grewal, D., & Bressgott, T. (2020). How artificial intelligence will change the future of marketing. *Journal of the academy of marketing science*, 48(1), 24–42. <https://doi.org/10.1007/s11747-019-00696-0>
- [2] Huang, M. H., & Rust, R. T. (2021). A strategic framework for artificial intelligence in marketing. *Journal of the academy of marketing science*, 49(1), 30–50. <https://doi.org/10.1007/s11747-020-00749-9>
- [3] 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>

- [4] De Bruyn, A., Viswanathan, V., Beh, Y. S., Brock, J. K. U., & Von Wangenheim, F. (2020). Artificial intelligence and marketing: Pitfalls and opportunities. *Journal of interactive marketing*, 51(1), 91–105. <https://doi.org/10.1016/j.intmar.2020.04.007>
- [5] Boerman, S. C., Kruikemeier, S., & Borgesius, F. J. Z. (2017). Online behavioral advertising: A literature review and research Agenda. *Journal of advertising*, 46(3), 363–376. <https://doi.org/10.1080/00913367.2017.1339368>
- [6] Bleier, A., & Eisenbeiss, M. (2015). Personalized online advertising effectiveness: The interplay of what, when, and where. *Marketing science*, 34(5), 669–688. <https://doi.org/10.1287/mksc.2015.0930>
- [7] 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>
- [8] Tucker, C. E. (2014). Social networks, personalized advertising, and privacy controls. *Journal of marketing research*, 51(5), 546–562. <https://doi.org/10.1509/jmr.10.0355>
- [9] Goldfarb, A., & Tucker, C. (2011). Online display advertising: Targeting and obtrusiveness. *Marketing science*, 30(3), 389–404. <https://doi.org/10.1287/mksc.1100.0583>
- [10] De Groot, J. I. M. (2022). The personalization paradox in Facebook advertising: The mediating effect of relevance on the personalization--brand attitude relationship and the moderating effect of intrusiveness. *Journal of interactive advertising*, 22(1), 57–74. <https://doi.org/10.1080/15252019.2022.2032492>
- [11] De Keyser, F., Dens, N., & De Pelsmacker, P. (2022). How and when personalized advertising leads to brand attitude, click, and WOM intention. *Journal of advertising*, 51(1), 39–56. <https://doi.org/10.1080/00913367.2021.1888339>
- [12] Lee, J., Kim, C., & Lee, K. C. (2022). Exploring the personalization-intrusiveness-intention framework to evaluate the effects of personalization in social media. *International journal of information management*, 66, 102532. <https://doi.org/10.1016/j.ijinfomgt.2022.102532>
- [13] White, T. B., Zahay, D. L., Thorbjørnsen, H., & Shavitt, S. (2008). Getting too personal: Reactance to highly personalized email solicitations. *Marketing letters*, 19(1), 39–50. <https://doi.org/10.1007/s11002-007-9027-9>
- [14] Beak, T. H., & Morimoto, M. (2012). Stay away from me, examining the determinants of individual avoidance of personalized advertising. *Journal of advertising*, 41(1), 59–76. <https://doi.org/10.2753/JOA0091-3367410105>
- [15] Ducoffe, R. H. (1996). Advertising value and advertising on the web. *Journal of advertising research*, 36(5), 21–35. <https://doi.org/10.1080/00218499.1996.12466626>
- [16] Li, H., Edwards, S. M., & Lee, J. H. (2002). Measuring the intrusiveness of advertisements: Scale development and validation. *Journal of advertising*, 31(2), 37–47. <https://doi.org/10.1080/00913367.2002.10673665>
- [17] Edwards, S. M., Li, H., & Lee, J. H. (2002). Forced exposure and psychological reactance: Antecedents and consequences of the perceived intrusiveness of pop-up ads. *Journal of advertising*, 31(3), 83–95. <https://doi.org/10.1080/00913367.2002.10673678>
- [18] Clee, M. A., & Wicklund, R. A. (1980). Consumer behavior and psychological reactance. *Journal of consumer research*, 6(4), 389–405. <https://doi.org/10.1086/208782>
- [19] Friestad, M., & Wright, P. (1994). The persuasion knowledge model: How people cope with persuasion attempts. *Journal of consumer research*, 21(1), 1–31. <https://doi.org/10.1086/209380>
- [20] Malhotra, N. K., Kim, S. S., & Agarwal, J. (2004). Internet users' information privacy concerns (IUIPC): The construct, the scale, and a causal model. *Information systems research*, 15(4), 336–355. <https://doi.org/10.1287/isre.1040.0032>
- [21] Phelps, J., Nowak, G., & Ferrell, E. (2000). Privacy concerns and consumer willingness to provide personal information. *Journal of public policy & marketing*, 19(1), 27–41. <https://doi.org/10.1509/jppm.19.1.27.16941>
- [22] Dinev, T., & Hart, P. (2006). An extended privacy calculus model for e-commerce transactions. *Information systems research*, 17(1), 61–80. <https://doi.org/10.1287/isre.1060.0080>
- [23] 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 personalization1. *Management information systems quarterly*, 30(1), 13–28. <https://doi.org/10.2307/25148715>

- [24] Smith, H. J., Dinev, T., & Xu, H. (2011). Information privacy research: an interdisciplinary review1. *Management information systems quarterly*, 35(4), 989-A27. <https://doi.org/10.2307/41409970>
- [25] Acquisti, A., Brandimarte, L., & Loewenstein, G. (2015). Privacy and human behavior in the age of information. *Science*, 347(6221), 509–514. <https://doi.org/10.1126/science.aaa1465>
- [26] Nissenbaum, H. (2004). Privacy as contextual integrity. *Washington law review*, 79, 119. <https://access.heinonline.com/HOL/LandingPage?handle=hein.journals/washlr79&div=16&id=&page=>
- [27] Kim, T., Barasz, K., & John, L. K. (2019). Why am i seeing this Ad? the effect of Ad transparency on Ad effectiveness. *Journal of consumer research*, 45(5), 906–932. <https://doi.org/10.1093/jcr/ucy039>
- [28] Sutanto, J., Palme, E., Tan, C. H., & Phang, C. W. (2013). Addressing the personalization–privacy paradox: An empirical assessment from a field experiment on smartphone users1. *Management information systems quarterly*, 37(4), 1141–1164. <https://doi.org/10.25300/MISQ/2013/37.4.07>
- [29] Martin, K. D., & Murphy, P. E. (2017). The role of data privacy in marketing. *Journal of the academy of marketing science*, 45(2), 135–155. <https://doi.org/10.1007/s11747-016-0495-4>
- [30] Martin, K. D., Borah, A., & Palmatier, R. W. (2017). Data privacy: Effects on customer and firm performance. *Journal of marketing*, 81(1), 36–58. <https://doi.org/10.1509/jm.15.0497>
- [31] Jaakkola, E. (2020). Designing conceptual articles: Four approaches. *AMS review*, 10(1), 18–26. <https://doi.org/10.1007/s13162-020-00161-0>
- [32] MacInnis, D. J. (2011). A framework for conceptual contributions in marketing. *Journal of marketing*, 75(4), 136–154. <https://doi.org/10.1509/jmkg.75.4.136>
- [33] Snyder, H. (2019). Literature review as a research methodology: An overview and guidelines. *Journal of business research*, 104, 333–339. <https://doi.org/10.1016/j.jbusres.2019.07.039>
- [34] Bleier, A., & Eisenbeiss, M. (2015). The importance of trust for personalized online advertising. *Journal of retailing*, 91(3), 390–409. <https://doi.org/10.1016/j.jretai.2015.04.001>
- [35] Castelo, N., Bos, M. W., & Lehmann, D. R. (2019). Task-dependent algorithm aversion. *Journal of marketing research*, 56(5), 809–825. <https://doi.org/10.1177/0022243719851788>
- [36] Logg, J. M., Minson, J. A., & Moore, D. A. (2019). Algorithm appreciation: People prefer algorithmic to human judgment. *Organizational behavior and human decision processes*, 151, 90–103. <https://doi.org/10.1016/j.obhdp.2018.12.005>
- [37] Luo, X., Tong, S., Fang, Z., & Qu, Z. (2019). Frontiers: Machines vs. Humans: The impact of artificial intelligence chatbot disclosure on customer purchases. *Marketing science*, 38(6), 937–947. <https://doi.org/10.1287/mksc.2019.1192>
- [38] Lee, S., Moon, W. K., & Song, Y. G. (2025). Is it transparent or surveillant? the effects of personalized advertising and privacy policy on advertising effectiveness. *Journal of interactive advertising*, 25(1), 1–21. <https://doi.org/10.1080/15252019.2024.2417422>
- [39] Toti, J. F., & Steils, N. (2024). Showing or hiding? the impact of visibility of general conditions of use on retargeted ad intrusiveness and perceived ethicality of mobile apps. *Journal of interactive marketing*, 59(3), 251–272. <https://doi.org/10.1177/10949968231221472>
- [40] Dietvorst, B. J., Simmons, J. P., & Massey, C. (2016). Overcoming algorithm aversion: People will use imperfect algorithms if they can (even slightly) modify them. *Management science*, 64(389), 1155–1170. <https://repository.upenn.edu/handle/20.500.14332/39569>
- [41] Shaddy, F., Friedman, E. M. S., & Toubia, O. (2026). Fairness perceptions in demographic targeting. *Journal of consumer research*, 53(1), 22–47. <https://doi.org/10.1093/jcr/ucaf048>
- [42] Chaudhuri, A., & Holbrook, M. B. (2001). The chain of effects from brand trust and brand affect to brand performance: The role of brand loyalty. *Journal of marketing*, 65(2), 81–93. <https://doi.org/10.1509/jmkg.65.2.81.18255>
- [43] Chen, C. J., & Cheng, J. M. S. (2026). When help feels creepy: A MOA model of AI smart-retail adoption via perceived intrusiveness and the moderating role of trust in AI. *Journal of retailing and consumer services*, 90, 104699. <https://doi.org/10.1016/j.jretconser.2025.104699>
- [44] Tangari, A. H., Bui, M. (Myla), Krishen, A. S., & Kemp, E. (2026). AI-driven personalization in social media advertising: How tracking awareness and eWOM Homophily shape trust and consumer responses. *Psychology & marketing*, 1–18. <https://doi.org/10.1002/mar.70166>