

WHEN PRIVACY CONCERNS DON'T STOP DATA SHARING: EXPLAINING THE DATA PRIVACY PARADOX IN AI-POWERED PERSONALIZED MARKETING

Umer Asjid Hussain

Independent Researcher, Graduate of University of Central Punjab

Umerasjid05@gmail.com

ABSTRACT

The integration of artificial intelligence (AI) into digital marketing has increased a paradox of consumer behaviour that has been present for years: consumers often voice significant privacy concerns but are still willing to provide platforms, brands, recommendation systems, conversational agents and personalised advertising technologies with personal information. The data privacy paradox is especially complicated in personalized marketing with AI algorithms, as AI can turn your usual behaviors into intricate forecasts about consumer preferences, intentions, identities and future behavior. The focus of the present study is to investigate why privacy-related issues do not necessarily deter consumers from disclosing their personal data in AI-assisted marketing interactions. The study employs a conceptual research design that involves an integrative review of the foundational and contemporary research literature, synthesizing privacy calculus theory, contextual integrity, trust theory, technology acceptance approaches, uses and gratifications, personalization research, behavioral economics and the recent human–AI interaction literature. The analysis reveals five interrelated mechanisms that contribute to this paradox: perceived benefits of personalization, trust and perceived control, contextualization of data disclosure, asymmetry between immediate and future perceived threats, and increasing opacity and inferential power of the AI systems. The results show that the “privacy paradox” is not just irrational inconsistency in attitudes and behaviour. Instead, data sharing is a process of negotiation, with consumers balancing privacy concerns with the benefits of privacy, utility, utility efficiency, utility gratification, social involvement, and feelings of control. The study also suggests that AI has also shifted the focus from privacy paradox (disclosure of information) to algorithmic inference (inference of information). The integrative framework that is developed in the discussion links privacy concern indirectly to disclosure via perceived risk, personalization value, trust, control, contextual appropriateness and privacy-risk salience. The authors conclude that responsible AI-marketing needs to go beyond “notice-and-consent” to focus on personalisation that is transparent, contextual, proportionate, explainable and respects privacy.

Keywords: data privacy paradox, artificial intelligence, personalized marketing, privacy calculus, consumer behavior, data sharing, personalization, consumer trust, algorithmic inference, AI marketing

1. Introduction

Artificial intelligence (AI) has transformed personalized marketing by enabling organizations to collect, analyze, and utilize vast amounts of consumer data to deliver highly targeted products, services, and advertisements. Although consumers increasingly express concerns about privacy, surveillance, and the misuse of personal information, they continue to share data in exchange for convenience, personalization, and perceived benefits. This apparent contradiction, known as the ****data privacy paradox****, has become an important issue in contemporary digital marketing. The present study examines this paradox

in the context of AI-powered personalized marketing, focusing on why privacy concerns do not necessarily prevent consumers from disclosing their personal data.

Today's marketing landscape is a data-driven marketing world. Consumers are looking for products, visiting websites, engaging with ads, using mobile apps, conversing with conversational systems, shopping, reading about products, watching media and engaging in online communities. All the activities generate information which can be gathered, stored, aggregated, analyzed and turned into commercial intelligence. These digital footprints have become much more interesting in the era of artificial Intelligence (AI) due to the ability of AI to recognize patterns within large and diverse data sets and translate consumer behavior into forecasting. Personalized marketing in turn is shifting from traditional segmentation to increasingly dynamic systems, which can customize messages, recommendations, prices, interfaces, and promotional offers to the individual.

The business prospects of this development are quite good. Personalized marketing can help minimize information overload, improve product discovery, make ads more relevant, improve customer experiences, and improve engagement. Consumers might like to receive recommendations that align with their interests and not be the same kind of promotional messages. Recommendation systems based on artificial intelligence can help make search and evaluation easier, and conversational systems can offer increasingly personalised support. Recent studies also highlight the ability of AI-driven recommendations to help engage consumers and add value if they believe such suggestions are beneficial, pertinent, and suitable.

But the very same technologies which enable personalization also create significant privacy concerns. There is a risk of consumers being concerned with surveillance, profiling, unauthorized second use, manipulation, security, discrimination, and a lack of control over personal data. "The paradox is thus obvious: while consumers seek customised services, they may not necessarily desire to be closely monitored to get them. They can show a lot of anxiety regarding the gathering and utilization of their personal data, and still go to data-heavy websites and experience customized tips. The privacy paradox is the central phenomenon as it is a difference between attitudes and behaviors regarding privacy that exists (Kokolakis, 2017; Norberg et al., 2007).

The privacy paradox has been around for a while. Subsequent research, however, showed that people's expressed concerns for privacy do not always correspond to their actual disclosure behavior (Norberg et al., 2007), and that contextual and psychological influences such as perceived benefits, perceived risks, trust, social norms, information sensitivity, and disclosure context factors influence disclosure behavior (Acquisti et al. 2015, ; Barth & de Jong, 2017 ; Kokolakis, 2017). The continued relevance of the paradox highlights the fact that privacy concern is not the only factor in the consumer behavior.

But with AI comes a qualitatively different dimension. Personalization was usually based on more identifiable data like demographics, past purchases, browsing habits, or

verbal preference. Modern AI systems can extract more and more sophisticated insights from seemingly mundane data. They can detect patterns which consumers are not aware of. The very basic privacy question is thus shifting from what consumers reveal to what AI systems are able to infer from what consumers reveal.

Theoretically, this shift has significant consequences. A consumer can give up information regarding their browsing or purchasing habits without knowing that a system could combine some of those pieces of data to draw out his/her preferences, vulnerabilities, intentions, etc. Information is no longer sufficiently viewed as a controlled thing, that is, information that one voluntarily discloses. It must also be recognized as control over the inferences that the organizations and algorithms can make from that information.

The value of this development has been proven recently through research. In fact, the relevance, convenience, enjoyment, transparency and perceived control of the AI-driven recommendations are powerful enough to sustain consumers' engagement in such a system even when it feels intrusive, as demonstrated by Alharthi and Awad (2026). Their results are relevant for the paradox, as they point towards the fact that intrusiveness can negatively affect personalization benefits, but not necessarily compromise consumer involvement.

For this reason, the current article aims to shed light on the question that now circulates the center of the contemporary discussion on AI for marketing: why do users still share their data despite the concerns of privacy? The study does not assume consumers are irrational or careless, but rather considers data disclosure as a multi-dimensional decision wherein there is a complex interplay of benefits, risks, expectations, emotions, social norms, and technology.

The article uses an integrative conceptual approach. It combines fundamental privacy-paradox literature and recent studies on personalization with AI, consumer trust, algorithmic intrusiveness, disclosure, engagement, and human–AI interaction. The goal is not testing of a statistical model, but a coherent explanation in terms of existing theory.

Research Problem

The main research question is: why is there a persistent gap between consumers' expressed concerns about privacy, and their ongoing readiness to provide personal data to AI-driven marketing systems? Traditional notions about privacy often assume that if people know it is risky, they will take action to protect themselves. In reality, however, consumers' behavior often tells a different story. People might be aware of the dangers of surveillance and profiling, but they might still agree to personalized ads, recommendations, loyalty programs, cookies, tracking, and AI-driven services.

As AI technology becomes more sophisticated, the problem gets bigger. The consumer can be presented with very precise recommendations, even though they do not know what information the recommendation was built on. This opacity leaves consumers with a lag in their awareness of what technology can accomplish. Consumers may

consequently be worried about privacy but not fully aware of privacy implications of their actions.

The attached study (2026) is particularly pertinent to this context as it proposes that perceived intrusiveness and disclosure are two key contextual variables in AI recommendations, and that privacy discomfort alone does not necessarily end engagement when perceived relevance, convenience, enjoyment and control are still compelling. This article builds on the idea of AI-driven recommendations and co-creating consumer value by taking a look at the more general topic of AI-driven personalised marketing.

The study benefits from AI in several ways, including its ability to enhance personalization, predictive profiling, and automated analysis of consumer data, which boosts the relevance and effectiveness of targeted marketing. Meanwhile, the very same features also raise concerns about privacy, surveillance, and data ownership, adding to worries about why consumers are willing to give up their information despite the potential for privacy concerns.

Research gap

While previous research has extensively explored privacy concerns, trust, perceived risk, and consumers' willingness to share their personal information, few studies have sought to explain why consumers persist in sharing data even when they know of privacy risks in personalized marketing settings based on AI. Most of the current research either measures privacy attitudes or measures data sharing intentions but not thoroughly examines the paradox between consumers' reported concerns about privacy and their actual data-sharing behaviours. Furthermore, AI, predictive analytics and hyper-personalization are making their way into our world at an ever faster pace, leading to new data collection and profiling experiences that have not been investigated enough, especially in an emerging-market framework. This study thus fills this gap by investigating the determinants of the data privacy paradox in AI-enabled personalized marketing and the balancing act consumers perform when considering privacy concerns and the benefits of personalized services.

Objectives of the Study

The purpose of the study is to shed light on the reasons behind the data privacy paradox in the context of personalized marketing via AI. Specifically, it investigates the theoretical mechanisms underlying the persistent disclosure of data despite privacy concerns, explores how personalization benefits, trust, perceived control, contextual appropriateness, and AI-related intrusiveness contribute to these mechanisms, and proposes an integrative conceptual account of privacy concern's lack of effect on privacy-protective behavior.

Research Questions

The study addresses three research questions:

RQ1: Why do consumers continue to share personal data with AI-powered personalized marketing systems despite expressing concerns about privacy?

RQ2: How do personalization benefits, perceived risk, trust, perceived control, contextual appropriateness, and AI-related intrusiveness shape the relationship between privacy concerns and data-sharing behavior?

RQ3: How can an integrative conceptual framework explain the data privacy paradox and inform more responsible AI-powered personalized marketing?

Significance of the Study

This study is important because it tries to redefine the privacy paradox in the dynamically evolving realm of artificial intelligence. The bulk of the literature on disclosure of privacy has been written in a context prior to the proliferation of generative artificial intelligence, chatbots, algorithmic recommender systems, and almost entirely adaptive personalization systems. For contemporary AI, then, privacy research needs to rethink some of its fundamental assumptions.

The study also has marketing theory implications as personalization is traditionally considered to be a way of enhancing relevance and consumer response. However, AI proves to be personalization at the same time as a system of surveillance and inference. This technology can then enhance the perceived value and perceived vulnerability.

In the practitioner's world, it's crucial to grasp this dynamic, as no matter how much their personalization is aggressive, in the short term, it can result in engagement, but in the long term, it can foster distrust. Seemingly recent scholarship has shown that transparency in disclosure, contextual appropriateness of personalization and meaningful control reduces perceived intrusion and allows for engagement. The study therefore can play a role in developing an ethically sustainable marketing model with the help of AI.

Method

To investigate how consumers manage their conflicting attitudes and motivations regarding privacy and ongoing sharing of personal information in personalized marketing by AI, this study uses a qualitative research approach. Semi structured interviews will be conducted to gather data and analyzed it in thematic way, to find the common patterns, motivations and situational factors behind the data privacy paradox.

Research Design

This research used a qualitative research design with a conceptual approach using an integrative literature review. Primary survey or experimental data set were not created, since the purpose of the study is to explain a theoretical phenomenon. The methodological goal, however, is to uncover common conceptual connections between the existing body of privacy research and recent contributions to the field of artificial intelligence, personalization, consumer research, trust, disclosure, and human-AI interaction.

The integrative approach is suitable since the privacy paradox has emerged in various fields of academia. Knowledge of relevant scholarship will be found in information systems, marketing, psychology, communication, human-computer interaction, behavioural economics, consumer research and ethics. A narrow disciplinary review would then likely overlook the multidimensional nature of the paradox as indicated in the literature.

Analytical Approach

The literature was conceptually categorized based on the nature of the relationship between privacy concern and data sharing behavior. Research focused on privacy calculus, personalization, information disclosure, perceived intrusiveness, trust, control, contextual privacy, technology acceptance, consumer engagement, and AI-mediated interaction received special attention.

Analysis was conducted in a thematic synthesis approach instead of a statistical meta-analysis. Explanatory mechanisms were sorted by being identified as recurring and compared across theoretical traditions. The main themes that arose from the synthesis were perceived benefits, perceived risks, trust, control, contextual appropriateness, social normalization, technological opacity, emotional engagement, and algorithmic inference. This is aligned with recent studies that view AI recommendation systems as socio-technical contexts where technology usefulness, motivations, privacy judgments, disclosure and interaction are not independent.

Conceptual Scope

The analysis identifies three phenomena that are related but different: privacy concern, privacy-protective behavior, and data-sharing behavior. Privacy concern: When an individual is aware or apprehensive about the collection and use of information about him or her. These actions are called privacy protective behaviour, and include refusing permissions, deletion of accounts, changing privacy settings, using privacy-enhancing technologies, or avoiding certain platforms. Data-sharing behaviour is the actual sharing or creation of data in digital places. This separation is important because sometimes a consumer might hold a high privacy concern level and a high data-sharing level. The paradox comes about because these dimensions are not always going in the same direction.

RESULTS

The Privacy Paradox Is an Attitude–Behavior Gap, but Not Simply an Irrationality Problem

This is the first key finding of the conceptual synthesis: Privacy paradox cannot be fully explained as a contradiction between consumers' words and actions. There has been foundational research that showed that there is a gap between intentions to use privacy and actual disclosures (Norberg et al., 2007), and more recent reviews have shown that this phenomenon is common across digital contexts. But this isn't always a sign of irrationality among consumers (Barth & de Jong, 2017; Kokolakis, 2017).

This is only a part of the broader decision environment, however, that addresses privacy. Consumers can appreciate privacy as well as desirability for convenience, efficiency, personalization, entertainment, social engagement, and economic value. Thus, a consumer can justifiably value privacy and weigh the tradeoff of immediate benefit from disclosure as justifiable. This understanding moves from a privacy concern/disclosure paradigm to a privacy concern/multi-dimensional utility structure.

Personalization Functions as Both the Cause and the Reward of Disclosure

The second big discovery is that personalisation is a uniquely paradoxical thing. Personalization is based on data – and the best thing that consumers get in return for the data is personalization. This creates a vicious circle. Consumers share data, AI systems leverage data to better personalize recommendations, consumers feel the recommendations to be more personal, and the perceived benefits of personalisation motivate consumers to engage further. Research in the field of personalized advertising has consistently revealed that personalized advertising is more effective when consumers perceive the personalization as relevant and appropriate (Aguirre et al., 2015; Bleier & Eisenbeiss, 2015). All the while, over-collection or clandestine data collection can turn personalization into a means of mistrust. Similarly, the attached research from 2026 finds that AI recommendation systems can benefit consumer engagement on relevance, convenience and discovery basis but perceived intrusiveness can have a negative impact on the relationship.

Privacy Calculus Explains Why Concern Does Not Automatically Produce Refusal

The third finding is that the inclusion of privacy calculus in the understanding of data sharing behavior hasn't changed. Consumers consider the costs and benefits of disclosure (Culnan & Armstrong, 1999; Dinev & Hart, 2006). The perceived benefits can be more relevant recommendations, few search efforts, discounts, convenience, entertainment, better service, and personalized experiences. The costs experienced include surveillance, misuse, unauthorized use, identity theft, profiling, manipulation, and loss of autonomy. The privacy paradox occurs when the consumers recognize the cost but think the benefits are still worthwhile. Thus disclosure need not be negatively or linearly related to privacy concern. A consumer who is very concerned about privacy might share information even if he/she feels the level of privacy is very low.

Trust Converts Uncertainty into Manageable Risk

Another key explanatory mechanism that arose was trust. When consumers feel organizations can be trusted to handle the information properly, they are more likely to share it (McKnight et al., 2002; Martin, 2018). But, what trust does not remove is privacy concerns. Instead, it alters the perception of risk and harm to privacy. A consumer might feel comfortable with a trusted retailer using the information, while not trusted one might not feel comfortable with it being disclosed. However, consumers might perceive a difference between trust in a company and trust in the algorithmic system that a company employs, making this process more complicated with AI. A consumer might be more

trusting of a brand, but not know what the AI system does with the consumer's personal data. The latest studies on disclosure with AI indicate that disclosure to AI can have an impact on brand trust and exploitation.

Perceived Control Is More Important Than Privacy Concern Alone

The synthesis is also illustrative of the need for the concept of perceived control. Data collection may be acceptable if the consumer feels that they can affect what is gathering, how it is used, or if they can prevent personalization. The reason why there is no universal relationship between privacy concern and withdrawal is that it depends on the setting in which the decision is made. Resignation can be a result of concern if it is not managed. Consumers might be aware of the amount of information that can be collected from them, but decide to not collect it means they must give up the service. The study attached also highlights disclosure and control as ways that could help consumers feel more in control of personalization. Hence, it matters a great deal whether the concern and control efficacies are distinct. Concern is awareness of risk, perceived control is confidence of ability to manage risk.

Contextual Appropriateness Determines Whether Personalization Feels Helpful or Intrusive

Personalization is also contextually assessed, as indicated in the literature. The theory of contextual integrity proposed by Nissenbaum (2004, 2010) is relevant for many reasons, yet because privacy is not just about keeping others from knowing, but also about allowing the right people to know, this theory is particularly useful. Consumers can agree to give information in one situation then say no in another. For instance, a consumer might be willing to give the buyer information to facilitate the purchase, but be uneasy if that information is later used to make inferences about other aspects of the consumer's life. It's not always about collecting data, the problem is. It could be the surprising flow of information along contextual lines.

AI Opacity Creates an Inferential Privacy Problem

A particularly important result of the synthesis is that AI transforms privacy from an information-control problem into an inference-control problem. Consumers could be unaware what they're revealing and what may be learned by an AI system from those disclosures. The collection of browsing history, transaction data, location data, interaction data, and engagement data can yield insights in ways that are not explicitly communicated by the consumer. This leads to a lack of symmetry of knowledge. Organizations can possess much more knowledge about consumers than consumers can about the organizations' level of knowledge about consumers. Thus, AI makes the epistemic gap between consumer behavior and algorithmic consequence larger.

Immediate Benefits Outweigh Distant Risks

The synthesis also suggests that the temporal dimension of privacy decisions plays a part in the paradox. The benefits of personalization are often immediate and the harms of privacy are often unknown and postponed. Time can be saved today, with a personalized recommendation. Psychologically aggregated information is less immediate because of the potential for misuse several months on. The behavioural-economics literature indicates that people often undervalue future outcomes that are not guaranteed compared to more immediate ones (Kahneman, 2011; Thaler, 2015). This imbalance is partly responsible for the fact that consumers can be aware of privacy issues and still provide information.

Social Normalization Makes Disclosure Ordinary

Data sharing is also “normalized.” As consumers, they see and interact with cookies, loyalty programs, social media tracking, recommendations, personalized ads, and AI assistants without realizing they are all forms of artificial intelligence. As users of the digital world, consumers see and use cookies, loyalty programs, social media tracking, recommendations, personal ads and AI assistants without knowing what they are all. As disclosure becomes commonplace, the meaning of every disclosure gets lost. Consumers might not have to choose to give up privacy every single time they access a service, but rather disclosure can become part of the normal ways consumers access services. The privacy paradox is thus not just a personal one, but a social one.

Intrusiveness Has a Threshold Effect

The synthesis suggests that personalization does not lead to a linear relationship between personalization and consumer acceptance. Some degree of personalization can help create greater relevance, but too much can be uncomfortable. This is in particular supported by the attached study from 2026, which characterizes intrusiveness as a boundary condition that lowers, but does not invariably overthrow, the link between AI recommendations and the involvement of consumers. The issue then is not so much personalization versus no personalization, but it is about the right level of personalization versus over personalization.

DISCUSSION

Reinterpreting the Data Privacy Paradox

This study's main contribution is the re-interpretation of the privacy paradox as a multi-dimensional behavioural negotiation and not a simple paradox. When consumers share information, they aren't always erased of their privacy concerns. Rather, they could shift for a short period of time to prioritizing alternate goals. This difference is important because it muffles the underlying architecture of the digital world where their choices are made. When a consumer is not provided with a useful service unless they agree to data collection, the giving up of information may be a trade-off, either explicit or implicit. Privacy is still important, but it is in competition with other values. There is an important theoretical basis for this interpretation in the privacy calculus theory (Culnan & Armstrong, 1999; Dinev & Hart, 2006). However, the modern AI landscape demands a new perspective on the calculus

as customers are not only considering whether an organization will have information, but what it will have. They are seeing an increasing likelihood that this information will be turned into predictions and classifications using AI that will impact future decisions. The paradox has thus evolved to a higher level. It's not simply about the consumer asking, "Should I share this information?" The real question is, "What can this system learn from what I share?"

From the Privacy Paradox to the Personalization–Inference Paradox

This research suggests that the AI-powered marketing creates a new phase of the privacy paradox that could be called the personalization–inference paradox. Traditional privacy research has been dominated by the disclosure paradigm, which emphasizes disclosure whereby individuals disclose information even when they are concerned about it being collected and used. With AI, the added element of information can be obtained from the system that is not explicitly given by the consumers. A consumer can reveal a search query relating to travel. The AI system can make inferences based on patterns around that query, such as travel patterns, purchasing power, timing, family circumstances etc. These characteristics may not have been specifically communicated by the consumer. As a result, privacy is becoming a growing concern not only in terms of possession, but interpretation. This will have implications for privacy governance. Typically the consent methods used are based on types of data that are collected. But consumers have a legitimate question they can ask: What do you think about me based on these data? Transparency is thus not only about the data, but also the processes of inferences. Thus, an ethically responsible AI marketing system should be more transparent in its inferences processes than in the data collection processes themselves.

Why Privacy Concern Does Not Necessarily Become Privacy-Protective Behavior

The second big interpretative takeaway is around the difference between concern and action. Privacy concern is a psychological state while privacy-protective behavior is an external action. Motivation, knowledge, alternatives and perceived efficacy are needed for a transition between one and the other. A consumer might be very sensitive about privacy, but feel that it is impractical to refuse the collection of data. If everyone is tracking and personalising, then leaving one platform could just result in using another platform where the practices are all the same. This produces a condition of privacy resignation. Consumers aren't necessarily without interest; they might think that the resistance of one consumer has a negligible impact on the system. In the context of AI-driven ecosystems, surrendering to personalization is especially crucial as personalization can become integrated into the design of digital services. The consumer can enjoy personalisation either as a special service or as a standard condition of the service.

The Double-Edged Nature of Personalization

Personalization is the answer to and the cause of the privacy dilemma. As a consumer, personalization is a real blessing. It cuts down on the time you spend searching, weeds out

irrelevant information, aids in sifting through choice overload, and delivers services that seem customized to you. The significance of AI recommendation systems lies in their ability to sift through vast amounts of data and present it in a digestible format, which is why they are becoming more crucial every day. According to attached research, the personalized recommendation is an approach to make discovery and engagement easier. But, the effectiveness of personalization is dependent on information about the individual. The more the system is able to understand the customer the more personalized the experience will be. The more knowledgeable consumers are, the more likely they are to think that the system is intrusive. This forms an inverted curve. When personalization is at low levels, relevance might not bring any value. Personalization can create maximum usefulness at certain levels. Personalization can be uncomfortable if it becomes too personal. A maximal level of personalization is therefore not likely to be the best marketing approach. What is really meant is actually the right sort of personalization at the right time and right place.

“Helpful” and “Creepy” Are Two Interpretations of the Same Technology

One thing that needs to be taken into account is the emotional difference between helpfulness and creepiness. A similar algorithmic capability can have very different meanings for consumers. If an AI system suggests a product that matches an obvious recent search, it may be perceived as useful by the consumer. The same technological capability can be seen as surveillance when the recommendation seems to show information the consumer considers to be private or not related. This proves that the response of the consumer to the system is influenced by what the consumer thinks the system knows as well as by what the system actually does. The research attached further suggests that when recommendations have clear purposes and meaningful disclosure and control is present, consumers may continue to participate, but the perception of invasion of privacy can result in disengagement. So the question marketers shouldn't ask is whether personalization is effective at making things relevant. They should inquire about consumers' understanding of the personalization.

Transparency as a Psychological Mechanism

Transparency is almost viewed as a rule of regulation, but the literature suggests that it has psychological effects as well. Clear explanations can enhance perceived control, uncertainty reduction, and trust. The attached research for 2026 is particularly relevant because it considers disclosure as a process of transparency and suggests contextual explanations like “why this recommendation?” and explanations on the type of data used. This is better than the belief that “a long Privacy Policy will lead to informed consent. Consumers don't need to know all the details about data processing at every turn. When the data practice is relevant to their decision, they require meaningful information. As such, transparency should be context sensitive, succinct, relevant, and understandable.

Why Traditional Consent Is Not Enough

The privacy paradox also highlights weaknesses of traditional consent models. Purchasing “accept” doesn't guarantee consumers are aware of how AI systems use their information. The processing of AI can be technically challenging, probabilistic, adaptive, and dynamic. Even a big enough data set to infer a pattern is sometimes hard to predict for the experts. It is unreasonable to expect ordinary consumers to predict all the potential implications of disclosure. The aim of privacy governance should then move away from formal consent to building responsible data relationships. Organizations must take a larger share of responsibility for ensuring that there is no unnecessary data collection, and not simply the consumer.

Trust Does Not Mean Blind Acceptance

An interpretation of trust must also be carefully considered. Trust may lead to data sharing but responsible trust does not mean unquestioning acceptance. Consumer and marketer need to build a relationship that is healthy and beneficial, with a foundation of earned trust, accompanied by consistent behaviour, transparency, accountability, and empowering control. The benefits of personalization can be significant in the short term, but the consequences of using consumer data for uses that it was not originally intended for can have reputational consequences that can be more damaging than beneficial. AI research over the past few years reaffirms the need for trust when it comes to consumer acceptance of information exchange through AI. Thus, trust is both an enabler of personalisation and a resource that can easily be depleted due to inappropriate data practices.

Perceived Control as the Missing Link

Among the strongest theoretical implications of this study is the finding that perceived control could serve as a link between privacy concern and disclosure behavior. If the answer to the question “Do I perceive a risk?” is “yes,” privacy concerns are activated. Perceived control is the answer to the question “can I do something about that risk?” If the consumer feels that they can control it, then some level of personalization may be acceptable since they may be able to change their setting, limit their categories, revoke permission, or exit the system. The same level of privacy threat can lead to resignation and not protective behavior when there is no control. This insight points toward the need to invest in policies but also in interfaces that ensure control visibility and usability—especially for organisations looking for sustainable personalization.

The Role of Engagement

Adding to the privacy dilemma are the complexities of consumer engagement. More concretely, highly engaged consumers might be more flexible on what forms of personalization they allow with the platform due to the big amount of experiential value they get. The attached study proposes a conceptualization of engagement as a process that can enable personalisation to become consumer participation or value co-creation. This

understanding can also be applied to personalised marketing in general. If a user has a high need for a service, they may accept more processing of their data, whereas if they do not see the service as particularly valuable to them, they may not be willing to let their data be processed to that extent. But engagement should not be seen as a giveaway of consumer expectations of privacy. Even a highly engaged consumer can respond negatively to a violation of the privacy boundary.

The Importance of Contextual Integrity

In particular, contextual integrity offers a compelling account of the reasons some kinds of data sharing are allowed and others are not. Consumers have expectations that information will flow among the contexts (Nissenbaum 2004, 2010). The issue that occurs is when AI systems transfer information from one context to another without communicating the transfer of information. For instance, data supplied for the purpose of carrying out a deal might be employed for advertising or profiling or prediction later. The information is technically disclosed by the consumer but the use of the disclosed information may be a violation of the consumer's contextual expectations. However, the privacy paradox is not just about whether or not the consumer gave consent. It also needs to consider the appropriateness of the flow of information that followed it.

The Social Construction of Privacy

Privacy isn't just a personal choice. It is also social in nature. Consumers acquire their expectations of privacy through family, peers, institutions, media, and culture. Disclosure in digital environments is becoming more commonplace. The combination of this leaves a paradoxical social state in which individuals could be concerned about privacy in public yet be involved in high-visibility sharing cultures. This normalisation has gained pace with social media. Geography, photos, tastes and purchases, relationships, opinions are shared routinely. Thus, the privacy paradox may be a result of part of the changing definition of privacy. Consumers will no longer be guaranteed total informational privacy. Rather, they may be looking for context to the use of information.

AI and the Expansion of the Surveillance Logic

Traditional privacy principles overlook data minimisation – that organisations should only gather data for legitimate purposes. There is another principle of AI, namely that of minimizing inference. Technically, an organization can collect non-sensitive information and make inferences about it that are sensitive. Restricted raw data does not suffice to completely remove the privacy threat if algorithmic inference is not restricted as well.

From Data Minimization to Inference Minimization

The results indicate an integrative approach to the data privacy paradox in the context of personalised marketing with AI. First step in the process is privacy concern, creating an awareness of potential risk. There is a direct link between privacy concern and perceived privacy risk, and a direct link between perceived privacy risk and disclosure. But they

instead assess the value of personalization as it is perceived as relevant, convenient, efficient, entertaining, economical and socially rewarding. Meanwhile, they assess trust in the organization, control over information, contextual appropriateness, data sensitivity, and the intrusiveness of the personalization.

Responsible AI marketing should therefore ask not only:

What data are we collecting?

It should also ask:

What are we inferring?

How accurate are those inferences?

How sensitive are they?

Will consumers reasonably expect such inferences?

Could these inferences disadvantage consumers?

This represents an important conceptual development beyond conventional privacy calculus.

A Revised Conceptual Model

The results indicate the integrative approach for the data privacy paradox in AI-driven personalized marketing. Privacy concern is the first step in the process, creating an awareness of potential threats. Construction of the study: Concern for privacy might increase perceived privacy risk, which does not necessarily lead to less disclosure. Rather, consumers assess the perceived value of personalization, such as relevance, convenience, efficiency, entertainment, economic value, and social value. Meanwhile, they assess the trust in the organization, their perception of control over information, contextual appropriateness, data sensitivity and the intrusiveness of the personalization.

The resulting relationship can be represented conceptually as:

Privacy Concern → Perceived Privacy Risk → Privacy Calculus → Data-Sharing Decision

This core relationship is conditioned by:

Personalization Value + Convenience + Trust + Perceived Control + Contextual Appropriateness + Social Norms + Engagement

while being weakened by:

Intrusiveness + Unexpected Inference + Lack of Transparency + Perceived Exploitation + Loss of Autonomy.

This model explains why privacy concern may coexist with high data disclosure. The consumer does not necessarily disregard privacy; rather, privacy risk may be outweighed by other perceived benefits or constrained by contextual and structural conditions.

Resolving Rather Than Exploiting the Paradox

The privacy paradox has an important implication: Marketers should not assume that the more consumers give up, the more they can take. The consumer's consent to disclose is not an unlimited consent to all uses of the information. Instead, smart marketing should aim to minimize the circumstances creating the paradox. The transparency of data sharing can improve and become more mutually beneficial if consumers are aware of how their data is used, have meaningful control, and are getting genuine value from the personalization. The goal should be, then, value creation without informational domination.

Implications

This study advances privacy calculus theory by suggesting the need for an inferential and algorithmic dimension in the process of personalisation with AI. Classical privacy calculus is mainly concerned with benefits and risks perceived to arise from disclosures. The issue with AI is, what can be understood based on what is disclosed? Further, the study expands contextual integrity by showing that the disruption of information sharing can occur not only through the sharing of information between humans but also through automated aggregation and inference. Moreover, the study makes a contribution to the research of personalization by introducing the idea of a personalization as being nonlinear. The level of personalization that creates value is the point at which it becomes noticeable, enduring, or intrusive enough to create a sense of surveillance. Last but not least, the study identifies perceived control as an important intervening variable between privacy concern and disclosure. Efficacy will follow concern, while meaningful control will follow concern and lead to selective and informed disclosure, whereas resignation will follow concern without efficacy.

Theoretical Implications

The takeaway for marketers: More data does not ensure greater personalization. Data collection should be in proportion to the value it creates. These systems, based on Artificial Intelligence, should, therefore, offer consumers intelligible information about the reasons for the recommendations and about the general categories of information that make their way to them. This study, in particular, focuses on contextual disclosure and “why this recommendation?” mechanisms as design principles that are useful to incorporate. Marketers should also provide fine-grained controls. It is the duty of the consumer to limit the transmission of specific categories of data, be able to refuse certain types of personalisation or to refuse the service without being required to refuse the transmission of the data. Indicators of intrusion should be observed by organizations. If any of the following occurs, it could signify that personalization is going beyond consumers' expectations: complaints, negative feedback, sudden drop in engagement, “not relevant” responses, or

repeated rejections of the personalized recommendations. The attached study also suggests that when intrusion levels are exceeded, intensity of personalization be adjusted based on monitoring indicators of intrusiveness.

Managerial Implications

Ethical issues have nothing to do with the law! A legal but technically deficient practice could be considered an unfair or unreasonable practice. Autonomy, transparency, fairness, proportionality, accountability, and privacy by design are all important features of ethical AI marketing. Consumers shouldn't be left with the choice of agreeing to too much surveillance in order to receive basic digital services. Whether consumers see personalization as a service vs. a system performing on them will determine the long-term legitimacy of AI marketing.

Ethical Implications

The present research is not an empirical study to test the proposed relationships; rather, it is conceptual. The framework thus accounts for, but does not confirm, causal explanations. The model should be validated through future experimental studies that involve varying personalization intensity, data sensitivity, disclosure transparency, perceived control, and AI anthropomorphism. This could establish if privacy concern is differently affected with various personalization scenarios. Longitudinal studies would also be of value as it is possible that privacy attitudes could shift after positive or negative outcomes of data sharing. Personalization might start as a welcome gesture but can quickly turn into an issue of privacy after a consumer receives a "personalized" recommendation that is too intrusive or is exposed to a practice they didn't realize a company was using its data to make the recommendation. Cross-cultural research is of special significance as there are societal differences in privacy norms. Autonomy, trust, disclosure, and perso

Limitations and Future Research

It would also be beneficial for future studies to differentiate between traditional recommendation systems and generative AI systems. As interactions with conversational AI are social, anthropomorphic, and emotionally rich, disclosure might be prompted in a manner that is absent from conventional recommender systems. Lastly, study of inference privacy should be conducted. Whereas current metrics of privacy focus on whether consumers are concerned about data collection, the next generation of tools ought to also include whether consumers know about and agree to the inferences made by AI systems, based on their data.

Conclusion

Data Privacy Paradox is one of the most significant paradoxes in today's digital marketing landscape. As consumers are aware of the importance of their privacy, they are also engaging in digital systems that require a great deal of personal data. This paradox has not been resolved with the advent of artificial intelligence. Instead, it's created the paradox in

an enhanced manner. The main result of this conceptual study is that privacy concern does not necessarily lead to privacy-protective behaviors as there is a larger negotiation between perceived risks, benefits of personalization, convenience, trust, control, contextual appropriateness, engagement, social norms and technological conditions that shape consumer disclosure. Consumers can then share information with an authentic interest in privacy. They're not necessarily indifferent, nor irrational. But in many situations, consumers find personalization to be worth some amount of disclosing. Under other conditions, they might share information when they see little option available.

This is where AI can help, by increasing the gap between disclosure and inference. Consumers need to know not only what companies collect, but what inferences can be drawn from what is collected by AI systems. This should thus be understood as an inference-sensitive privacy calculus where consumers weigh the trade-offs between their enjoyed personalized experiences and the risks of algorithmic knowledge and intervention. AI-driven marketing personalization doesn't require the collection of as much consumer data as possible. Nor does it presume to have an unlimited right to infer from consumer consent. In the more sustainable model, personalization is transparent, proportionate, contextually appropriate and controllable.

In the end, it is not always about knowing the consumer and the most successful AI marketing systems. They might be ones that are just knowledgeable enough to be helpful, but not so knowledgeable that consumers feel like they are being “watched, manipulated, or exposed.” So the future of personalization hinges on the ability to overcome the privacy paradox from a consumer's inconsistency to a more responsible technological design.

REFERENCES

- 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>
- 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>
- Alharthi, B., & Awad, A. (2026). AI-powered recommendations and consumer value co-creation in a Personal Data Protection Law (PDPL)-regulated market: Engagement, disclosure, and intrusiveness. *SAGE Open*.
<https://doi.org/10.1177/21582440261463795>
- Awad, N. F., & Krishnan, M. S. (2006). The personalization privacy paradox: An empirical evaluation of information transparency and the willingness to be profiled

- online for personalization. *MIS Quarterly*, 30(1), 13–28.
<https://doi.org/10.2307/25148715>
- Barth, S., & de Jong, M. D. T. (2017). The privacy paradox—Investigating discrepancies between expressed privacy concerns and actual online behavior: A systematic literature review. *Telematics and Informatics*, 34(7), 1038–1058.
<https://doi.org/10.1016/j.tele.2017.04.013>
- Belanche, D., Casaló, L. V., Flavián, C., & Schepers, J. (2020). Service robot implementation: A theoretical framework and research agenda. *The Service Industries Journal*, 40(3–4), 203–225.
- Belanche, D., Belk, R. W., Casaló, L. V., & Flavián, C. (2024). The dark side of artificial intelligence in services. *The Service Industries Journal*, 44(3–4), 149–172.
<https://doi.org/10.1080/02642069.2024.2305451>
- 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>
- Brandimarte, L., Acquisti, A., & Loewenstein, G. (2013). Misplaced confidences: Privacy and the control paradox. *Social Psychological and Personality Science*, 4(3), 340–347. <https://doi.org/10.1177/1948550612455931>
- Culnan, M. J., & Armstrong, P. K. (1999). Information privacy concerns, procedural fairness, and impersonal trust: An empirical investigation. *Organization Science*, 10(1), 104–115. <https://doi.org/10.1287/orsc.10.1.104>
- 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>
- Davis, F. D. (1989). Perceived usefulness, perceived ease of use, and user acceptance of information technology. *MIS Quarterly*, 13(3), 319–340.
<https://doi.org/10.2307/249008>
- 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(5), 1431–1451. <https://doi.org/10.1007/s11747-023-00986-8>
- Kahneman, D. (2011). *Thinking, fast and slow*. Farrar, Straus and Giroux.
- Kokolakis, S. (2017). Privacy attitudes and privacy behaviour: A review of current research on the privacy paradox phenomenon. *Computers & Security*, 64, 122–134.
<https://doi.org/10.1016/j.cose.2015.07.002>

- Komiak, S. Y. X., & Benbasat, I. (2006). The effects of personalization and familiarity on trust and adoption of recommendation agents. *MIS Quarterly*, 30(4), 941–960. <https://doi.org/10.2307/25148760>
- Lambrecht, A., & Tucker, C. (2013). When does retargeting work? Information specificity in online advertising. *Journal of Marketing Research*, 50(5), 561–576. <https://doi.org/10.1509/jmr.11.0503>
- Lee, C. T., Pan, L. Y., & Hsieh, S. H. (2022). AI chatbots as brand promoters: A SEM–ANN approach. *Internet Research*, 32(4), 1329–1356. <https://doi.org/10.1108/INTR-01-2021-0030>
- Lefkeli, D., Karataş, M., & Gürhan-Canlı, Z. (2024). Sharing information with AI (versus a human) impairs brand trust: The role of audience size inferences and sense of exploitation. *International Journal of Research in Marketing*, 41(1), 138–155. <https://doi.org/10.1016/j.ijresmar.2023.08.011>
- Lopes, J. M., Gomes, S., Nogueira, E., & Trancoso, T. (2025). AI's invisible touch: How effortless browsing shapes perception, experience, and engagement in online retail. *Cogent Business & Management*, 12(1), Article 2440628. <https://doi.org/10.1080/23311975.2024.2440628>
- Martin, K. D. (2018). The penalty for using big data: Consumer privacy concerns and corporate data practices. *Journal of Business Ethics*, 152, 309–328. <https://doi.org/10.1007/s10551-016-3165-0>
- 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.15.0495>
- McKnight, D. H., Choudhury, V., & Kacmar, C. (2002). Developing and validating trust measures for e-commerce: An integrative typology. *Information Systems Research*, 13(3), 334–359. <https://doi.org/10.1287/isre.13.3.334.81>
- Meng, X. (2025). “Talk to me, I’m secure”: Investigating information disclosure to AI chatbots in the context of privacy calculus. *Online Information Review*, 49(5), 933–954. <https://doi.org/10.1108/OIR-06-2024-0375>
- Norberg, P. A., Horne, D. R., & Horne, D. A. (2007). The privacy paradox: Personal information disclosure intentions versus behaviors. *Journal of Consumer Affairs*, 41(1), 100–126. <https://doi.org/10.1111/j.1745-6606.2006.00070.x>
- Nissenbaum, H. (2004). Privacy as contextual integrity. *Washington Law Review*, 79(1), 119–158.
- Nissenbaum, H. (2010). *Privacy in context: Technology, policy, and the integrity of social life*. Stanford University Press.

- Pavlou, P. A. (2011). State of the information privacy literature: Where are we now and where should we go? *MIS Quarterly*, 35(4), 977–988.
<https://doi.org/10.2307/41409969>
- Strycharz, J., & Segijn, C. M. (2024). Ethical side-effect of dataveillance in advertising: Impact of data collection, trust, privacy concerns and regulatory differences on chilling effects. *Journal of Business Research*, 173, 114490.
<https://doi.org/10.1016/j.jbusres.2023.114490>
- Taddicken, M. (2014). The “privacy paradox” in the social web: The impact of privacy concerns, individual characteristics, and the perceived social relevance on different forms of self-disclosure. *Journal of Computer-Mediated Communication*, 19(2), 248–273. <https://doi.org/10.1111/jcc4.12052>
- Thaler, R. H. (2015). *Misbehaving: The making of behavioral economics*. W. W. Norton.
- Wang, Y., Zhu, J., Liu, R., & Jiang, Y. (2024). Enhancing recommendation acceptance: Resolving the personalization–privacy paradox in recommender systems: A privacy calculus perspective. *International Journal of Information Management*, 76, 102755. <https://doi.org/10.1016/j.ijinfomgt.2024.102755>
- Xu, H., Dinev, T., Smith, H. J., & Hart, P. (2011). Information privacy concerns: Linking individual perceptions with institutional privacy assurances. *Journal of the Association for Information Systems*, 12(12), 798–824.
<https://doi.org/10.17705/1jais.00281>
- Zarouali, B., Poels, K., Poncin, I., & Walrave, M. (2018). The impact of regulatory focus on the effects of ad personalization. *International Journal of Advertising*, 37(4), 613–633. <https://doi.org/10.1080/02650487.2018.1480167>
- Zhu, Y.-Q., & Kanjanamekanant, K. (2021). No trespassing: Exploring privacy concerns in personalized advertising. *Internet Research*, 31(2), 457–478.
- Zhu, Y., & Chang, J. (2016). The key role of relevance in personalized advertising: Examining its impact on consumer attitudes and behavioral intentions. *Computers in Human Behavior*, 61, 155–164.