

From Recommendation to Purchase: AI Personalization, Trust, and Privacy in E-Commerce

Tvisha Ahuja

Abstract

AI-driven recommendation systems increasingly shape online purchasing decisions, yet the extent to which trust, privacy concern, and age influence consumers' willingness to act on these recommendations remains unclear. Using survey data from 151 online shoppers- a sample skewing older, with 74.8% aged 35 and above - this study tested three hypotheses: that trust predicts purchase intention (H1), that privacy concern predicts purchase intention (H2), and that attitude toward personalization differs by age group (H3). H1 was supported: trust was significantly and positively associated with purchase intention ($r = .332$, $p < .001$) and remained a significant independent predictor when entered alongside privacy concern, $F(2, 148) = 9.29$, H2 was not supported; privacy concern showed only a weak negative association that did not reach significance ($r = -.156$, $p = .056$). H3 was not supported either, with no significant difference in attitude toward personalization between younger and older respondents (Welch's $t(64.19) = 1.43$, $p = .157$). Despite being the highest-rated construct overall, privacy concern did not exhibit a significant direct relationship with purchase intention. These findings suggest that the personalization–privacy trade-off may be better understood through the lens of consumer trust, highlighting trust as a potentially more important factor in translating AI-driven personalization into purchasing intentions.

Keywords: Social and Behavioral Science, Consumer Behavior, E-Commerce, AI-Driven Personalization, Consumer Trust

1.1 Background

Over the past two decades, e-commerce has transformed the way consumers discover and purchase products, it has grown from a niche convenience into a dominant mode of retail worldwide. What began as a small number of online storefronts has expanded into a global marketplace generating trillions of dollars in annual sales. The global e-commerce market size was valued at USD 33.9 trillion in 2025 and is projected to grow from USD 39.7 trillion in 2026 to USD 155.9 trillion by 2033, at a CAGR of 21.6% from 2026 to 2033, (Grand View Research, 2024). These numbers have been accelerated further by improvements in internet access, mobile technology, and digital payment systems, as well as by shifts in consumer behaviour during and after the COVID-19 pandemic. In rapidly growing markets such as India, the expansion has been especially pronounced, with rising smartphone usage and affordable data bringing millions of first-time shoppers online. India's e-commerce market was valued at around USD 137 billion in 2025 and is projected to reach USD 363 billion by 2030 (Mordor Intelligence, 2025). India has also surpassed the United States to become the world's second-largest e-retail market, with roughly 280–300 million online shoppers in 2025 (IBEF, 2025).

As online catalogues continuously expanded to include millions of products, platforms faced a new challenge: helping each shopper find what they want within an overwhelming amount of choice. This problem, often described as "choice overload," can leave consumers unable to decide and ultimately less likely to purchase [In a landmark 2000 study, Iyengar and Lepper found that shoppers presented with 24 jams were far less likely to buy than those shown only 6 -the classic demonstration of "choice overload.", Iyengar, S. S., & Lepper, M. R. (2000)]

The solution has been the rise of AI-driven recommendation engines. These systems analyse a user's browsing history, past purchases, searches, and behaviour to suggest products tailored to that individual. Early recommendation systems relied on relatively simple techniques such as collaborative filtering, which suggests items based on the preferences of similar users, and content-based filtering, which matches products to a user's past choices

More recently, advances in machine learning and deep learning have allowed these systems to process vast quantities of behavioural data and generate highly individualised suggestions in real time. Major platforms such as Amazon, Flipkart, and Alibaba Group have now built sophisticated recommendation systems that generate the familiar "recommended for you" and "customers also bought" suggestions that shape much of online retail. The commercial impact of these systems is considerable: industry analyses suggest that a substantial share of revenue on leading platforms is driven by recommended products [About 35% of Amazon's revenue is generated by its recommendation engine. Furthermore, 71% of shoppers expect personalized experiences and 76% feel frustrated when they don't get them, McKinsey & Company]

The effectiveness of personalization, however, depends on a continuous and expanding flow of personal data. To generate relevant recommendations, platforms must collect detailed information about what users view, click, search, purchase, and even how long they linger on particular items. This data-intensive model is precisely what makes personalization powerful, but it is also the source of growing unease. The same mechanisms that allow a platform to be helpfully attentive can equally be experienced as intrusive surveillance, and the boundary between the two is often unclear to the consumer.

However, the enthusiasm surrounding AI-enabled personalisation exists in tension with escalating consumer apprehensions regarding data privacy and security. Contemporary research indicates that only forty-seven per cent of individuals globally express confidence in AI companies' capacity to safeguard personal information, with this figure continuing to decline amid high-profile data breaches and revelations concerning algorithmic surveillance practices. In developed markets, including the United States, approximately seventy per cent of consumers familiar with AI technologies report minimal or absent trust in organisational handling of AI collected data, underscoring a significant credibility deficit that potentially constrains commercial adoption and consumer engagement. Closely related to privacy is the issue of algorithm transparency. Most consumers have little understanding of how recommendations are generated, why particular products appear, or what data underlies each suggestion. In fact, one survey found that 65% of consumers do not realise that the product recommendations they receive from online retailers are powered by AI (DataStax, 2023). This opacity, sometimes described as the "black box" nature of AI systems, can make personalization feel unpredictable or even manipulative, particularly when recommendations appear unusually well-informed about a user's private interests. Notably, when AI is used to make decisions about them, 83% of Americans say they would want to know exactly what information the system used (Consumer Reports, 2024)- a desire for transparency that the black-box nature of these systems rarely satisfies.

These competing forces give rise to a well-documented contradiction known as the privacy paradox: although consumers frequently express strong concerns about privacy, their actual behaviour often does little to protect their personal data (Barth & de Jong, 2017). People may report discomfort with being tracked, yet still click on recommendations, accept default settings, and complete purchases prompted by algorithmic suggestions, and understanding how this tension between stated concern and observed behaviour resolves itself in real purchasing decisions is central to evaluating the true impact of AI-driven personalization. Importantly, these responses are unlikely to be uniform across the population, as factors such as age, digital literacy, and shopping experience may shape how individuals perceive and react to recommendations; younger consumers, who have largely grown up with digital platforms, tend to be more comfortable with data-driven personalization and more inclined to act on recommendations, whereas older consumers often approach the same systems with greater caution. No doubt that 41% of Gen Z consumers say online convenience is worth the privacy trade-off, compared with 29% of Boomers, while 51% of Gen Z feel they have "become the product" in the digital economy, against 74% of Boomers (Usercentrics, 2025).

Such demographic variation suggests that the relationship between personalization, trust, privacy, and purchasing behaviour is more nuanced than a single, universal pattern. Taken together, these developments establish e-commerce personalization as a powerful but contested phenomenon: it offers clear benefits in convenience, relevance, and commercial value, while simultaneously raising concerns about privacy, transparency, and trust that may differ across consumer groups. It is within this context that the present study examines how trust and privacy concerns jointly influence consumers' purchase intentions in response to AI-driven recommendations, and whether these effects vary by age.

1.2 Problem statement

Although personalized recommendations are widely used to improve the relevance and convenience of online shopping, it remains unclear whether these benefits come at an acceptable cost to the consumer. The same systems that make shopping easier depend on the continuous collection, storage, and analysis of personal data, and in doing so they raise unresolved concerns about how that data is used, whether recommendation algorithms are transparent, and whether suggestions are designed to serve the consumer or to maximise spending. Each of these concerns bears directly on consumer trust - yet it is not established how strongly they actually shape consumer attitudes and behaviour, nor whether they affect all consumers to the same degree.

The specific problem this study addresses is therefore twofold. First, it is unclear whether privacy, transparency, and trust concerns are significant enough to offset the convenience and relevance gains that personalization offers - that is, whether personalization ultimately strengthens or weakens consumers' willingness to accept and act on recommendations. Second, it is unclear whether this relationship is uniform across consumers or is moderated by demographic factors such as age, with different age groups potentially weighing convenience against privacy in different ways. This study aims to resolve these questions by examining how consumers perceive and respond to personalized recommendations, and how demographic factors, age in particular, influence that response. Establishing this is necessary because, without it, platforms and policymakers cannot tell whether current personalization practices build or erode the very trust on which continued engagement depends.

1.3 Research gap

Two distinct streams of research bear on this study, and each has advanced considerably, but along separate tracks.

The first stream has established the commercial effectiveness of personalization. Work in this area consistently finds that recommendation accuracy and relevance raise engagement, satisfaction, and sales. Amazon's recommendation engine is estimated to drive roughly 35% of its sales (McKinsey, as cited in industry reporting, 2023), 88% of consumers report being more likely to buy when engagement is personalized in real time, and 71% will abandon a purchase that does not feel relevant (Twilio, 2025). What this stream demonstrates is that personalization works as a revenue lever. What it largely brackets out, however, is cost: it measures the upside of tailoring without weighing it against the privacy and trust concerns that the same data collection generates, and it tends to treat the consumer as a single, undifferentiated responder.

The second stream has established the privacy and trust dimensions of personalization. Scholars have documented the "personalization–privacy paradox," in which consumers simultaneously desire tailored experiences and guard their personal data (Awad & Krishnan, 2006; Barth & de Jong, 2017); shown how privacy concerns and privacy policies shape consumer trust (Wu et al., 2012); and mapped the defensive behaviours consumers adopt when those concerns are heightened (Bandara et al., 2021; Smith et al., 2011). This work confirms that privacy concerns and trust are real, persistent forces in how consumers respond to data-driven services. What it largely stops short of, however, is the purchase decision itself: it explains attitudes toward data-sharing and toward platforms in general, but rarely connects those attitudes to whether a consumer actually acts on a specific personalized recommendation - the outcome that matters commercially.

The gap lies precisely at the intersection these two streams leave open. The effectiveness literature measures purchase behaviour but omits privacy and trust; the privacy literature measures trust and concern but omits purchase behaviour - so the two central forces are rarely modelled together, as simultaneous and possibly competing influences on the same outcome. The specific tension this creates - a consumer who holds genuine privacy concerns yet continues to trust and buy from recommendation systems, therefore remains under-examined at this combined level. A second, compounding gap is demographic: both streams tend to treat consumers as homogeneous, leaving it unclear whether age moderates the trade-off, even though younger and older consumers plausibly weigh convenience against privacy very differently. This study addresses both gaps by modelling trust and privacy concern as joint predictors of purchase intention in response to personalized recommendations, and by testing whether that relationship differs between younger and older consumers. In doing so it integrates two literatures that have developed in parallel, offering a unified view of the personalization trade-off that single-factor studies - measuring effectiveness or privacy, but not their interaction, cannot provide.

1.4 Research Objectives

RO1

To examine the relationship between trust in recommendation systems and purchase intention.

RO2

To analyze the association of privacy concerns on purchase intention.

RO3

To investigate whether age influences attitudes toward personalized recommendations.

RO4

To understand consumer perceptions regarding adaptive product recommendations over time.

1.5 Research Questions

1. Does trust in recommendation systems increase purchase intention?

A significant amount of evidence shows trust as the central mechanism converting recommendations into purchases. In a large social-shopping study using 1,219 survey responses, trust in product recommendations was found to raise consumers' intention to purchase from the website (Hsiao et al., 2010). A meta-analysis pooling effect sizes from nineteen studies reached the same conclusion, confirming that trust positively influences consumers' purchase intentions, with trust in the seller exerting the strongest effect (Wang et al., 2022). More recent work focused specifically on personalization reports a very strong positive association between consumer trust and purchase intention, reinforcing that the more consumers believe a recommendation is competent and acting in their interest, the more willing they are to act on it. This body of work directly motivates Hypothesis H1

2. Do privacy concerns reduce purchase intention?

A lot of literature establishes privacy concern as a suppressing force for purchase intention. An early study of 148 online customers found that privacy concerns negatively affected both enjoyment and purchase intention in personalized online shopping (Pappas, Giannakos, & Chrissikopoulos, 2012). A study of Indian consumers similarly found that customers' intention to use online personalization was negatively associated with their degree of privacy concern (Priyanka & Siraj, 2021). Specifically, the study reported a significant negative correlation between intention to use online personalization and privacy concern ($r = -0.609$, $p < 0.001$). Evidence from social commerce also supports this relationship: Selem et al. (2023), using a sample of 490 social-media users in Pakistan, found that privacy concerns significantly affected purchase intention within a social-commerce context. Together, these findings provide empirical support for the proposed negative relationship between privacy concerns and consumers' purchase-related intentions and motivate Hypothesis H2.

3. Do younger consumers respond more positively to personalized recommendations than older consumers?

Evidence on age is somewhat consistent in direction, but it involves several factors. Industry surveys show younger cohorts are markedly more receptive: 47% of Gen Z and 40% of Millennials say personalization is important to them (Security Magazine, 2025), while only about 7% of older consumers say personalized recommendations would encourage them to shop more (Lightspeed, 2024). Older consumers also report greater privacy resistance - roughly one in ten baby boomers avoid online shopping specifically out of concern that brands will share their personal data, whereas over half of Millennials and Gen Z are not opposed to such data sharing (Euromonitor, 2023). However, an experimental study of 239 shoppers complicates the picture, finding that older adults are not always less receptive to novel recommendations, and under some conditions show heightened receptiveness (Zhao & Fu, 2024) - a caution worth citing, as it shows the age effect is real but not absolute. This evidence motivates Hypothesis H3.

4. How do consumers perceive evolving recommendation systems?

As recommendation systems grow more AI-driven and sophisticated, consumer perceptions reveal quite mixed feelings - and, importantly, those perceptions appear to feed directly into how far consumers are willing to trust the system. A qualitative study using 20 in-depth interviews surfaced a clear privacy paradox: most participants valued the convenience of AI recommendations, yet a majority reported "surveillance anxiety" when suggestions felt too personal, alongside a strong demand for transparency and control over their data profiles. Experimental work shows this tension operating at the level of awareness itself: greater algorithm awareness increases perceived usefulness but simultaneously intensifies skepticism, underscoring the need for transparent, user-controllable systems. Synthesising both forces, a

2026 study concludes that trust and perceived relevance act as the key enablers of AI-driven personalization, while privacy concern is the single strongest inhibiting factor

Read together, these findings point beyond description to a specific relationship: the way consumers perceive adaptive recommendations - how relevant, transparent, and well-tuned the suggestions feel - is itself a determinant of whether they come to trust the system. Consumers who experience recommendations as high in quality extend greater trust, while those who find them intrusive or opaque withhold it. This is precisely the relationship RO4 isolates - a pattern this study explores conceptually but does not test as a formal hypothesis, since the survey instrument was built around H1-H3.

1.6 Hypotheses

H1

Trust in recommendation systems positively influences purchase intention.

H2

Privacy concerns negatively influence purchase intention.

H3

Younger consumers demonstrate more positive attitudes toward personalized recommendations than older consumers.

1.7 Conceptual Framework

The conceptual framework for this study draws on the technology-acceptance and privacy-calculus traditions, which propose that a consumer's response to a data-driven service is shaped by weighing its perceived benefits against its perceived risks (Awad & Krishnan, 2006). Applied to personalized recommendations, the framework identifies three independent variables - trust in recommendation systems, privacy concerns, and age - and examines how they relate to the dependent variable, purchase intention, along with consumers' broader attitude toward personalized recommendations.

The model proposes three relationships. First, **trust in recommendation systems** is expected to have a positive influence on purchase intention: the more a consumer believes recommendations are designed to genuinely help them, the more willing they are to act on them (H1). Second, **privacy concerns** are expected to have a negative influence on purchase intention: the more a consumer worries about how their data is collected and used, the less willing they are to purchase recommended products (H2). Third, **age** is expected to moderate attitudes toward personalization, with younger consumers holding more positive attitudes than older consumers (H3). Together, these variables capture the central tension of the study - the pull of trust and convenience against the push of privacy concern - and allow it to be examined within a single integrated model rather than in isolation.

2. Literature Review - ROLs

2.1 Evolution of Personalized Recommendation Systems

Recommendation systems originated as rule-based and statistical models designed to assist users in decision-making under information overload. Early systems relied heavily on explicit user ratings and simple similarity measures. As data availability increased, these systems evolved toward learning-based approaches capable of handling larger datasets and more complex preference patterns. The integration of AI marked a significant shift, allowing recommendation systems to move beyond static assumptions toward adaptive and predictive personalization. ([Lawal G. Anand](#), 2023)

The idea of filtering information for individuals came even before the concept of modern e-commerce. The earliest systems introduced collaborative filtering (CF), in which recommendations are derived from

the preferences of like-minded users; the Tapestry system (Goldberg et al., 1992) and the GroupLens architecture for filtering Usenet news (Resnick et al., 1994) are widely regarded as the foundational works, with Resnick and Varian (1997) formalising the term "recommender systems." Early CF was user-based, but scalability problems led to item-based collaborative filtering, which computes similarities between items rather than users (Sarwar et al., 2001). A further leap came with latent-factor models: matrix-factorization techniques, popularised during the Netflix Prize competition (Bennett & Lanning, 2007), decomposed the user-item matrix to predict preferences at scale (Koren et al., 2009).

Running parallel to CF is content-based filtering, which recommends items resembling those a user has previously liked, based on item attributes rather than peer behaviour (Pazzani, 1999; Lops et al., 2011). Each approach has well-documented weaknesses-CF struggles with the "cold-start" problem for new users and items, while content-based methods tend toward over-specialisation- which motivated hybrid recommendation models that combine the two to offset their respective limitations (Burke, 2002, 2007). Most recently, the field has shifted toward AI and machine-learning methods: early neural approaches such as restricted Boltzmann machines for CF (Salakhutdinov et al., 2007) gave way to a broad family of deep-learning recommenders that learn complex, non-linear preference patterns from large behavioural datasets (Zhang et al., 2019). This trajectory- from peer-based heuristics to opaque, self-learning AI-frames why questions of trust, transparency, and privacy (Sections 3.3–3.4) have become increasingly pressing.

2.2 Personalization in E-Commerce

Building on these technical foundations, e-commerce has increasingly adopted AI-powered personalization, which can reshape how businesses engage with consumers by tailoring product recommendations, content, and user experiences to individual preferences (Vidhya et al., 2023). AI-powered personalization refers to the use of advanced algorithms and machine-learning techniques to tailor content, product recommendations, and user experiences to individual preferences (Haleem et al., 2022).

The key principles underlying AI-powered personalization involve the analysis of vast datasets -including user behaviour, preferences, and historical interactions - to generate insights that enable platforms to predict and deliver highly relevant content. The goal is to create a customised and engaging experience for each user, fostering a sense of personal connection with the e-commerce platform. These principles also encompass continuous learning and adaptation: as users interact with the platform, the AI algorithms gather data, refine their understanding of individual preferences, and dynamically adjust recommendations (Venkatachalam & Ray, 2022).

This capability has clear commercial value. By improving the relevance of what shoppers see, personalization reduces the "choice overload" that can deter purchasing (Iyengar & Lepper, 2000), and relevance has shifted from a bonus to a baseline expectation -McKinsey reports that 71% of consumers now expect personalized interactions and 76% are frustrated when they do not receive them (McKinsey, 2021). This expectation translates into measurable engagement and loyalty, with personalized experiences linked to higher purchase and repeat-purchase intention.

However, the relationship between personalization and purchasing is not purely positive. Research shows that relevance and usefulness raise purchase intention only up to a point -over-personalization can itself create discomfort and lower acceptance -and the broader privacy-personalization paradox means that concerns around privacy, data security, and consumer skepticism persist, potentially hindering trust in AI

applications (Wang et al., 2023). It is this tension between the benefits of personalization and the costs to privacy and trust that the following sections examine in detail.

2.3 Trust in Recommendation Systems

Across most current research, definitions of trust fall into various categories, and a solid definition for it, in many cases can be quite elusive. Marsh's work in [[Formalising Trust as a Computational Concept](#)] goes some way towards formalising trust in a computational sense, taking into account both its social and technological aspects. This study is concerned primarily with how trust functions in recommendation systems -both context-specific trust (in a particular recommendation) and system trust (in the platform). Three themes recur. First, trust reduces uncertainty: because online purchasing involves risk and information asymmetry, trust lowers perceived risk and improves attitudes toward online shopping (Handoyo, 2024), with a recommendation's perceived ability, benevolence, and integrity acting as the key drivers of that trust (Hsiao et al., 2010). Second, trust increases reliance: users' willingness to engage with recommender systems is influenced by their trust in the system, particularly perceptions of its competence and the quality of its recommendations (Saxborn, Pan, & Said, 2024). Third, trust enhances purchase decisions: a meta-analysis of nineteen studies found trust positively influences purchase intention, strongest for trust in the seller (Wang et al., 2022). Together, this positions trust as the pivotal variable converting a recommendation into a purchase - motivating Hypothesis H1.

2.4 Privacy Concerns and Consumer Behavior

A recurring observation in the privacy literature is that people will trade personal information for surprisingly small rewards. Carrascal et al. (2013), for instance, estimated that internet users valued their browsing history at roughly €7 - about the price of a fast-food meal. Yet surveys of internet users' attitudes simultaneously show that they are highly concerned about their privacy and about how their personal information is collected and used (TRUSTe, 2014; Pew Research Center, 2014). This contrast between stated privacy attitudes and actual behaviour has been termed the "privacy paradox" (Brown, 2001; Norberg et al., 2007) -or, more precisely, the "information privacy paradox." The e-commerce sector is very largely impacted by this paradox: While people say they care about privacy but often don't act to protect it. This encourages companies to collect more personal data and makes it harder to push for stricter privacy regulations. (Kokolakis, 2017).

Importantly, these concerns do not arise in a vacuum; they are shaped by how responsibly firms and regulators handle personal data. As Bandara, Fernando, and Akter (2021) argue, a lack of corporate privacy responsibility and regulatory protection can deprive consumers of privacy empowerment and damage consumer trust, leading to major privacy concerns and subsequent defensive responses. This directly connects privacy to trust - the two opposing forces at the centre of this study - and establishes privacy concern as the principal inhibiting factor in the proposed model, motivating Hypothesis H2.

2.5 Purchase Intention in Online Shopping

Purchase intention refers to the consumer's willingness or likelihood of buying a particular product or service. Rooted in models such as the Theory of Reasoned Action and the Theory of Planned Behaviour, it represents the final stage of the consumer decision-making process, and several studies point to different factors that influence it.

Various well known researchers have claimed that trust is the major catalyst impacting the consumer's willingness to purchase. As per the reasoned action theory, online shopping action can be stated as an intentional action which is strongly affected by consumer trust (Jeong et al., 2006). Curras-Perez et al.

(2017) argued that in the conditions of high perceived risk, there is a greater impact of trust on the purchase intention.

Perceived usefulness is a second key determinant. Many studies indicate that other factors like perceived usefulness also has a positive impact on technology acceptance, highlighting its key role in the consumer's final choice.(Rahmayanti, 2021). Perceived usefulness is defined as the extent to which customers believe that using a technology will provide significant value (Olaleye et al., 2018). In technology-adoption research, perceived usefulness is associated with users' attitudes and intentions to adopt mobile applications, with evidence suggesting that usefulness can contribute to positive attitudes toward the technology (Gupta & Arora, 2017).

Personalization quality can improve the customer experience and strengthen satisfaction and purchase-related outcomes when recommendations and interactions are appropriately tailored to individual preferences (Vass, 2023). Finally, privacy concerns act in the opposite direction, suppressing intention when consumers feel their data is being over-collected or misused. The present study draws these competing determinants - trust, personalization quality, perceived usefulness and privacy concerns -into a single model of purchase intention.

2.6 Demographic Differences in Personalization Acceptance

Consumer responses to personalization are not uniform, and generational differences are among the most consistently reported. Younger cohorts are markedly more receptive: industry research finds that 47% of Gen Z and 40% of Millennials consider personalization important (Security Magazine, 2025), whereas only about 7% of older consumers say personalized recommendations would encourage them to shop more (Lightspeed, 2024). Older consumers also show greater privacy resistance - roughly one in ten baby boomers avoid online shopping out of concern that brands will share their data, while over half of Millennials and Gen Z are not opposed to such sharing (Euromonitor, 2023).

This pattern is often framed through the lens of digital natives versus digital immigrants: consumers who grew up immersed in digital platforms tend to accept data-driven personalization more readily than those who adopted these technologies later in life, hence technology acceptance is more among young consumers. However, the relationship is not absolute -an experimental study of 239 shoppers found that older adults are not always less receptive to novel recommendations, and under some conditions show heightened receptiveness (Zhao & Fu, 2024). This evidence, while directionally clear, motivates the more careful test set out in Hypothesis H3.

2.7 Literature Review Summary

The reviewed literature converges on several major findings. Recommendation systems have evolved from early rule-based and collaborative-filtering models into sophisticated, self-learning AI engines (Section 3.1); AI-powered personalization delivers clear commercial value by improving relevance, engagement, and repeat purchasing (3.2); trust consistently emerges as the pivotal force that converts a recommendation into a purchase (3.3); privacy concern operates as the principal opposing force, expressed most sharply in the privacy paradox (3.4); purchase intention is jointly determined by trust, perceived usefulness, personalization quality, and privacy concern (3.5); and acceptance of personalization varies systematically across age groups, with younger "digital native" consumers generally more receptive (3.6).

The literature also contains important contradictions. The most fundamental is the privacy paradox itself - consumers simultaneously desire personalization and worry about the data collection that makes it possible, saying they value privacy while behaving as though they do not. A second tension appears in the demographic evidence: although younger consumers are generally more accepting of personalization,

experimental findings show that older consumers are not uniformly resistant and can, under certain conditions, be equally or more receptive to novel recommendations (Zhao & Fu, 2024). Personalization itself is shown to be double-edged, raising purchase intention up to a point and then depressing it once it is perceived as intrusive.

These tensions expose clear research gaps. Trust and privacy have largely been studied in isolation, each linked separately to purchase intention rather than modelled together as competing forces. Demographic factors such as age are frequently treated as secondary controls rather than as central moderators. Few studies integrate trust, privacy concern, and age into a single model, and fewer still do so in the context of present-day, AI-driven recommendation systems, where the tension between personalization and privacy is most acute.

This constitutes the justification for the current study. By examining how trust and privacy concerns jointly shape purchase intention -and whether this balance shifts across age groups -the present research directly addresses the integration gap the literature leaves open, offering a more complete account of how consumers move "from recommendation to purchase."

3. Methodology

3.1 Research Design

This study adopts a **quantitative research design**. A quantitative approach is one in which data is collected in numerical form and analysed using statistical techniques to test relationships between variables. It is appropriate here because the study's aim is not to explore why individuals feel a certain way in open-ended depth, but to measure the strength and direction of specific relationships - namely, how trust and privacy concerns relate to purchase intention, and whether these differ by age. Since the hypotheses (H1, H2, H3) are stated as measurable relationships between defined variables, a quantitative design allows them to be tested objectively, with results that can be expressed as scores, correlations, and comparisons rather than interpretation alone.

3.2 Research Approach

This research uses a **primary study** approach, meaning the data is collected first-hand by the researcher directly from respondents, rather than relying on data gathered by others. This is distinct from a secondary study, which reuses existing datasets, reports, or published statistics. A primary approach was chosen because no existing dataset captures the exact combination of variables this study examines - individual consumers' trust in recommendation systems, their privacy concerns, their purchase intentions, and their age - measured together. By designing and distributing an original questionnaire, the study generates data tailored precisely to its research questions and conceptual framework, ensuring the variables are measured in the specific way the hypotheses require. The approach is also **deductive**: it begins with theory and hypotheses drawn from the literature review and then collects data to test them, rather than building theory up from observation.

3.3 Data Collection Method

Data was collected through an **online questionnaire (survey)**, administered using **Google Forms**. A structured questionnaire is well-suited to a quantitative study because it presents every respondent with identical questions in a fixed order, producing standardised, directly comparable responses. The questionnaire is organised into six sections, mapped directly to the study's variables. It opens with a **demographic and behavioural section** (five multiple-choice items covering age group, gender, online

shopping frequency, platforms used, and a screening question confirming that the respondent has noticed personalized recommendations such as "recommended for you" or "you may also like")

The remaining sections measure the study's core variables using a **five-point Likert scale** (1 = Strongly Disagree to 5 = Strongly Agree), so that each attitude can be scored numerically and correlated:

- **Perceptions of personalization** - capturing how relevant, convenient, and useful respondents find recommendations, including one **reverse-coded** item
- **Trust** - measuring belief that recommendations and platforms act in the consumer's interest and handle data responsibly, including one **reverse-coded** item
- **Privacy concerns** - capturing unease about data collection, tracking, misuse, and lack of control.
- **Purchase intentions** - measuring willingness to act on and buy recommended products.
- **The trade-off** - directly probing the privacy paradox by asking respondents to weigh convenience against privacy, forming the bridge between the trust and privacy constructs central to this study.

The two **reverse-coded items** are deliberately worded in the opposite direction to detect inattentive responding; their scores are reversed before averaging so that, across every item, a higher number consistently indicates a stronger presence of the variable being measured.

3.4 Target Population

The target population consists of **consumers who shop online regularly and have used e-commerce platforms that employ recommendation systems**. This population is appropriate because the study concerns responses to personalized recommendations, so respondents must have actually encountered them. Relevant platforms include Amazon, Flipkart, Myntra, Ajio, and Meesho - all of which use "recommended for you" and similar features. A screening question at the start of the survey (confirming that the respondent shops online and has noticed personalized recommendations) ensures that only members of this relevant population are included, improving the validity of the responses.

3.5 Sampling Technique

The study used **convenience sampling**, in which respondents were recruited through an accessible pool - primarily personal and social networks reached via a shared survey link - rather than drawn from a defined sampling frame. This approach was adopted for three reasons. First, it made data collection feasible within the project's timeframe and resources: reaching a fully random sample of online shoppers would have required a sampling frame this study did not have access to. Second, it allowed the survey to reach respondents who had genuinely encountered personalised recommendations in practice, since the link circulated among people already active online. Third, it kept the process straightforward to administer through a single Google Form, consistent with the study's quantitative design. The trade-off is that convenience sampling does not guarantee that respondents' ages, attitudes, or behaviours reflect the wider population of online shoppers, and the sample's actual age distribution (Section 5.1) shows this limitation directly - respondents aged 35 and above make up 74.8% of the sample. This is addressed further in Section 4.8 (Study Limitations).

3.6 Sample Size

A target sample of 150 was set, which is adequate for detecting moderate correlations of the kind anticipated between trust, privacy concern and purchase intention, while remaining feasible within the project's timeframe. As noted in the limitations, the uneven age split that emerged left the group comparison less well powered.

3.7 Ethical Considerations

Because this study collects personal data and opinions from human participants, it was designed and cond-

ucted in accordance with standard ethical research principles. The following safeguards were applied. Participation was entirely **voluntary**, and **informed consent** was obtained before any data was collected. Respondents were told the purpose of the study, what their participation involved, and that they could withdraw at any point without giving a reason and without any consequence. No personally identifying information -such as name, contact details, or account information -was collected, ensuring respondent **anonymity**; where any identifying detail was unavoidable, it was kept **confidential** and accessible only to the researcher.

The study posed **no foreseeable physical, psychological, or financial risk** to participants; it involved only their opinions about a familiar everyday activity. Data was reported only in aggregate form, so that no single respondent's answers could be isolated.

A consideration specific to this study deserves note. Because the research itself concerns consumer privacy and the responsible handling of personal data, it was important that the study practise the standard it examines -collecting only the data necessary to answer the research questions, being transparent with respondents about how their information would be used, and giving them control over their participation. Conducting a study on privacy in a way that did not respect participants' own privacy would undermine its credibility. Respondents under 18 provided parental consent prior to participation

3.8 Study Limitations

While the study was designed to answer its research questions rigorously, several limitations should be acknowledged, both to interpret the findings responsibly and to guide future research.

First, the study uses a **cross-sectional design**: data was collected from each respondent at a single point in time. This allows the study to identify associations between trust, privacy concern, and purchase intention, but it cannot on its own establish **causation** or the direction of influence. This limitation is especially relevant to the adaptive-recommendation construct, which captures consumers' perception that recommendations improve over time rather than tracking that adaptation as it actually unfolds; a longitudinal design would be needed to observe change over time directly.

Second, the study relies on self-reported data, which is subject to well-known biases, including social-desirability bias (respondents answering as they feel they should) and recall inaccuracy. In particular, the dependent variable measures purchase intention rather than actual purchase behaviour, and stated intention is an imperfect proxy for what consumers do in practice. The scale is also not uniformly an intention measure: the item "I often click on recommended products" asks respondents to report past behaviour rather than future intent, so the construct blends the two rather than isolating intention cleanly.

Third, because all variables were measured using the same self-report questionnaire, the findings may be affected by **common-method bias**, whereby correlations between variables are inflated by the shared measurement method rather than reflecting a true underlying relationship.

Fourth, the sample was recruited through convenience sampling rather than a random or stratified method, which constrains the extent to which the findings can be generalised to the wider population of online consumers. As the demographic profile in Section 5.1 shows, the sample skews heavily toward older respondents (74.8% aged 35 and above) and toward male respondents (60.9%), so the results reflect this particular group more than online consumers as a whole. This skew also bears directly on the age comparison: 28 of the 32 respondents in the younger group were under 18, an age at which online purchases are frequently made with or by a parent rather than from independent discretionary income. The younger category therefore captures adolescents rather than young adult consumers, and the absence of a significant age difference should be interpreted with that composition in mind.

Finally, the study examines a defined set of constructs -trust, privacy concern, perceived quality, and age- and does not capture every factor that may shape purchase intention (such as price sensitivity, brand loyalty, or platform reputation). The model is therefore an intentionally focused account of the personalization trade-off, not a complete explanation of consumer purchasing behaviour.

3.9 Variables

Independent Variables

1. Trust in Recommendation Systems
2. Privacy Concerns

Dependent Variable

1. Purchase Intention

3.10 Questionnaire Structure

Section A: Demographics

- Age
- Gender
- Education
- Frequency of online shopping

Section B- Perceptions of personalization

- I trust product recommendations shown to me.
- Recommended products are generally relevant.
- Recommendation systems help me make better purchase decisions.

Section C- Trust

- I trust the platforms I shop on to act in my best interest.
- I believe recommendations are meant to help me, not just make me spend more.
- I trust that platforms handle my personal data responsibly.
- I feel recommendations are designed to manipulate me into buying. (reverse-coded)

Section D - Privacy concerns

- I am concerned about how much data shopping platforms collect about me.
- It bothers me that recommendations are based on my browsing and purchase history.
- I worry my personal data could be misused or shared without my consent.
- I feel I have little control over the data used to personalize my experience.

Section E - Purchase intentions

- I am more likely to buy a product if it is recommended to me.
- I often click on recommended products.

Section F - The trade-off

- The convenience of personalization is worth the privacy trade-off for me
- If I could turn off tracking, I would- even if recommendations got worse.
- . I would buy a recommended product even from a platform I don't fully trust.

Use a **5-point Likert Scale**:

1 = Strongly Disagree

5 = Strongly Agree

3.11 Reliability Testing

Before testing the hypotheses, the internal consistency of each composite measure was checked using Cronbach's alpha. This was needed because none of the four constructs were measured with a single

question. Trust, privacy concerns, purchase intention and positive attitude were each measured using a small group of Likert items which were then averaged into one score. Cronbach's alpha shows whether those items were actually measuring the same idea, and therefore whether averaging them was justified. One negatively worded items was reverse-coded first: Q14 ("recommendations are designed to manipulate"). Each composite was then calculated as the mean of its items, keeping the 1-5 scale so that every mean can be compared to the neutral midpoint of 3.00.

The usual benchmark for alpha is .70 (Nunnally, 1978), and this was the main criterion. However, alpha depends on the number of items as well as how strongly they correlate, so short scales tend to produce lower values even when they measure the same thing. Three of the four constructs here used only two or three items, so a second criterion of .60 was treated as acceptable for short exploratory scales (Nunnally, 1967; Hair et al., 2019). Scales in the .60–.69 band were kept but interpreted more cautiously.

The decision rule applied is summarised below.

Cronbach's alpha	Interpretation	Decision rule applied
$\geq .80$	Good internal consistency	Retain; interpret without qualification
.70 – .79	Acceptable internal consistency	Retain; interpret without qualification
.60 – .69	Marginal; acceptable for short exploratory scales	Retain; interpret with explicit caution
$< .60$	Unacceptable	Revise scale or exclude construct

Table 1. Cronbach's alpha and its interpretation

3.12 Data Analysis Techniques

All analyses were run on the 151 complete responses using SPSS. Timestamp and name fields were removed to keep responses anonymous. Items were coded 1 = strongly disagree to 5 = strongly agree, and a significance level of $\alpha = .05$ was used with two-tailed tests throughout.

Descriptive Statistics

Frequencies and percentages were calculated for the categorical variables (age group, gender, and whether respondents had noticed recommendation prompts). For each composite variable, the mean, standard deviation, minimum and maximum were calculated. Means were compared against the midpoint of 3.00, and the spread of each variable was checked for floor or ceiling effects. This matters because a variable with very little variation in the sample will produce a smaller correlation than it otherwise would, which affects how a non-significant result should be read.

Hypothesis Testing

Three hypotheses were tested. Because the study design is cross-sectional and observational, all hypotheses were operationalised as tests of association or of group difference rather than of causal influence, and the reporting language throughout reflects this.

H1 and H2 were both tested using Pearson correlation, the standard test for a linear relationship between two continuous variables, with Spearman correlation as a robustness check. Spearman does not assume linearity or interval-level measurement, so it helps confirm that a Pearson result is not just an artefact of those assumptions.

H3 was tested by splitting respondents into a younger group (24 or below) and an older group (25 or above) and comparing their mean attitude scores. Levene's test checked equality of variances. Welch's independent-samples t-test was used as the main test rather than Student's t-test, because the groups were very unequal in size ($n = 32$ and $n = 119$) and Welch's test does not assume equal variances or group sizes. Three further checks were run: a Mann-Whitney U test (non-parametric), a one-way ANOVA across all five original age brackets, and a Spearman correlation between the ordinal age variable and attitude. Hedges' g was used as the effect size, since it is preferred over Cohen's d for small, unequal groups.

Supplementary Multiple Regression

Trust and privacy concerns are related ideas and were expected to correlate with each other, so a multiple regression was also run with purchase intention as the dependent variable and both entered together as predictors. This shows whether each predictor still has its own relationship with purchase intention once the other is controlled for.

4. Results and Analysis

4.1. Demographic Profile

A total of 151 complete responses were obtained. There were no missing Likert responses and no exact duplicate response patterns, so all 151 cases were retained for analysis.

Age group	Frequency	Percent
Under 18	28	18.5%
18 – 24	4	2.6%
25 – 34	6	4.0%
35 – 44	45	29.8%
45 and above	68	45.0%
Total	151	100.0%

Table 2. Distribution of respondents by age group.

The age distribution is markedly uneven and this feature of the sample is consequential for the interpretation of H3 Respondents aged 35 and above account for 74.8% of the sample, while the 18-24 and 25-34 brackets together account for only 10 respondents (6.6%). The younger analytical group used to test H3 is therefore composed almost entirely of respondents aged under 18 (28 of 32 cases), meaning that the age comparison in this study is more accurately characterised as a comparison between adolescent and adult consumers than between young adult and older adult consumers.

Sample characteristic	Frequency	percent
Male	92	60.9%
Female	59	39.1%
Noticed recommendation prompts: Yes	142	94.0%
Noticed recommendation prompts: No	9	6.0%

Table 3. Gender composition and exposure to personalised recommendations.

The sample was male-skewed at 60.9%. Notably, 94.0% of respondents reported having noticed personalised recommendation prompts while shopping online, confirming that the phenomenon under study is near-universally encountered within this population and that respondents were answering from direct experience rather than hypothetically.

5.2 Descriptive Statistics

Composite scores were computed as the mean of each construct's items, with the two reverse-coded items recoded beforehand. All items used a five-point scale, so 3.00 represents the neutral midpoint. Table 4 reports the descriptive statistics together with the reliability figures discussed in Section 5.3.

Scale	Items	Mean	SD	Min	Max	Cronbach's α
Trust in recommendation systems	3	2.848	0.771	1.00	4.67	.630
Privacy concerns	4	4.046	0.659	2.00	5.00	.783
Purchase intention	2	3.010	0.793	1.00	4.50	.838
Positive attitude toward recommendations	3	3.517	0.625	1.33	5.00	.643

Table 4. Descriptive statistics and internal consistency of the composite scales (N = 151)

Three features of the distribution bear on the hypothesis tests. Privacy concern was the highest-scoring construct ($M = 4.05$), and no respondent averaged below 2.00, meaning concern about data collection was near-universal rather than an attribute separating one group of respondents from another. Trust was the only construct to fall below the midpoint ($M = 2.85$), while positive attitude sat well above it ($M = 3.52$): respondents evidently found recommendations useful while remaining unconvinced about the intentions of the platforms producing them. Purchase intention rested almost exactly at the midpoint ($M = 3.01$), indicating no overall tendency either toward or away from recommendation-led buying.

Table 5. Pearson intercorrelations among the principal study variables (N = 151)

Variable	1	2	3	4	5
1. Trust	1.000				
2. Privacy concerns	-0.362	1.000			
3. Purchase intention	0.332	-0.156	1.000		
4. Positive attitude	0.389	-0.177	0.339	1.000	
5. Age	0.012	0.204	-0.052	-0.139	1.000

Note. Age was coded ordinally across the five age brackets used in the questionnaire.

The strongest association in the matrix is the negative relationship between trust and privacy concern ($r = -.362$): respondents who worried more about data collection trusted the recommending platform less. Age was effectively unrelated to trust ($r = .012$) and to purchase intention ($r = -.052$) but was positively related to privacy concern ($r = .204$), a pattern returned to in Section 6.3.

5.3 Reliability Analysis

Applying the decision rule set out in Section 4.9, purchase intention ($\alpha = .838$) and privacy concerns ($\alpha = .783$) exceeded the .70 benchmark and are interpreted without qualification. Trust ($\alpha = .630$) and positive attitude ($\alpha = .643$) fell into the .60 to .69 band and are retained with explicit caution. Both are three-item scales containing a reverse-coded item, and alpha is depressed by short scale length and by reverse-coded items alike, so the lower figures reflect how the scales were constructed as much as the coherence of the underlying constructs. The practical consequence is that any correlation involving trust or attitude is likely to understate the true relationship, since measurement error attenuates observed coefficients.

5.4 Hypothesis Testing

All tests were two-tailed at $\alpha = .05$. Because three hypotheses were tested on the same dataset, a Bonferroni-corrected threshold of .0167 is also reported where a result sits close to the conventional boundary. Correlations for H1 and H2 are reported together in Table 6, and the joint regression bearing on both is reported in Table 7.

Table 6. Correlations of trust and privacy concerns with purchase intention (N = 151)

Predictor	Test	Coefficient	Sig. (2-tailed)	95% CI
Trust	Pearson $r(149)$	0.332	< .001	[0.182, 0.467]
	Spearman ρ	0.300	< .001	Not applicable
Privacy concerns	Pearson $r(149)$	-0.156	.056	[-0.308, 0.004]
	Spearman ρ	-0.166	.042	Not applicable

Table 7. Multiple regression coefficients predicting purchase intention

Predictor	B	SE B	β	t	Sig.	95% CI for B
Constant	2.282	0.546	–	4.180	< .001	[1.203, 3.361]
Trust	0.326	0.085	0.317	3.814	< .001	[0.157, 0.495]
Privacy concerns	-0.050	0.100	-0.041	-0.495	.621	[-0.247, 0.148]

Note. Model summary: $R = .334$, $R^2 = .112$, adjusted $R^2 = .100$, $SE = 0.752$; $F(2, 148) = 9.29$, $p < .001$.

5.4.1 H1 Results: Trust and Purchase Intention

Trust was positively and significantly associated with purchase intention, $r(149) = .332$, $p < .001$, 95% CI [.182, .467]. The confidence interval sits well clear of zero, the Spearman check agreed ($\rho = .300$, $p < .001$), and the result remains significant under Bonferroni correction. In the regression, trust remained a significant predictor once privacy concern was statistically controlled ($\beta = .317$, $p < .001$), so the bivariate correlation is not an artefact of privacy concern operating in the background. The unstandardised coefficient ($B = 0.326$) gives the relationship a concrete size: a one-point increase in trust is associated with roughly a one-third-point increase in purchase intention on the same five-point scale. H1 was supported.

5.4.2 H2 Results: Privacy Concerns and Purchase Intention

Privacy concern ran in the predicted negative direction but did not reach significance, $r(149) = -.156$, $p = .056$, with a 95% confidence interval of [-.308, .004] that includes zero. The Spearman coefficient was nominally significant ($\rho = -.166$, $p = .042$) but does not survive Bonferroni correction, and where the primary test and the robustness check disagree at the margin the conservative reading applies. Within the regression, privacy concern made no unique contribution once trust was entered ($\beta = -.041$, $p = .621$). H2 was not supported.

Two qualifications attach to this null result. First, the smallest correlation this sample could detect at $\alpha = .05$ with $N = 151$ is approximately $|r| = .16$, and the observed coefficient of $-.156$ falls just below that threshold; the finding is therefore inconclusive rather than evidence that no relationship exists. Second, privacy concern correlated with trust at $r = -.362$ and lost its independent contribution once trust entered the model, which is the pattern expected if privacy concern shapes purchase intention through trust rather than acting on it directly. Establishing this would require a formal mediation analysis, which the present design did not include.

5.4.3 H3 Results: Age Differences in Attitudes Toward Personalisation

Table 8. Age comparison of positive attitudes toward personalised recommendations

Group / Test	qn or statistic	Mean or df	Sig.	Note
Younger (24 or below)	$n = 32$	$M = 3.635$ (SD 0.489)	–	Mean difference = 0.150

Older (25 or above)	n = 119	M = 3.485 (SD 0.655)	–	
Levene's test	2.528	–	.114	Unequal variance not detected
Welch's t-test	1.433	df = 64.19	.157	Hedges' g = 0.24 (small)
Mann-Whitney U	2145.5	–	.261	Non-parametric check
One-way ANOVA (5 brackets)	1.967	df = 4, 146	.103	No overall group difference
Spearman age trend	-0.179	df = 149	.028	Not significant after Bonferroni

Younger respondents did report more positive attitudes than older respondents, which is the direction H3 predicted, but the gap of 0.15 scale points was not statistically significant, Welch's $t(64.19) = 1.43$, $p = .157$, and the effect size was small, Hedges' $g = 0.24$. The supplementary tests agreed: the one-way ANOVA across all five age brackets was non-significant ($p = .103$), as was the Mann-Whitney U test ($p = .261$). H3 was not supported.

Two qualifications are important. Given an effect of $g = 0.24$ and group sizes of 32 and 119, the achieved power of the test was approximately .22, and detecting a difference of that size at the conventional .80 standard would have required roughly 274 respondents per group. The test was therefore unable to detect a small difference, which is not the same as having shown that none exists. Further, as noted in the demographic profile, the younger group consisted almost entirely of respondents under 18, so the comparison performed is closer to adolescents against adults than young adults against older adults. The Spearman trend across the full five-category age ordering was negative and nominally significant ($\rho = -.179$, $p = .028$) though not Bonferroni-significant; since dichotomising an ordinal variable discards information, this trend may be the more informative of the two analyses.

6. Discussion

This section interprets each hypothesis in turn and then draws the three results together. Because the design is cross-sectional and observational, the discussion describes associations and group differences rather than causal effects, and the language used reflects this throughout.

6.1 Discussion of H1

The mechanism usually offered for the trust-purchase relationship is uncertainty reduction. Online purchasing requires the consumer to act on incomplete information: the product cannot be handled, the seller's motives cannot be observed, and the reasoning behind an AI-generated recommendation is invisible. Trust substitutes for the information the shopper does not have, lowering perceived risk so that a recommendation can be treated as advice rather than as an attempt to sell. The content of the trust items supports this reading, since they asked whether platforms handle data responsibly, whether recommendations are meant to help rather than simply to increase spending, and whether they are designed

to manipulate. What is being measured is close to perceived benevolence: whether the system is understood to be working for the consumer or on them, which is precisely the ability-benevolence-integrity structure identified by Hsiao et al. (2010).

The result aligns with the established position of trust in this literature, including the meta-analysis of nineteen studies reporting that trust positively influences purchase intention (Wang et al., 2022) and the theoretical claim that trust matters most where perceived risk is high (Curràs-Pérez et al., 2017). The magnitude observed here, $r = .332$, sits toward the lower end of the reported range. Two features of the present study plausibly account for this: the trust scale had only moderate reliability, and attenuation means the underlying relationship is likely somewhat stronger than the coefficient suggests; and respondents were asked about recommendation systems in general rather than about a single platform they actively used, which introduces variability a platform-specific study would avoid.

6.2 Discussion of H2

H2 followed from the privacy calculus, in which consumers weigh the benefit of personalisation against the perceived cost of the disclosure that makes it possible (Awad & Krishnan, 2006). The tension is structural rather than incidental, since a recommendation system becomes more accurate as it collects more behavioural data, and that collection is exactly what concerns the shopper. The data here pointed in the predicted direction but did not confirm the prediction, and three explanations are worth considering.

The privacy paradox. The gap between stated privacy concern and protective behaviour is well documented (Norberg et al., 2007; Barth & de Jong, 2017), with Kokolakis (2017) noting that consumers say they value privacy while behaving as though they do not. This sample fits that description closely: 94% had noticed recommendation prompts, mean privacy concern was 4.05 out of 5, and yet purchase intention sat at the midpoint and was not meaningfully lower among the more concerned. A concern that is widely held but does not register in behavioural intention is the paradox itself rather than an anomaly in the data.

Restricted variance. Privacy concern had the highest mean and the narrowest observed range of the four scales, with no respondent scoring below 2.00. A variable on which almost everyone scores highly has little variation left with which to explain variation in something else, so privacy concern may matter a great deal to consumers and still correlate weakly with purchase intention in a sample where nearly all of them share it. This is a property of the sample rather than a statement about the construct.

Indirect operation through trust. The most substantively interesting reading is visible in the regression. Privacy concern correlated with trust at $r = -.362$ and its contribution vanished once trust was entered, which is what would be expected if privacy concern shapes purchase intention by first shaping trust. This is consistent with Bandara et al. (2021), who position weak privacy protection as damaging consumer trust and only then producing defensive responses. On that reading privacy concern is not irrelevant; it is an input into trust, and trust is what the consumer acts on. The correct conclusion is therefore narrow: in this sample, self-reported privacy concern did not by itself predict purchase intention to a degree the study could detect.

6.3 Discussion of H3

H3 rested on the digital natives argument set out in Section 3.6: consumers who grew up with algorithmic curation should treat personalisation as an ordinary convenience, while those who encountered it later should weight the perceived risk more heavily. The data supported this weakly and no further, with a difference in the predicted direction that none of the four tests could distinguish from chance. Three considerations explain the outcome, and they differ in what they imply.

The first is the composition of the younger group, which consisted of 28 respondents under 18 and only 4 aged 18 to 24, so the category does not represent young adult consumers with independent purchasing power. The comparison the test performed is not quite the comparison H3 was framed to make. The second is statistical power, which at approximately .22 left the study with roughly a one-in-five chance of detecting a difference of the observed size even if one genuinely exists; a non-significant result under these conditions is close to uninformative. The third is substantive rather than methodological: familiarity is the mechanism through which age is supposed to operate, and 94% of the sample across every age bracket reported noticing recommendation prompts. If exposure to personalisation has become near-universal, the generational gap in familiarity that the digital natives argument assumes may simply have narrowed. That would be a finding of interest in itself, but it cannot be claimed from an underpowered test on an unbalanced sample.

One nuance is worth preserving. Age was not inert in this dataset: it correlated positively with privacy concern ($r = .204$) and showed a weak negative trend with attitude across the full age ordering. The suggestion is that age relates more to how consumers feel about the data collection underlying personalisation than to how positively they judge the recommendations themselves. This is consistent with Zhao and Fu (2024), whose experimental finding that older adults are not uniformly less receptive already cautioned that the age effect is real in direction but not absolute.

6.4 Overall Interpretation

Taken together, the three results describe the personalisation-privacy trade-off differently from the way it is usually framed. The standard account runs in three steps: personalisation requires data, data collection generates concern, and concern depresses engagement. This study is consistent with the first two links and none for the third. Concern was near-universal at $M = 4.05$ and did not translate into lower purchase intention, whereas trust tracked purchase intention both as a correlation ($r = .332$) and as a regression predictor ($\beta = .317$). The more accurate description of the trade-off in this dataset is therefore not personalisation against privacy but personalisation against trust, with privacy concern operating as one of the things that erodes trust rather than as a direct brake on purchasing. The gap between the attitude and trust means sharpens the point: respondents rated recommendations as useful ($M = 3.52$) while rating the platforms behind them below the midpoint on trustworthiness ($M = 2.85$). They found the output helpful and doubted the motive, and usefulness alone did not appear to be enough.

The size of the effect should be kept in proportion. The model explained 11.2% of the variance in purchase intention, leaving the large majority unaccounted for by price, product need, brand familiarity, reviews and prior experience, none of which were measured here. The defensible claim is that trust is one meaningful lever among several rather than the determinant of online purchasing. Three implications follow for e-commerce firms.

- Interventions should target trust rather than privacy concern directly. Reducing stated concern is difficult, and on this evidence it may not change purchase intention even if it succeeds.
- Explanation is likely to matter more than accuracy alone. The trust items turned on whether recommendations were meant to help or to extract spending, and stating plainly why an item was recommended addresses that doubt at low cost. Visible data controls serve the same function, since controls buried in settings menus do not register with users and cannot contribute to trust.
- Older customers should not be treated as a segment that dislikes personalisation. Age did not predict attitudes toward recommendations; what it predicted was privacy concern, which calls for reassurance about data handling rather than a different recommendation experience.

7. Implications

7.1 Managerial Implications

The results point to one main conclusion for e-commerce firms: trust, not privacy concern, was the variable that tracked purchase intention ($\beta = .317$, $p < .001$, against $\beta = -.041$, $p = .621$). Effort is better spent building trust than trying to reduce stated concern, which was near-universal in this sample and did not by itself depress willingness to buy.

- **Explain the recommendation.** The trust items turned on whether recommendations were seen as helping the shopper or extracting spending. Stating why an item was suggested addresses that doubt directly and at low cost, which matters given that 65% of consumers do not realise recommendations are AI-generated (DataStax, 2023).
- **Make data controls visible.** One privacy item concerned perceived lack of control. Controls that exist but sit buried in settings menus cannot contribute to trust.
- **Treat privacy policy as a trust instrument.** Privacy concern correlated with trust at $r = -.362$ and lost its independent effect once trust was controlled, so stronger data practices are most plausibly valuable through the trust they build rather than through any direct effect on purchasing.
- **Do not segment by age.** Age did not predict attitudes toward recommendations ($p = .157$) but did correlate with privacy concern ($r = .204$). Older customers need reassurance about data handling, not a different recommendation experience.
- **Keep expectations proportionate.** The model explained 11.2% of variance in purchase intention. Trust is one lever among several, alongside price, need and brand.

7.2 Consumer Implications

Respondents rated recommendations as useful ($M = 3.52$) while rating the platforms producing them below the midpoint on trustworthiness ($M = 2.85$), and 94% had encountered recommendation prompts. Consumers are therefore acting on systems they use constantly and do not fully trust. The findings suggest two things: awareness of what data drives personalisation is uneven, and stated concern is not translating into protective behaviour. Informed decision-making requires consumers to recognise that a recommendation is an output of their own behavioural data, and to use the privacy settings platforms provide rather than accepting defaults.

7.3 Policy Implications

- **Data governance.** Consent mechanisms that produce high stated concern alongside unchanged behaviour are not functioning as meaningful consent. Regulation focused on defaults and data minimisation is likely to be more effective than disclosure alone.
- **Ethical AI and explainability.** The gap between perceived usefulness and trust suggests a case for requiring platforms to disclose that recommendations are algorithmic and what data informs them.
- **Consumer protection.** Since privacy concern appears to operate through trust, protection frameworks should be assessed on whether they restore consumer confidence, not only on compliance.

8. Limitations

The limitations set out in Section 4.9 apply throughout. Four further limitations emerged from the analysis itself.

- **Sample composition.** Respondents aged 35 and above made up 74.8% of the sample, and the younger group used to test H3 was composed almost entirely of respondents under 18 (28 of 32). The age

comparison is closer to adolescents against adults than young adults against older adults, which is not the comparison H3 was framed to make. The sample was also male-skewed (60.9%).

- **Insufficient statistical power.** Achieved power for H3 was approximately .22; roughly 274 respondents per group would have been needed at the conventional .80 standard. Similarly, the smallest correlation detectable at $N = 151$ was $|r| = .16$, and H2's coefficient of $-.156$ fell just below it. Both null results are inconclusive rather than evidence of no relationship.
- **Restricted variance in privacy concern.** With a mean of 4.05 and no respondent averaging below 2.00, there was little variation for privacy concern to explain, which independently weakens any correlation involving it.
- **Moderate scale reliability.** Trust ($\alpha = .630$) and positive attitude ($\alpha = .643$) fell below the .70 benchmark, so correlations involving them are likely attenuated.

9. Future Research Directions

- **Mediation testing.** The clearest extension: formally test whether privacy concern affects purchase intention through trust, which is the pattern the regression suggested but could not confirm.
- **Longitudinal design.** RO4 concerned how perceptions of adaptive recommendations evolve, which a single-point survey cannot capture. Tracking the same consumers over time would allow direction of influence to be established.
- **Age-balanced sampling.** A replication with adequate representation in the 18-24 and 25-34 brackets would give H3 a fair test, and would resolve whether the near-universal exposure found here has genuinely narrowed the generational gap.
- **AI explainability as a variable.** Given that trust was the operative predictor, manipulating how much explanation accompanies a recommendation would test whether transparency actually raises trust.
- **Gender and cross-country comparison.** Neither was tested here; a male-skewed Indian sample cannot speak to either.
- **Behavioural outcomes.** Replacing stated intention with observed clicks or purchases would test whether the privacy paradox found here holds when real money is involved.

10. Conclusion

This study set out to examine how trust and privacy concerns jointly relate to purchase intention in response to AI-driven recommendations, and whether that relationship varies by age (RO1-RO3). Survey data from 151 respondents produced a clear pattern across the three hypotheses. H1 was supported: trust in recommendation systems was positively and significantly associated with purchase intention, $r(149) = .332$, $p < .001$, and remained significant when privacy concern was controlled. H2 was not supported: privacy concern ran in the predicted negative direction but did not reach significance, $r(149) = -.156$, $p = .056$, and contributed nothing unique once trust entered the model. H3 was not supported: younger respondents held marginally more positive attitudes, but the difference was small and non-significant across all four tests.

The study's contribution follows from modelling both forces together, which the literature has rarely done. Doing so reframes the personalisation-privacy trade-off. The standard account holds that data collection generates concern and concern depresses engagement; this study found the first link and not the second. Concern was near-universal yet did not lower purchase intention, while trust tracked it consistently. The trade-off, on this evidence, is better described as personalisation against trust, with privacy concern acting

as one input into trust rather than as a direct brake on purchasing. This offers the privacy paradox literature a possible mechanism for why stated concern so rarely predicts behaviour, and offers the effectiveness literature a reason to treat trust rather than relevance as the binding constraint on recommendation-led purchasing.

Bibliography

1. Awad, N. F., & Krishnan, M. S. (2006). The personalization privacy paradox. *MIS Quarterly*, 30(1), 13–28.
2. Bandara, R., Fernando, M., & Akter, S. (2021). Managing consumer privacy concerns and defensive behaviours in the digital marketplace. *European Journal of Marketing*, 55(1), 219–246. <https://doi.org/10.1108/EJM-06-2019-0515>
3. Barth, S., & de Jong, M. D. T. (2017). The privacy paradox. *Telematics and Informatics*, 34(7), 1038–1058.
4. Bennett, J., & Lanning, S. (2007). The Netflix Prize. *Proceedings of KDD Cup and Workshop*.
5. Carrascal, J. P., Riederer, C., Erramilli, V., Cherubini, M., & de Oliveira, R. (2013). Your browsing behavior for a Big Mac. *Proceedings of WWW '13*.
6. Goldberg, D., Nichols, D., Oki, B. M., & Terry, D. (1992). Using collaborative filtering to weave an information tapestry. *Communications of the ACM*, 35(12), 61–70.
7. Hair, J. F., Black, W. C., Babin, B. J., & Anderson, R. E. (2019). *Multivariate Data Analysis* (8th ed.). Cengage.
8. Handoyo, S. (2024). Purchasing in the digital age: A meta-analytical perspective on trust, risk, security, and e-WOM in e-commerce. *Heliyon*, 10(8), e29714. <https://doi.org/10.1016/j.heliyon.2024.e29714>
9. Hsiao, K.-L., Lin, J. C.-C., Wang, X.-Y., Lu, H.-P., & Yu, H. (2010). Antecedents and consequences of trust in online product recommendations. *Online Information Review*, 34(6), 935–953.
10. Iyengar, S. S., & Lepper, M. R. (2000). When choice is demotivating. *Journal of Personality and Social Psychology*, 79(6), 995–1006.
11. Jeong, C. H., Jeong, Y. S., & Lee, K. (2006). Research about the effects of internet shopping mall user trust on royalty: Comparison analysis between Korea and China trust determinant factors. *Spring Semi-Annual Conference of KMIS*.
12. Kokolakis, S. (2017). Privacy attitudes and privacy behaviour. *Computers & Security*, 64, 122–134.
13. Koren, Y., Bell, R., & Volinsky, C. (2009). Matrix factorization techniques for recommender systems. *Computer*, 42(8), 30–37.
14. Marsh, S. P. (1994). Formalising trust as a computational concept [Doctoral dissertation, University of Stirling].
15. Norberg, P. A., Horne, D. R., & Horne, D. A. (2007). The privacy paradox. *Journal of Consumer Affairs*, 41(1), 100–126.
16. Nunnally, J. C. (1978). *Psychometric Theory* (2nd ed.). McGraw-Hill.
17. Resnick, P., Iacovou, N., Suchak, M., Bergstrom, P., & Riedl, J. (1994). GroupLens: An open architecture for collaborative filtering of netnews. *Proceedings of ACM CSCW*.
18. Resnick, P., & Varian, H. R. (1997). Recommender systems. *Communications of the ACM*, 40(3), 56–58.

19. Salakhutdinov, R., Mnih, A., & Hinton, G. (2007). Restricted Boltzmann machines for collaborative filtering. Proceedings of ICML.
20. Sarwar, B., Karypis, G., Konstan, J., & Riedl, J. (2001). Item-based collaborative filtering recommendation algorithms. Proceedings of WWW '01.
21. Saxborn, M., Pan, Y., & Said, A. (2024). Trust through recommendation in e-commerce. Proceedings of the 2024 ACM SIGIR Conference on Human Information Interaction and Retrieval (CHIIR '24), Sheffield, UK. <https://doi.org/10.1145/3627508.3638294>
22. Wang, J., Shahzad, F., Ahmad, Z., Abdullah, M., & Hassan, N. M. (2022). Trust and consumers' purchase intention in a social commerce platform: A meta-analytic approach. SAGE Open, 12(2). <https://doi.org/10.1177/21582440221091262>
23. Wu, K.-W., Huang, S. Y., Yen, D. C., & Popova, I. (2012). The effect of online privacy policy on consumer privacy concern and trust. Computers in Human Behavior, 28(3), 889–897. <https://doi.org/10.1016/j.chb.2011.12.008>
24. Zhang, S., Yao, L., Sun, A., & Tay, Y. (2019). Deep learning based recommender system. ACM Computing Surveys, 52(1), 1–38.
25. Zhao, L., & Fu, B. (2024). Assessing the impact of recommendation novelty on older consumers: Older does not always mean the avoidance of innovative products. Behavioral Sciences, 14(6), 473. <https://doi.org/10.3390/bs14060473>
26. Gupta, A., & Arora, N. (2017). Understanding determinants and barriers of mobile shopping adoption using behavioral reasoning theory. Journal of Retailing and Consumer Services, 36, 1–7. <https://doi.org/10.1016/j.jretconser.2016.12.012>
27. Olaleye, S. A., Sanusi, I. T., & Adepoju, B. (2018). Actual use and continuous use of retail mobile app: A model comparison perspective. Advances in Science, Technology and Engineering Systems Journal, 3(6), 151–158. <https://doi.org/10.25046/aj030619>
28. Rahmayanti, P. L. D., Widagda, I. G. N. J. A., Yasa, N. N. K., Giantari, I. G. A. K., Martaleni, Sakti, D. P. B., Suwitho, & Anggreni, P. (2021). Integration of technology acceptance model and theory of reasoned action in predicting e-wallet continuous usage intentions. International Journal of Data and Network Science, 5(4), 649–658. <https://doi.org/10.5267/j.ijdns.2021.8.002>
29. Venkatachalam, P., & Ray, S. (2022). How do context-aware artificial intelligence algorithms used in fitness recommender systems? A literature review and research agenda. International Journal of Information Management Data Insights, 2(2), 100139. <https://doi.org/10.1016/j.jjime.2022.100139>
30. Pappas, I. O., Giannakos, M. N., & Chrissikopoulos, V. (2012). Personalized Services in Online Shopping: Enjoyment and Privacy. International Conference on Information Society (i-Society 2012). The study used **148 online customers** and found that privacy negatively affected both enjoyment and purchase intention.
31. Priyanka & Siraj (2021). Unlocking the Personalized-Privacy Paradox: An Empirical Study of Customer Intention to Use Digital Personalized Services. Academy of Strategic Management Journal, 20(Special Issue 6). The study of Indian consumers reports $r = -0.609$, $p < 0.001$ for intention to use online personalization and privacy concern.

32. **Selem, K. M., Shoukat, M. H., Shah, S. A., et al. (2023).** The dual effect of digital communication reinforcement drivers on purchase intention in the social commerce environment. *Humanities and Social Sciences Communications*, 10, 47. The study collected **490 responses from Pakistan's social-media users** and found that privacy concerns significantly affected purchase intention.