

Real-Time AI Personalisation, Dynamic Pricing, and the Consumer Comfort Threshold in E-Commerce

Himani Taneja

Student, Bachelor of Business Administration (Marketing), NMIMS Navi Mumbai

Abstract

The technology of artificial intelligence (AI) has transformed the field of personalisation in e-commerce from employing somewhat visible recommendation methods to relying on cross-context decision-making. The use of information inferred from sources beyond a shopper's direct platform activity substantially changes existing tensions surrounding the use of specific personalisation technologies. The current research intends to define the limits of e-commerce personalisation. In this regard, a survey was conducted from July 7 to July 16, 2026, which included 51 respondents. The survey was organised using a within-subjects scenario design and included five types of personalisation, including history-based recommendations and personalised targeting. The data were analysed applying various statistical techniques, including the Friedman test, Wilcoxon signed-rank test, Mann–Whitney U test, and Spearman rank correlation. The data revealed differences in comfort across the five scenarios ($\chi^2(4) = 35.57$, $p < 0.001$, Kendall's $W = 0.17$), a pattern described as the comfort cliff. The generic scenario was much more comfortable than the other scenarios (all $p < 0.001$), but no comparison between the other four scenarios was significant. These results favour the origin-based understanding of the personalisation-privacy paradox, as it seems to hinge on whether people think the platform uses information they produced or information obtained in other ways. Trust was strongly correlated with comfort ($\rho = 0.57$, $p < 0.001$), whereas prior knowledge of dynamic pricing was not.

Keywords: AI Personalisation, E-Commerce, Dynamic Pricing, Privacy, Consumer Trust, Personalisation-Privacy Paradox, Explainable AI, Price Fairness, Algorithmic Decision-Making

1. Introduction

1.1 Background

The field of e-commerce has evolved from basic online shops to sophisticated systems that can identify users' preferences. Recommender systems use collaborative filtering and content-based algorithms, contextual information, and machine learning in their operations. In a paper by Raza et al. [9], the authors state that recommender systems have evolved from simple filtering methods to complex, large-scale methods used across various application domains and real-time environments.

The concept behind the enterprise is very basic: relevance lowers search costs, speeds up the search process, and increases consumer engagement. Bleier and Eisenbeiss [1] conducted research and discovered that personalisation enhances ad effectiveness, though its success is affected by various aspects including timing, placement, and the degree of personalisation.

The use of artificial intelligence (AI) in this context increases the tension between the data available for analysis and how consumers react to analysis made using it. The underlying data composition has advanced to the point where signals once considered incompatible can be integrated. Consumers might feel ambivalent when they realise an e-commerce company remembers their search history. Still, they might have a completely different perspective when they see an ad based on their private conversations, location, or activities on the social network.

Mehmood et al. [8] assert that AI-enabled personalisation is of great importance to consumers and society, especially given the need to understand how consumers interpret such systems. Therefore, personalisation is a matter of technological optimisation, but also of trust and governance. Evaluation of the personalisation experience depends on the perceived trade-off between the benefits gained and the risks involved.

1.2 Problem Statement

The main issue studied in this research is that existing studies on personalisation tend to describe buyer reactions as a process: the level of personalisation increases relevance until privacy issues outweigh the gains. Researchers concluded that consumer responses may not be uniform. For instance, Xu and others [13] differentiated between overt and covert personalisation in the context of privacy calculation. They showed that the perceived nature of personalisation affects the perceived level of risk and advantages. Boerman et al. [2] proved that personalisation can exceed restrictions imposed on a person, depending on the type of data used, whether the data is shared, and whether dynamic pricing is used.

A recent study provides a particularly pertinent conceptual framework for the discussion. The authors of that study, Kim and Han [6], conducted an experiment utilising low-, moderate-, and highly personalised situations to demonstrate that information personalisation can backfire when privacy worries are triggered. De Keyser et al. [3] have also established a turning point between low- and moderate-personalisation phenomena when customers exhibit more “creepy” perceptions. Therefore, consumers may not straightforwardly react to increasingly progressive personalisations.

Thus, the researchers behind the current study will raise a more precise issue: when personalisation becomes more of an intrusion than useful service, what factors influence consumers’ views: level of personalisation or the context with which data was collected? The study will involve dynamic pricing methods, as personalised pricing allows for introducing the second mode of personalisation being applied – this time, personalisation will not only influence what a consumer sees. Still, it will also alter the economic conditions of purchase.

1.3 Research Objectives and Questions

The research intends to achieve three goals. The first goal is to conduct a study to determine whether customer comfort varies systematically among different levels of personalisation. The second goal is to identify the point of maximum change in comfort and investigate whether it coincides with the distinction between intentionally collected data and covertly collected inferences. The third goal is to check whether knowledge of dynamic pricing and customers’ initial trust in the platforms have any bearing on customers’ comfort with high-intensity personalisation.

The research questions are:

RQ1: Do customers feel comfortable with the idea of AI-mediated personalisation?

RQ2: Where is the cut-off point between acceptable and intrusive levels of personalisation?

RQ3: Is knowledge about dynamic pricing important for customer comfort levels associated with the use of AI?

1.4 Significance of Study

The significance of the research paper can be interpreted from three different aspects. Firstly, it enables scholars in the marketing field to determine whether the concept of the personalisation–privacy paradox can be understood as a continuous trade-off or a threshold phenomenon. Secondly, the paper demonstrates the relevance of data provenance and context integrity to practitioners when formulating recommendations, advertisements, and personalised prices. Lastly, the paper helps users and persons engaged in policy formation to understand the significance of transparency and substantial control, rather than only greater accuracy in personalisation.

More importantly, the research does not establish a universal numeric threshold that applies to all users or e-commerce marketplaces. The sample is small and based on convenience sampling of participants in the research. Nonetheless, the main idea is conceptual and anchored in empirical findings: within this sample, the primary dividing line is not between “low,” “medium,” and “high” personalisation in the abstract.

2. Literature Review

2.1 Real-Time Data Pipelines and AI Infrastructure

Advanced recommender systems yield a promising technical basis for large-scale personalisation. Raza et al. [9] trace the evolution of recommender systems from content filtering and joint filtering to advanced approaches, including deep learning, graph-based systems, reinforcement learning, and large language models. Their review notes the importance of scalable, real-time, and reliable recommender systems. In the e-commerce sphere, a customer interaction may be introduced into the recommendation system almost in real time.

The real-time architecture changes the way consumers act in two major ways. First, it minimises the time gap between people's behaviour and personalisation processes. A person looking for a product may receive recommendations or ads almost immediately. Secondly, the systems can combine several signals at once, which means the personalisation provided to a customer may be more precise than any individual piece of information.

The application of artificial intelligence in personalised recommendations leads to an information imbalance. Other parties can access information that consumers are unaware of. Mehmood et al. [8] provide evidence that consumers' interpretations of AI-based personalisation are subjective and vary among cultures and social contexts. This is important because perceived personalisation matters not only for the effectiveness of the recommendation but also for the consumer's understanding of how the system collects the necessary information.

The technology itself is not considered problematic, as personalised recommendations that reflect consumers' previous purchases and browsing experience can be treated as useful and reliable. The problem arises when technological capability surpasses contextual expectations. A recommendation may be accurate but not perceived positively if the consumer does not understand how the data was collected.

This difference prompted the current research on data provenance. Rather than presuming that only the level of personalisation causes inconvenience, this study considers various situations that differ in what is known about the source of information and how it was obtained. The first situation is that of a recommendation based on the user's history. The following situations involve the use of overheard conversations, location, price discrimination, and targeting with personal information. Thus, the design shifts from the technical question of how companies use data to the consumer question of which kinds of data use are perceived as acceptable.

2.2 Dynamic Pricing and Price-Fairness Theory

Dynamic pricing is understood to mean the variation of prices along the axes of time, demand conditions, customer features, or quantitative information, all of which are subjected to algorithmic processing. The emergence of the digital economy and the ability to gather finer levels of information increases the potential for greater price differentiation.

Regarding price fairness, Xia et al. [12] note that consumers assess prices relative to one another and evaluate the reasons for price differences. Hence, a price might be seen as unfair by consumers if the basis of the variance is considered to violate a principle of fairness, even if the firm's pricing decision can be economically justified.

New evidence complicates the assumption that people find algorithmic pricing less fair. In a series of experiments conducted by Duani et al. [5], consumers sometimes indicated that demographic price discrimination is fairer when implemented by an algorithm than by a human. This might be explained by a perception that algorithms do not make judgments as personally as people do and do not exploit consumers as much. Therefore, the perceived fairness of algorithmic pricing depends on how it is done, why consumers are charged differently, and what they think about the decision-maker.

The presented research uses a very specific comparison involving interpersonal interactions: a consumer learns that their friend paid less for the same product they bought at the same time through the same platform. This situation brings the fairness aspect to the surface since the consumer has a reference point. It also provides an opportunity to analyse algorithmic pricing through the lens of wider personalisation.

The data indicating that 58.8% of those surveyed found algorithm-led pricing unfair compared to human-led pricing should be treated as sample-based findings rather than a universal principle. It does not mean that, from the respondents' sample standpoint, the use of algorithms should lend legitimacy to the process. Instead, the group holds a negative attitude toward algorithm-led pricing, which aligns with the study's well-known conclusion that fairness depends on the process.

2.3 Personalisation-Privacy Paradox and the Backfire Effect

The personalisation-privacy paradox refers to the conflict between consumers' need for relevant, convenient personalisation of offers and their concerns about the use of their private information. According to Dinev and Hart [4], privacy decisions are viewed as a rational calculation in which privacy risk is balanced against benefits, trust, and motivations. Consumers may accept some degree of data collection to enjoy the desired convenience despite their stated privacy worries.

Xu et al. [13] further develop the idea and argue that personalisation may work differently when done transparently compared to when performed covertly. It is important to note that this difference is significant for the current study. It is easier for consumers to understand a recommendation based on their own inputs than one based on information collected from other sources.

The study conducted by Bleier and Eisenbeiss [1] has shown that the effect of personalisation may decrease at excessive levels of personalisation. This study has also stated that the effectiveness of personalisation depends on the context and timing of its use, and that understanding privacy issues can clarify why personalisation does not always outperform non-personalised tools.

The research done by Kim and Han [6] is theoretically close to this research. They have divided the extent of personalisation into three levels, and the level of privacy matters, which they call the intrusiveness ladder. The results of this study showed that personalisation may be ineffective due to privacy concerns, thereby clarifying the strategy's limits.

De Keyzer et al. [3] wrote that perceived creepiness acts as a mediator between the degree of personalisation and subsequent consumer outcomes. It is important to note that the authors found that the critical degree of personalisation is reached within the range of low and medium personalisation, denoting a sudden consumer reaction rather than a gradual one. Boerman et al. [2] also identify limits to the acceptability of personalised advertising, such as the use of individual-specific information, information sharing, and high prices for personalised products.

The present research expands the established body of literature by posing the question of whether the threshold can be treated as an intensity threshold or a provenance threshold. If, during high-intensity situations, consumers feel less comfortable than in each lower-intensity situation, then the intensity model is warranted. If one situation is distinct from the others in the group of high-intensity situations, then the provenance or environment model applies.

2.4 Trust, Transparency, and Explainable AI

Consumers' participation in data-dependent commerce depends on their trust. Smith, Milberg, and Burke [10] have made a considerable contribution to the study of individuals' anxiety regarding organisations' information practices, such as collection, secondary use, unauthorised access, and error. Such a study remains relevant because consumers would assess how organisations use data rather than only whether data have been collected.

Dinev and Hart [4] provide evidence that trust mitigates some of the privacy worries that may arise in privacy calculus. Tucker [11] provides real-world evidence supporting the assertion that the perception of control over one's personal information influences the efficiency of personalised marketing. This study is based on a field experiment in which enhanced perceived control over privacy increased the effectiveness of targeted advertising.

Transparency may also diminish uncertainty. Kim et al. [7] explored the "Why am I seeing this ad?" concept and found that increased transparency may improve the performance of advertising messages. However, this does not necessarily mean that every 'explanation' would lead to greater trust. Rather, if a user is provided with a clear explanation of the message selection, that would influence the interpretation of said message.

In line with those premises, the present study included one question on whether respondents believe that the "Why am I seeing this?" explanation would improve their level of trust. Although the overwhelming majority (74.5%) of respondents selected either "Yes" or "Maybe," the exploratory analysis did not detect any statistically significant differences in trust levels among the three respondent groups. Therefore, even though consumers are in favour of transparency, this does not necessarily imply that it would translate into trust.

Research by Boerman et al. [2] and De Keyzer et al. [3] further strengthens the view of perceived creepiness and its limits. If personalisation is considered inappropriate, an explanation will have a positive effect only if it is easily understood and the approach followed is also accepted. Transparency cannot validate what consumers see as totally inappropriate.

2.5 Research Gap

The literature on personalisation establishes several important propositions: personalisation can increase perceived relevance; privacy worries reduce acceptance of personalisation; feelings of creepiness can arise when personalisation crosses perceived boundaries; feelings of trust and control over data can moderate these responses; and personalisation in pricing may lead to perceptions of unfairness. Less established is

how these various mechanisms interact inside a single test of personalisation, comparing different forms of personalisation in the same respondent.

Much of the prior research examines each of these areas separately. Experimental studies often manipulate the level of personalisation while holding the data context constant. This study, however, places different data-provenance cues side by side within the same experiment: personalisation based on behavioural history, information overheard in conversations, personalisation based on price disparities, physical location, and personal details. This allows a within-subjects comparison of whether respondents feel comfortable with personalisation based on all these data points, or whether they feel differently about each of them. The study fills a specific gap within the literature on personalisation and privacy. The finding that stands out in this study is the difference in comfort levels between personalisation based on behavioural history (self-disclosed to the system) and personalisation based on data inferred or covertly collected from the user.

3. Methodology

3.1 Research Design

In this research, a quantitative cross-sectional design with a within-subjects methodology has been employed. Each respondent rated their comfort level in five personalisation situations, enabling comparison of responses across several personalisation methods used for the same subject. The chosen design is appropriate for the study, as the aim of this research is to discover the relative threshold of consumers' comfort about personalisation rather than to assess causal effects on the whole population.

3.2 Sample Characteristics

Table 1: Age Distribution of Respondents

Age group	n	%
Under 18	1	2.0%
18–24	32	62.7%
25–34	11	21.6%
35–44	4	7.8%
45 and above	3	5.9%

Figure 1: Age Distribution of the Sample (N=51)




Figure 1 visualises the strong concentration of respondents in the 18–24 age group. This composition is important when interpreting the results: the findings are best viewed as a study of the surveyed online-shopping sample rather than as a representative estimate of all e-commerce users.

Table 2: Gender Distribution of Respondents

Gender	n	%
Female	29	56.9%
Male	21	41.2%
Prefer not to say	1	2.0%

3.3 Sample and Procedure

Data were obtained using an anonymous, self-administered online questionnaire between 7 July and 16 July 2026 with convenience sampling. 51 responses were deemed valid. The sample consisted of young Indian online shoppers (Table 1). Of the fifty-one responses, twenty-nine (56.9%) selected female, twenty-one (41.2%) male and one either male or female did not respond. Of the age group, 32 (62.7%) were aged 18–24, 11 (21.6%) were 25–34; 4 (7.8%) were 35–44, 3 (5.9%) were 45 and above, and 1 (2.0%) was under 18.

The survey shows that the respondents are frequent users of Myntra, Amazon, social commerce platforms, Flipkart, and Nykaa. Because the sample is gathered through convenience sampling and is thus demographically biased, the results do not reflect the entire population but serve as evidence for the studied sample. The data also include an under-18 respondent; thus, the study offers only descriptive results and does not claim the possibility of generalising the findings.

3.4 Measures

Table 3: Measures and Operationalisation

Construct	Operationalisation
Dynamic-pricing awareness	Yes / No / Not sure

Construct	Operationalisation
History-based recommendation comfort	5-point scale; 1 = very uncomfortable, 5 = very comfortable
Belief in tracking without purchase	Yes / No / Not sure
Perceived effect of personalisation	Much worse → much better
Scenario 1	Generic streaming recommendation based on prior activity
Scenario 2	Ad based on an overheard conversation
Scenario 3	Different price from a friend for the same item
Scenario 4	Store notification triggered by physical proximity
Scenario 5	Name plus personal-detail targeting
Algorithmic price fairness	More fair / No difference / Less fair
Explainability	Yes / Maybe / No
Trust	1 = no trust, 5 = high trust
Open-ended experience	Description of creepy, unfair, or invasive personalisation
Media literacy	Prior exposure to privacy-related documentaries/articles and self-reported effect

The five scenario items comprised the main repeated-measures outcome. These were worded rather concretely (not abstractly) to enable respondents to assess the subjective source of personalisation and what they may have imagined it related to. Additionally, across the 5 scenarios, they differed in communication media, channels, places, and prices. This allows the analysis to assess whether the predicted reaction pattern was evident among different forms of personalisation.

3.5 Ethical and Data-Handling Considerations

The survey was anonymous and self-administered. Respondents were not asked to provide any personal data for this study. The data contains only one respondent who is less than 18 years of age. Because of the anonymous nature of the survey and the fact that the information reported here is in the aggregate, the under-age response is retained solely for its illustrative use and not for purposes of making inferential claims regarding age. The presence of the under-age respondent is explicitly disclosed to clarify the structure of the sample and the parameters of the ethical/data-handling discussion. Only brief samples of the open-ended comments are presented.

3.6 Hypotheses

H1: The level of consumer comfort varies greatly in different scenarios of personalisation.

H2: The greatest drop in comfort level can be found between personalisation based on the data provided by the consumer and personalisation based on information that is not provided readily by the consumer.

H3: Awareness of dynamic pricing practices lowers comfort level for personalisation options that involve intense personalisation.

3.7 Data Analysis

In accordance with the ordinal nature of the five scenario scores and the repeated nature of the data collected from the respondents, the Friedman test has been chosen as the appropriate test for H1. This nonparametric method is applicable for comparisons on more than two connected samples without the parametric condition of normality in the population. Where nonparametric tests yield statistically significant results, Wilcoxon signed-rank tests have been employed for all 10 pairwise comparisons. The Bonferroni correction has been applied, resulting in an adjusted significance level of $\alpha = 0.005$.

Spearman's rank correlation was employed for an exploratory analysis of relationships between ordinal variables such as trust and high-intensity comfort. The Mann–Whitney U test was used to compare the groups, including participants who were aware of dynamic pricing and those who were not. Unless otherwise noted, all statistical tests were two-tailed at a significance level of $\alpha = 0.05$.

4. Results

4.1 Descriptive Findings

Comfort with history-based recommendations was moderate ($M = 3.39/5$), while trust in the platforms' responsible use of consumer data was lower than moderate ($M = 2.43/5$). The result is in line with the personalisation–privacy conflict: users can recognise the beneficial aspects of personalisation and use it without showing much trust in the firm's data-usage policies.

Participants showed a relatively high level of knowledge about dynamic pricing techniques, with 74.5% of respondents indicating that e-commerce websites apply dynamic pricing to different users for the same item. At the same time, about 58.8% think that dynamic prices are unfair compared to prices determined by people. Hence, knowledge does not lead to acceptance.

It was evident that personalisation had a favourable effect as well. The results indicated that 76.5% of respondents felt that personalised offers boosted their shopping experience. Indeed, 60.8% said that it was more enjoyable, while 15.7% reported that their shopping experience became much more pleasant thanks to personalisation. This information is relevant to the privacy paradox discussed further in the paper, given that consumers realise the considerable value of personalisation but believe that companies do not really protect their information.

Furthermore, the level of trust differed by age group. According to the descriptive cross-tabulation, the respondents aged 25–34 reported the lowest level of trust regarding data management by the companies ($M = 1.91/5$), whereas the mean trust was $M = 2.66$ for people aged 18–24 years, $M = 2.00$ among respondents aged 35–44 years, and $M = 2.33$ for respondents aged 45 and older.

Transparency was considered satisfactory, as 74.5% of respondents said "Yes" or "Maybe" when asked, "Why am I seeing this?" Thus, even if transparency is perceived as important, exploratory analysis shows it does not have a significant impact on trust.

4.2 Hypothesis Testing: The Personalisation Tipping Point

4.2.1 H1 — Omnibus Test

The Friedman test identified a statistically significant difference in comfort across the five scenarios, $\chi^2(4, N = 51) = 35.57$, $p < 0.001$, Kendall's $W = 0.17$. H1 is therefore supported. The results show that respondents did not treat all personalisation mechanisms as equally acceptable.

Table 4: Mean Comfort Ratings Across Personalisation Scenarios

Scenario	M	SD
Streaming recommendation (generic, history-based)	3.59	1.17
Overheard-conversation ad	2.39	1.20
Geofenced store notification	2.39	1.52
Price discrepancy vs. a friend	2.27	1.47
Name + personal-detail targeting	1.94	1.42

Figure 2: Comfort Across Personalisation Scenarios




Figure 2 shows the overall pattern of responses from the study participants. The participants felt significantly more comfortable with the first scenario than with the other four, which clustered at substantially lower comfort levels. The error bars indicate the standard deviations of the responses, illustrating the heterogeneity of the individual responses to the scenarios.

4.2.2 H2 — Locating the Threshold

Applying post hoc Wilcoxon signed-rank tests with a Bonferroni correction at $\alpha = 0.005$ indicates that the recommendation based on historical data is rated as more comfortable than any of the four other scenarios examined (all $p < 0.001$). However, no pairs of the other four scenarios achieved statistical significance in the comparisons.

This is the key finding of the research. If discomfort had been explained solely by personalisation levels, a continuous decline in comfort levels would have been expected. However, the data show that the first scenario is quite comfortable, while the remaining four scenarios are grouped into low-comfort categories. Thus, H2 is supported with modification: the results indicate a sharp comfort decline between the history-based scenario and the four other scenarios, rather than a gradual decline across increasing levels of personalisation. The findings therefore suggest that the source and context of the information used for personalisation may be more important than personalisation intensity alone.




Figure 3 summarises the “comfort cliff” by contrasting the generic history-based scenario with the average of the four higher-intensity scenarios. The latter average is approximately 2.25, compared with 3.59 for the generic scenario, a difference of roughly 1.34 points on the five-point scale.

4.2.3 H3 — Awareness and Comfort

The hypothesis was rejected in the study. The Mann-Whitney test results show no significant differences in comfort regarding instances of personalisation between individuals who have heard of dynamic pricing ($M = 2.28$) and those who have not ($M = 2.15$; $U = 186.5$, $p = 0.939$).

This indicates that mere knowledge of dynamic pricing does not account for the discomfort associated with extensive personalisation practices. This implies that consumers might be responding to actual instances rather than to heuristic knowledge of the practices associated with dynamic pricing.

4.3 Exploratory Findings

The age-by-trust cross-tabulation reveals additional information about the phenomenon under discussion. The 25-34-year-old cohort demonstrated the lowest average level of trust ($M = 1.91/5$). Although the process described does not support an age-group hypothesis test due to the small sample size, it highlights an important subgroup for future research on privacy issues.

Trust significantly correlated with the threshold response. Overall trust in the use of data by platforms was found to be positively correlated with comfort in difficult situations ($\rho = 0.57$, $p < 0.001$). Moreover, trust appeared to correlate negatively with the comfort difference between the neutral & difficult situations ($\rho = -0.41$, $p = 0.003$). This means that low-trust respondents showed a sharp decline when they crossed the data-source sphere.

Media literacy resulted in a similar null pattern as H3. Surveyed respondents who reported any exposure to media reports on privacy practices had an average high-intensity confidence level of 2.05, compared with 1.91 for those without such exposure. However, it is important to note that the previous result should be interpreted with caution due to the self-reported nature of the measurement and the small sample size.

The “explainability” variable also received no results. Trust averaged 2.35 for respondents who believed that explanations would raise their trust level, 2.50 for those who thought that explanations might raise their trust level, and 2.46 for those who did not believe that explanations would raise their trust level. Since this variable measures anticipated rather than actual effects, it should be emphasised that the study's results do not indicate that the “explainability” feature is ineffective. It just means that the endorsement of “explainability” does not correspond to trust-level indices in this particular sample.

4.4 Qualitative Results

Qualitative responses provide useful background information to interpret the quantitative threshold. A recurring point in many responses is tracking the consumer's behaviour across several platforms after they identify a product. One respondent reports noticing the dress on a shopping app, and several ads for this item later appear in other apps. The respondent feels uncomfortable since it appears that he is being followed everywhere. Another respondent states that after having a conversation about something with her friends, she starts seeing posts about this content on Instagram and Myntra.

These experiences are important because they indicate that relevance is not the only factor influencing comfort. In both examples, the consumer is aware of the data path beyond the immediate interaction. It was not just a case of incorrect recommendations; the respondent could not connect what they had been recommended to the information they believed they had provided to the platform.

5. Discussion

5.1 Interpreting the Personalisation Threshold Pattern

The primary conclusion from this research is that consumer comfort with personalisation does not decline gradually as personalisation intensifies. Instead, results indicate a clear threshold, or “comfort cliff.” Consumers were quite comfortable with general recommendations based on shopping history ($M=3.59$); however, comfort drops significantly when personalisation becomes more precise, especially when the message contains inferred personal information, location, private conversations, or other sensitive contextual information. For the four scenarios with high levels of personalisation, the mean comfort scores range from 1.94 to 2.39. Finally, the Friedman test yielded statistically significant results ($p < 0.001$), indicating that participants did not assign the same evaluation to the personalisation scenarios.

This pattern needs to be emphasised because it challenges the traditional assumption that increased personalisation leads to greater consumer value. While personalisation can enhance the relevance level of a certain service, it can also increase the feeling of surveillance or intrusion. The outcomes indicate that consumers can distinguish between useful personalisation and overly personalised services. Products that consumers have looked at or bought before can be perceived as normal and helpful examples of online services. In contrast, recommendations that rely on information consumers did not consciously provide may provoke discomfort.

Consequently, the findings lend credence to the notion that the relationship between personalisation and consumer acceptance is not linear. Instead of proceeding linearly from an acceptance standpoint, consumers can tolerate personalisation to a certain extent, then experience a sharp drop in comfort as soon as the perceived limit is breached. This finding is consistent with studies suggesting that personalisation becomes counterproductive when consumers perceive it as intrusive or when the information used is incommensurate with the benefit provided [1, 6].

It should be noted that the outcomes do not imply that respondents completely oppose personalisation. In fact, 76.5% of respondents reported that the personalised recommendations were helpful to their shopping

experience (60.8% said it was “somewhat improved” and 15.7% said it was “greatly improved”). However, mean trust toward the responsible use of personal information was only 2.43 out of 5. The combination of these two findings emphasises the personalisation-privacy paradox: consumers may appreciate the benefits of personalisation while distrusting the data practices that underlie it.

Also, it must be noted that consumers' awareness does not equal their trust. 88.2% of respondents believe that e-commerce services monitor their browsing activities even when they make no purchases. This means that people know their information is being collected; still, their average trust score is rather low, suggesting that awareness may lead to doubts rather than positive feelings about digital tracking.

To sum up, findings indicate that absolute personalisation does not mean the most successful personalisation. Personalisation is only possible when consumers understand it, when it provides them with something useful, and when it enables them to control the information used.

5.2 Theoretical Implications

The study contributes to the literature on personalisation and privacy by providing empirical evidence that the limit of acceptance for personalisation depends on the specific context. Prior investigations show that consumers may be affected by the use of personalisation in marketing both positively and negatively. Thus, personalisation helps enhance the relevance of communication and the efficiency of advertising, while excessive personalisation leads to intrusiveness and violations of privacy perceptions [1, 13].

This research goes beyond this argument by regarding personalisation as a continuum rather than a dichotomy. Rather than merely asking whether people wanted personalised or non-personalised experiences, the survey presented them with situations in which the level of personalisation varied. The results show that various scenarios contribute significantly to differences in data sensitivity and in the types of data used.

This is especially important in light of the personalisation–privacy paradox. Privacy calculus theory argues that people weigh the perceived benefits of disclosing personal information against the perceived risks of losing privacy when deciding whether to consent to the use of their personal data [4]. The current findings illustrate this balance in the personalisation experience. Consumers seem willing to be personalised when the benefits are apparent, but their willingness declines when the information appears more personal than expected.

The results are also in line with the findings of Kim and Han [6] about personalisation backfire. Their research shows that personalisation can elicit negative reactions from consumers when privacy concerns are great. The present research represents the related empirical data: respondents feel much less comfortable when personalisation goes beyond usual behavioural information.

It thus follows that the pertinent managerial query is not simply how much personalisation a platform can offer; rather, it should be at what point extra personalisation ceases to be advantageous and becomes annoying to the user. The transition from the intensity of personalisation to its boundaries provides a good starting point for future studies.

5.3 Real-World Illustrations

The concept of the threshold discussed in this research can be illustrated with typical e-commerce activities. Amazon and other big online shops often use data on past purchases, web browsing behaviour, and product interactions to generate relevant recommendations for customers. It is simple for customers to understand how their behaviours influence the recommendations generated.

Another good example is streaming platform recommendations. Movie suggestions based on what else has been watched are logical for viewers, since they can easily perceive a link between their previous behaviour and the recommendations they receive.

When personalisation begins to draw on information sources unrelated to the immediate shopping context, things can change. For example, a consumer may feel uneasy if a marketing communication seems to relate to a private conversation, or reveals an intimate detail or an exact location. Even if the underlying technology legitimately has access to such information, a consumer may see it as an inappropriate use of that information.

This distinction aligns with the problem of “creepiness” or intrusiveness identified in studies of personalised advertising. The level of personalisation can, thus, exert different effects depending on how the information is viewed, how it was collected, and whether the consumer had any expectations regarding its use.

One should note that these examples serve to illustrate rather than present empirical findings from this research.

5.4 Managerial Recommendations

The findings provide several practical implications for e-commerce platforms, marketers, and designers of AI-driven recommendation systems.

1. Design personalisation around a “comfort zone”

It is important for platforms not to assume that maximum personalisation is their main goal. Recommendation engines should differentiate between personal information that consumers may find acceptable and information that may be considered too sensitive or private.

For instance, recent browsing habits, purchase history, or preferences are likely to be more acceptable for personalisation purposes than more delicate information and unexpected contextual signals.

2. Explain why a recommendation was shown

The discovery that 74.5% of surveyed individuals responded “Yes” or “Maybe” to the question of whether a “Why am I seeing this?” description would improve their trust provides clear evidence of the usefulness of explanations.

Basic explanations such as:

“Recommended after examining similar products”

help clarify the personalisation process without inundating consumers with excessive technical jargon.

Hence, it is possible to use explainability as a type of trust infrastructure. Rather than conceal the personalisation mechanism, the sites can show how the recommendation was produced based on the user's actions.

3. Provide meaningful user control

Control is crucial to accompany transparency. Consumers should be able to change recommendation preferences, decline specific types of personalisation, and understand the categories of data that influence recommendations.

This is particularly important since awareness itself does not seem to generate trust. The 88.2% response rate on tracking awareness indicates that consumers have some understanding of tracking; however, the overall trust score remained low.

4. Treat algorithmic pricing differently from recommendations

The findings show that 58.8% of respondents believe prices set by algorithms are unfair compared to prices set by humans, while 74.5% know that prices can vary from standardisation, depending on user or

market features.

This means that consumers tend to accept algorithmic personalisation for recommendations more readily than for pricing. In the case of a recommendation, the consumer feels they receive something additional. In contrast, with personalised pricing, the perception may arise that two customers pay different amounts for the same product.

As a result, platforms should be wary of combining personalisation with price discrimination. It might be necessary to provide a clear explanation of the pricing rules and implement concrete measures to prevent discrimination in order to maintain the impression of fairness.

5. Segment trust-building strategies carefully

There is an additional possible indication for management based on the age analysis of responding participants. The mean trust score corresponds to the 25-34 age bracket ($M = 1.91/5$), the lowest among the age brackets. However, it is more appropriate to use the obtained statistical data as a basis for new hypotheses rather than to claim that the age brackets correspond to levels of trust.

6. Optimise for perceived value, not data intensity

To sum up, it is essential to assess personalisation based on value-to-intrusion ratio. This means that if additional data only improves recommendation quality slightly but causes far too much intrusion, the additional personalisation would do more harm than good. Thus, the idea should be to achieve proper personalisation and not maximum personalisation.

6. Conclusion

6.1 Summary of Findings

The present study investigated how consumers reacted to varying levels of AI-driven personalisation in e-commerce, with a special focus on the moment when personalisation shifts from being perceived as beneficial to being considered bothersome.

Results reflect an undeniable clash between the usefulness of personalisation and consumers' trust in the processes behind it. On the one hand, 76.5% of the respondents indicated that the use of personalisation improved their buying experience. At the same time, the average level of trust in e-commerce platforms regarding the responsible use of personal data was only 2.43 out of 5. It indicates that consumers appreciate the benefits of personalisation without fully trusting the systems that make it work.

The scenario-based analysis offers the key revelation of the research as a whole. Personalisation based on common shopping behaviour leads to a high level of comfort, which decreases quickly as the personal sensitivity to the scenarios increases. The Friedman test, which was performed and proved significant ($p < 0.001$), suggests that the differentiation in terms of personalisation explored across different scenarios is valid from a statistical point of view.

This tendency explains the proposition of a threshold of personalisation comfort. It can be stated that customers can endure a certain level of personalisation before the effect is perceived as overly invasive. Therefore, it should be stressed that personalisation should not be treated as a constant that adds value to recipients.

The research revealed considerable understanding of data retrieval techniques. According to the survey, 88.2% of the respondents agree that e-marketing sites monitor web performance irrespective of the buying process. Nonetheless, this knowledge does not mean that people felt safe regarding data usage.

The findings regarding price personalisation contributed to the emergence of another significant difference. The survey has shown that 58.8% of respondents view prices set by computer programs as

unfair compared with human-made prices. This indicates that consumers are more willing to accept AI when it improves product search, as opposed to instances where high prices result from AI applications. Ultimately, the concept of transparency has been recognised as a potentially useful tool for restoring trust. According to a survey, 74.5% of respondents said they would be more convinced if presented with a feature that allows users to find out the reason behind a given advertisement. It implies that transparency and control over processes may enable businesses to reap the benefits of personalisation while avoiding violations of user privacy.

6.2 Limitations

The interpretation of the outcomes should take into account certain constraints of this research.

To begin with, this investigation used a rather limited sample of 51 participants, which was adequate for identifying noticeable trends in the data collected but does not allow generalising the findings to the entire e-commerce community.

Second, this study used convenience sampling to obtain data, which implies that the sample group does not provide a representative demographic or behavioural profile for the whole community. Therefore, the results obtained in this study must not be considered population-level findings.

Third, the present research relied on self-reports from the subjects, which may differ from actual consumer behaviour. In this case, people were asked to provide their opinions about their trust, convenience, awareness, and fairness.

It's important to keep in mind that, when comparing age groups, not all groups had an adequate number of respondents. For instance, stating that people aged 25 to 34 had the lowest average trust score ($M = 1.91/5$) would be more descriptive than evidence of an age effect.

Moreover, since the current research is largely concerned with people's perceptions of comfort and trust, rather than how much they actually purchase, future studies could investigate the extent to which the identified personalisation threshold affects click rate, purchase intention, platform abandonment, and information disclosure.

Finally, results of the study represent perceptions at a certain point in time and, therefore, are likely to change as people become more familiar with recommendations created by computers, dynamic pricing, conversational commerce, and other methods of programming personalisation.

These limitations don't suggest that the research findings are wrong but rather define the range within which they should be evaluated.

6.3 Future Research

The findings create several opportunities for future research.

Larger and More Representative Samples

Further research on the topic should be conducted to obtain the same results using larger, more diverse samples. The implementation of probability sampling may allow verification of whether the personalisation limit discovered in the current project applies to various groups of people and markets.

Experimental Testing of Personalisation Intensity

Future research could conduct a controlled experiment in which participants are exposed to the same products personalised to varying degrees. Such an experiment could examine changes in participants' trust, purchase intention, perceived intrusiveness, and willingness to share personal data. This approach would provide a stronger basis for determining whether the "comfort cliff" observed in the present study is a causal or non-causal phenomenon.

Generative AI and Conversational Commerce

The rapid growth of generative AI ushers in a new area of research. AI shopping assistants trained on generative AI can recall user preferences, interpret user intent, and use conversational techniques to communicate with consumers. This raises important questions of whether conversational personalisation will be perceived as more helpful or more intrusive than traditional recommendation systems.

This opens the door for future research on whether consumer privacy thresholds change when it comes to conversational AI personalisation as opposed to conventional recommendation systems.

Spatial and Immersive Commerce

New forms of spatial computing, augmented reality, and virtual reality could enhance the amount of information that retailers consider useful. Factors such as location, motion, and gaze direction, among others, may become data inputs in personalised commerce systems.

The concept raises an interesting question for future investigation: Does personalisation become more invasive when artificial intelligence can observe not only customers' clicks but also their actions within the digital environment?

Personalisation and Dynamic Pricing

Future research should explore in depth the differences between personalised recommendations and personalised pricing. The results of the current study show that consumers may be receptive to AI recommendations while remaining sensitive to algorithm-determined prices.

Experimental studies can assess whether the fairness of the pricing process is perceived to be greater when information about the pricing algorithm is known, as well as whether consumers respond differently when price variations stem from actual cost differences versus personal characteristics.

Explainability as a Trust Mechanism

The positive response to the “Why am I seeing this?” concept provides a strong basis for further research into explainable AI. Future experiments could compare different forms of explanation, ranging from simple statements such as “Based on your recent browsing” to more detailed explanations of the data and algorithmic factors involved.

Researchers could then determine whether more explanation necessarily leads to greater trust, or whether excessive technical information creates additional confusion.

Acknowledgement

The author sincerely thanks all the respondents who took the time to participate in the survey and share their honest experiences. The author also thanks family and friends for their support and for helping to circulate the questionnaire.

References

1. A. Bleier, M. Eisenbeiss, "Personalized Online Advertising Effectiveness: The Interplay of What, When, and Where", *Marketing Science*, 2015, 34 (5), 669–688.
2. S. C. Boerman, S. Kruikemeier, N. Bol, "When Is Personalized Advertising Crossing Personal Boundaries? How Personalisation Affects Consumer Responses", *Computers in Human Behaviour Reports*, 2021, 4, 100144.
3. F. De Keyser, G. Van Noort, S. Kruikemeier, "Going Too Far? How Consumers Respond to Personalised Advertising from Different Sources", *Journal of Electronic Commerce Research*, 2022, 23 (3), 138–159.

4. T. Dinev, P. Hart, "An Extended Privacy Calculus Model for E-Commerce Transactions", *Information Systems Research*, 2006, 17 (1), 61–80.
5. N. Duani, A. Barasch, V. G. Morwitz, "Demographic Pricing in the Digital Age: Assessing Fairness Perceptions in Algorithmic versus Human-Based Price Discrimination", *Journal of the Association for Consumer Research*, 2024, 9 (3), 257–268.
6. H. Kim, S. Han, "Triggering the Personalisation Backfire Effect: The Moderating Role of Situational Privacy Concern", *Behavioral Sciences*, 2025, 15 (10), 1323.
7. T. Kim, K. Barasz, L. K. John, "Why Am I Seeing This Ad? The Effect of Ad Transparency on Ad Effectiveness", *Journal of Consumer Research*, 2019, 45 (5), 906–932.
8. K. Mehmood, K. Verleye, A. De Keyser, B. Lariviere, "The Transformative Potential of AI-Enabled Personalisation Across Cultures", *Journal of Services Marketing*, 2024, 38 (6), 711–730.
9. S. Raza, M. Rahman, S. Kamawal, A. Toroghi, A. Raval, F. Navah, A. Kazemeini, "A Comprehensive Review of Recommender Systems: Transitioning from Theory to Practice", *arXiv preprint*, 2024, arXiv:2407.13699. <https://arxiv.org/abs/2407.13699>
10. H. J. Smith, S. J. Milberg, S. J. Burke, "Information Privacy: Measuring Individuals' Concerns About Organisational Practices", *MIS Quarterly*, 1996, 20 (2), 167–196.
11. C. E. Tucker, "Social Networks, Personalised Advertising, and Privacy Controls", *Journal of Marketing Research*, 2014, 51 (5), 546–562.
12. L. Xia, K. B. Monroe, J. L. Cox, "The Price Is Unfair! A Conceptual Framework of Price Fairness Perceptions", *Journal of Marketing*, 2004, 68 (4), 1–15.
13. H. Xu, X. R. Luo, J. M. Carroll, M. B. Rosson, "The Personalisation Privacy Paradox: An Exploratory Study of Decision Making for Location-Aware Marketing", *Decision Support Systems*, 2011, 51 (1), 42–52.