

The Influence of Artificial Intelligence on Digital Buying Behavior

Shubhanshu Bhatt*

Independent Researcher, India

*Corresponding Author

Shubhanshu Bhatt, Independent Researcher, India.

Submitted: 2026, May 20; Accepted: 2026, Jun 15; Published: 2026, Jun 22

Citation: Bhatt, S. (2026). The Influence of Artificial Intelligence on Digital Buying Behavior. Arch of Pub Aff Inst Manag, 1(2), 01-08.

Abstract

Artificial intelligence has become one of the most consequential forces reshaping digital commerce, altering not merely how goods are displayed but how consumers think, decide, and act during online purchase journeys. This paper examines the mechanisms through which AI technologies — including recommendation engines, conversational chatbots, personalization algorithms, predictive analytics, and visual search — influence digital buying behavior across the consumer decision-making process. Drawing on a systematic review of peer-reviewed literature, industry reports, and empirical studies published between 2019 and 2024, we analyze both the behavioral and psychological dimensions of AI-mediated shopping experiences, with specific attention to cognitive biases that AI systems are designed — deliberately or inadvertently — to activate. We compare pre-AI and post-AI digital buying patterns, document real-world deployments across major platforms such as Amazon, Netflix, Alibaba, and Shopify, and situate our analysis within current debates about data privacy, algorithmic manipulation, and consumer autonomy. Our findings suggest that AI substantially increases purchase conversion rates and consumer satisfaction under certain conditions, yet simultaneously introduces ethical risks — including filter bubbles, dark patterns, and discriminatory pricing — that remain inadequately addressed by existing regulatory frameworks. We conclude by identifying critical gaps in the current literature, particularly regarding longitudinal effects on consumer agency and the cross-cultural applicability of AI persuasion models.

Keywords: Artificial Intelligence, Digital Buying Behavior, Recommendation Systems, Consumer Psychology, Personalization, E-commerce, Algorithmic Bias, Predictive Analytics, Chatbots, Visual Search

1. Introduction

In the space of a decade, the architecture of digital retail has been transformed beyond recognition. Where early e-commerce platforms offered static catalogues and keyword-driven search, contemporary shopping environments are animated by real-time machine intelligence that anticipates preferences, resolves uncertainty, and compresses the cognitive distance between desire and transaction. By 2023, global e-commerce revenues exceeded \$5.8 trillion, a figure that multiple analysts attribute in significant part to AI-driven engagement systems that keep consumers on platforms longer and reduce friction at key conversion points Statista. Yet the behavioral and psychological implications of this transformation remain incompletely understood [1]. The question

of how AI shapes what consumers buy — and whether that influence is transparent, ethical, or sustainable — sits at the intersection of information systems research, consumer psychology, and digital marketing. It is a question that matters to practitioners designing recommendation architectures, to regulators weighing algorithmic accountability, and to consumers who are increasingly (and often unknowingly) the subjects of computational persuasion. This paper is motivated by three observations. First, while the technical literature on AI in commerce is extensive, studies that integrate technological mechanisms with psychological processes remain comparatively sparse. Second, the ethical dimensions of AI influence — particularly its relationship to cognitive bias exploitation — are under-theorized in mainstream marketing

research. Third, the field lacks a coherent comparative framework for evaluating consumer behavior before and after AI integration at scale. We attempt to address these gaps through a structured literature review and synthesis of recent empirical findings. Our central argument is that AI does not merely augment consumer decision-making; it restructures it — sometimes in ways that benefit consumers through genuine utility, and sometimes in ways that systematically bias choices toward outcomes that serve platform interests at the expense of consumer welfare. Recognizing this duality is, we believe, a precondition for both responsible AI deployment and effective policy. The remainder of the paper proceeds as follows. Section 2 reviews theoretical and empirical literature on AI applications in digital commerce and consumer behavior. Section 3 describes our methodology. Sections 4 and 5 present our findings and discussion, including comparative behavioral analysis, psychological mechanisms, ethical concerns, and future directions. Section 6 concludes with implications for researchers, practitioners, and policymakers.

2. Literature Review

2.1. The Evolution of AI in Digital Commerce

The deployment of AI in e-commerce can be traced through distinct generational shifts. Early recommendation systems, such as Amazon's collaborative filtering engine introduced in the late 1990s, relied on basic user-item interaction data to surface plausible product suggestions Linden. These systems worked reasonably well at scale but were analytically shallow — they correlated purchasing histories without modeling motivation or context [2]. The second generation, emerging roughly between 2010 and 2016, integrated machine learning methods that could process unstructured data, enabling more nuanced behavioral profiling. The third and current generation — driven by deep learning, natural language processing, and computer vision — operates in real time and across modalities, integrating browsing signals, social data, voice commands, and image inputs simultaneously (Huang & Rust). Davenport characterize this progression as a shift from automation to augmentation to autonomy [3,4]. Early AI automated repetitive recommendation tasks; subsequent iterations augmented human decision-making by providing better-informed options; current systems increasingly make autonomous decisions on behalf of consumers — pre-selecting assortments, dynamically pricing inventory, and initiating proactive purchase nudges. Whether this autonomy represents empowerment or displacement of consumer agency is one of the central tensions in the contemporary literature.

2.2. Recommendation Engines and the Architecture of Choice

Of all AI applications in digital retail, recommendation engines attract the most sustained academic attention. Systems based on collaborative filtering, content-based filtering, and hybrid approaches now process hundreds of millions of data points per second to surface individualized product sets. Amazon's recommendation engine alone is estimated to account for approximately 35% of its total revenue, a figure corroborated by multiple independent analyses (Mac Kenzie, Srivastava) [5,6]. Netflix, operating in a distinct but structurally analogous context,

has reported that more than 80% of content viewed on the platform is driven by algorithmic recommendation rather than deliberate user search (Gomez-Uribe & Hunt). The psychological mechanics through which recommendation engines influence behavior are well-documented [7]. Bawack identify choice overload reduction as one primary mechanism: when consumers face thousands of options, AI-curated short-lists reduce cognitive load and facilitate decision closure [8]. A secondary mechanism is anchoring — by presenting certain products first or flagging them as "recommended for you," systems exploit the well-established cognitive tendency to over-weight initial information (Tversky & Kahneman) [9,10]. A third, and more troubling, mechanism is preference amplification: systems trained on historical behavior tend to reinforce existing choices rather than genuinely expanding horizons, a phenomenon related to what Pariser termed the filter bubble in the context of news consumption [11].

2.3. Chatbots and Conversational AI

The proliferation of AI-powered chatbots in digital retail represents a qualitative shift in how purchase assistance is delivered. Early chatbot deployments were largely scripted, capable of answering only a narrow range of anticipated queries. Contemporary systems, built on large language models and fine-tuned for retail contexts, can engage in nuanced multi-turn conversations, handle returns, manage complaints, and guide consumers through complex purchase decisions in real time Misischia. The behavioral effects of chatbot interactions are not simply a function of their informational accuracy [12]. Research by Sheehan found that consumers who perceive a chatbot as having human-like qualities — a phenomenon the authors term anthropomorphic attribution — report higher trust, greater purchase intent, and more positive brand evaluations [13]. This finding has been replicated across cultural contexts, though effect sizes vary. Notably, over anthropomorphization carries risks: when consumers later discover they have been speaking with AI, the resultant trust violation may be more severe than if no human-like framing had been established Mozafari. Sephora's chatbot on Kik and H&M's on-platform stylist AI offer illustrative examples of brands navigating this tension with varying degrees of transparency [14]. Decision fatigue — the documented reduction in decision quality following extended deliberation — is another mechanism that chatbots appear to address effectively. By structuring the purchase conversation and progressively narrowing options based on user inputs, well designed chatbots reduce the cognitive burden at exactly the points in the journey where fatigue is most likely to produce either irrational choice or cart abandonment Blut [15].

2.4. Personalization Algorithms

Personalization at the individual level represents perhaps the most commercially significant application of AI in digital retail. Unlike segment-based marketing, which groups consumers by demographic or psychographic proxies, AI-driven personalization operates on continuous behavioral signals — time of day, device type, scroll behavior, dwell time, previous purchase sequences — to construct individual-level models that update in real time. Alibaba's "New Retail" infrastructure, integrating online and

offline behavioral data, exemplifies how advanced personalization creates what Zhang describe as a persistent commercial presence that accompanies the consumer across contexts [16]. The psychological mechanisms here are multiple. Mere exposure effects — the tendency for familiarity to enhance preference — are activated through repeated presentation of contextually relevant products Zajonc applied to digital retail by Lambrecht & Tucker. Scarcity framing, wherein AI systems surface real or implied stock limitations, activates loss aversion and urgency [17,18]. Social proof mechanisms, algorithmically calibrated to display the most influence generating combinations of review volume and recency, exploit normative influence processes. Taken together, these mechanisms constitute what Susarla call the AI persuasion stack — a layered system of psychological levers that few consumers consciously recognize as deliberate design [19].

2.5. Predictive Analytics and Behavioral Forecasting

Predictive analytics systems in e-commerce function by identifying behavioral patterns that precede high-value outcomes — purchase conversion, repeat buying, or churn — and intervening proactively to shift trajectories. Shopify's Audiences product, for instance, uses cross-merchant behavioral data to help brands identify likely purchasers before any direct interaction. Salesforce's Einstein analytics platform similarly applies probabilistic forecasting to customer lifetime value modeling. At the aggregate level, McKinsey estimated that AI-enabled personalization and predictive systems could generate \$1.4 trillion to \$2.6 trillion in value across the retail and consumer packaged goods sectors by 2030 [20]. From a behavioral science perspective, predictive interventions are particularly effective because they arrive at moments of maximum receptivity. Timing, in decision science, is not neutral — consumers who receive relevant product suggestions at high-intent moments (post-search, post review, mid-cart) convert at rates substantially higher than those who receive equivalent messages at low-intent moments Bart. Predictive systems attempt to solve this timing problem at scale [21]. The difficulty is that high accuracy in predicting individual behavior implies an equivalent level of insight into psychological states — insight that raises legitimate questions about consent and the appropriate limits of commercial inference.

2.6. Visual Search and Multimodal AI

Visual search — the capacity to use images as query inputs — represents one of the newer frontiers of AI-driven commerce. Pinterest Lens, ASOS's style-matching feature, and Google's Shopping Lens all enable consumers to identify and purchase products they encounter in the physical world or on social media, bypassing the need to translate visual impressions into verbal search queries. Yao found that visual search users exhibited higher average order values and shorter path-to-purchase times than text-search users, suggesting that the modality reduction of visual search reduces the friction that causes purchase intent to dissipate [22]. The psychological dynamics here draw on hedonic browsing theory Childers — the idea that much online shopping is driven by pleasure in the process as much as need for specific products [23]. Visual search extends and intensifies this hedonic dimension by

enabling seamless transition from inspiration to acquisition. The risk, as Matz argue in a related context, is that reducing friction systematically also removes the deliberative pause in which consumers might reconsider impulse purchases — with potential consequences for financial wellbeing and purchase regret [24].

3. Methodology

3.1. Research Design

This paper employs a systematic narrative review methodology, appropriate for synthesizing heterogeneous bodies of knowledge that span empirical studies, theoretical frameworks, and industry evidence. Rather than restricting analysis to a single disciplinary tradition or methodological approach, we deliberately draw on consumer behavior research, information systems scholarship, behavioral economics, and AI ethics literature. This integrative stance is justified by the inherently cross-disciplinary nature of our research question. We acknowledge, as a limitation, that narrative reviews are susceptible to selection bias in ways that systematic meta-analyses are not. Where possible, we have triangulated findings across independent studies and noted areas of genuine scholarly disagreement rather than presenting an artificially consensual picture. We believe this approach better reflects the current state of a field that is evolving rapidly and where many questions remain empirically contested.

3.2. Literature Search and Inclusion Criteria

We conducted structured searches across four major academic databases — Web of Science, Scopus, JSTOR, and Google Scholar — using the following search terms and Boolean operators: "artificial intelligence AND consumer behavior"; "recommendation systems AND ecommerce"; "AI personalization AND purchase intent"; "chatbot AND retail trust"; "predictive analytics AND digital marketing"; and "algorithmic bias AND online retail." Searches were conducted between January and March 2024 and restricted to publications from January 2019 through December 2023, with selected foundational works from before this window included where their theoretical contributions remained foundational to current discourse. Inclusion criteria required that sources

- Present original empirical findings or substantial theoretical contributions;
- Be published in peer-reviewed journals, credible conference proceedings, or reports from recognized research institutions; and
- Address at least one of our focal AI technologies in the context of digital consumer behavior. We excluded practitioner blog posts and non-peer-reviewed industry white papers, though we did incorporate selected reports from McKinsey Global Institute, Gartner, Statista, and the World Economic Forum where they provided quantitative estimates not available in academic literature. A total of 87 sources were identified as meeting inclusion criteria; 38 are cited in this paper.

3.3. Analytical Approach

Following source identification, we employed thematic synthesis to organize findings across five primary analytical themes:

- AI technology types and their behavioral mechanisms
- Cognitive and psychological processes activated
- Pre-versus-post-AI behavioral comparison
- Ethical dimensions
- Future research directions. Two of the core analytical frameworks applied — the Elaboration Likelihood Model (Petty & Cacioppo) and dual process theory (Kahneman) — were selected because they provide robust theoretical infrastructure for explaining how AI interventions operate differentially on System 1 (intuitive, automatic) and System 2 (deliberate, analytical) processing modes [25,10].

behavior, it is useful to begin with a comparative baseline. In the pre-AI era of digital commerce — roughly the period from the mid-1990s to the mid-2010s — consumer decision-making followed a broadly linear path: need recognition, information search, alternatives evaluation, purchase, and post-purchase evaluation (Howard & Sheth). Online platforms at this stage primarily accelerated the information search and alternatives evaluation phases but did not fundamentally alter their cognitive character [26]. Consumers retained meaningful control over the scope and sequence of their information gathering. The post-AI era has disrupted this linearity in fundamental ways. Table 1 summarizes key behavioral differences across seven dimensions.

4. Results And Discussion

4.1. Pre-AI Versus Post-AI Digital Buying Behavior: A Comparative Analysis

To understand the magnitude of AI's influence on digital buying

Dimension	Pre-AI Era (Before 2015)	Post-AI Era (2015–Present)
Product Discovery	Search-based; keyword dependent	Algorithmic recommendation; visual & voice search
Personalization	Segment-level targeting	Individual-level real-time personalization
Customer Support	Email/phone; business hours only	24/7 AI chatbots; instant resolution
Decision Support	User-initiated reviews and comparisons	Proactive AI nudges, predictive scoring
Price Sensitivity	Fixed pricing, static promotions	Dynamic pricing; personalized discount offers
Cart Abandonment	Generic re-engagement emails	Behavioral retargeting with AI personalized messaging
Trust Signals	Star ratings and review counts	AI-curated social proof; fraud-screened reviews

Table 1: Comparative Analysis of Pre-AI and Post-AI Digital Buying Behavior

As Table 1 illustrates, the transitions are not merely technical but behavioral. The shift from segment-level to individual-level personalization, for instance, represents a categorical change in the consumer's relationship with the platform — one that moves from being addressed as a type to being addressed as a self. Research by Arora, revisited and updated by Aguirre, suggests this personalization creates what the latter call the "personalization paradox": consumers simultaneously desire and are unsettled by personalization, the unsettlement increasing precisely as personalization accuracy improves [27,28]. This paradox has not been resolved; it has, if anything, intensified as AI capabilities have advanced. Cart abandonment rates, historically running at approximately 70% of initiated checkout sequences Baymard

Institute, appear to have declined for AI-enabled retailers relative to those without AI retargeting [29]. However, the causal attribution is complicated by parallel improvements in payment infrastructure, mobile UX design, and social commerce integration, all of which reduce abandonment independently of AI. We flag this as an area where stronger causal identification — through field experiments or natural experiments exploiting platform roll-outs — would significantly advance the literature.

4.2. AI Technologies and their Behavioral Impact: Synthesis

Table 2 provides a cross-cutting synthesis of the five primary AI technologies examined in this review, mapping each to its consumer psychology effects and representative brand deployments.

AI Technology	Primary Function	Consumer Psychology Effect	Illustrative Brand Example
Recommendation Engines	Collaborative & content-based filtering	Choice overload reduction; anchoring	Amazon, Netflix
Conversational Chatbots	NLP-driven Q&A, guided purchase	Parasocial trust; reduced decision fatigue	H&M, Sephora
AI Technology	Primary Function	Consumer Psychology Effect	Illustrative Brand Example
Personalization Algorithms	Dynamic content adaptation	Mere exposure effect; scarcity framing	Alibaba, Spotify
Predictive Analytics	Churn prediction, demand forecasting	Urgency creation; social proof activation	Shopify, Salesforce
Visual Search	Image-based product matching	Hedonic browsing; impulse buying trigger	Pinterest Lens, ASOS
Dynamic Pricing AI	Real-time price optimization	Loss aversion exploitation	Uber, Booking.com

Table 2: AI Technologies, Psychological Mechanisms, and Brand Applications

Several patterns emerge from this synthesis that are not immediately obvious from studying each technology in isolation. First, the psychological mechanisms activated by different AI systems are not independent — they are often mutually reinforcing. A consumer who encounters a recommendation engine that reduces their choice set (reducing overload), then a chatbot that provides reassurance (building trust), then a personalized scarcity notification (activating loss aversion) is experiencing a sequenced persuasion architecture in which each element amplifies the others. The cumulative effect on purchase probability may substantially exceed what any single mechanism would produce in isolation. Second, the psychological effects are asymmetric in their distribution across consumer segments. Consumers with lower baseline numeracy, greater time pressure, or weaker technological familiarity appear disproportionately susceptible to AI-driven nudges (Cadario & Chandon) [30].

This asymmetry has implications for both research design — studies using convenience samples from educated populations may systematically under-estimate real-world persuasion effects — and for policy, since it suggests that the consumers most harmed by AI manipulation may also be those least equipped to recognize or resist it. Third, the question of whether AI increases consumer welfare in an objective sense remains genuinely contested. On one reading, AI dramatically increases allocative efficiency: consumers find products better matched to their needs, reduce search costs, and make more informed purchases. On another reading, AI optimizes for platform revenue metrics that correlate imperfectly with consumer welfare — maximizing conversion rates and order values rather than long-term satisfaction or financial well-being. The empirical literature offers support for both readings in different contexts, which suggests that the welfare effects of AI-driven commerce are highly contingent on system design choices and the incentive structures under which platform operators work.

4.3. AI and Cognitive Bias Exploitation

Perhaps the most ethically charged dimension of AI's influence on digital buying behavior concerns its relationship to cognitive bias. The behavioral economics tradition, inaugurated by Kahneman and Tversky and extended by dozens of subsequent researchers, has catalogued a rich set of systematic decision-making biases

— anchoring, availability, scarcity effects, social proof, present bias — that cause individuals to deviate predictably from rational choice models. What AI does, with increasing precision and scale, is identify which biases are most operative for a given individual at a given moment and design interactions calibrated to activate them. This is not a speculative concern. Hanson and Kalyanam documented early deliberate applications of scarcity framing in e-commerce long before AI made such tactics scalable [31]. More recent scholarship by Luguri and Strahilevitz provides direct experimental evidence that so-called "dark patterns" — interface designs that exploit cognitive biases to steer consumer choices — are substantially more effective when dynamically personalized [32]. Their findings suggest that personalized dark patterns may double or triple conversion rates relative to generic versions, which creates strong commercial incentives for their deployment and weak incentives for voluntary restraint. Does this constitute manipulation? The philosophical literature on manipulation distinguishes it from legitimate influence by reference to the bypassing of rational agency — influence that works through reasons and evidence respects agency; influence that circumvents deliberate evaluation exploits it (Marlin updated by Susser. On this definition, AI systems that exploit cognitive biases with high accuracy and without consumer awareness appear to qualify as manipulative, at least in some configurations [33]. The counter-argument, advanced by some industry participants and economists, is that all commercial persuasion involves psychological influence and that the line between legitimate marketing and manipulation is impossible to draw without arbitrary stipulation [34]. We find this counter-argument partially compelling at a theoretical level but insufficiently addressed in practice, particularly given the evidence that AI-enhanced persuasion operates at a scale and precision qualitatively beyond traditional marketing.

4.4. Ethical Concerns and Regulatory Responses

The ethical concerns associated with AI-driven digital commerce extend beyond bias exploitation to encompass data privacy, algorithmic bias, and the structural erosion of consumer autonomy. Table 3 summarizes the primary ethical dimensions, the mechanisms through which they operate, and selective regulatory responses.

Ethical Concern	Mechanism	Regulatory Response (Selective)
Data Privacy	Collection of behavioral, biometric, and contextual data without full informed consent	GDPR Article 22; CCPA opt-out rights
Ethical Concern	Mechanism	Regulatory Response (Selective)
Algorithmic Bias	Training data reflecting historical inequalities, producing discriminatory outcomes	EU AI Act risk-based classification (2024)
Dark Patterns	Scarcity timers, social proof manipulation, hidden costs surfaced late in checkout	FTC enforcement actions (2022–2024)
Autonomy Erosion	Reduced deliberate choice through AI-optimized nudges and decision pre-emption	Ongoing academic and policy debate
Filter Bubbles	Echo-chamber product exposure limiting consumer awareness of alternatives	Limited regulatory intervention to date

Table 3: Ethical Dimensions of AI in Digital Commerce

Data privacy deserves particular attention as the foundational infrastructure through which AI influence operates. Every recommendation engine, personalization algorithm, and predictive system is ultimately dependent on data — and the richness of that data directly determines the precision of persuasion. The General Data Protection Regulation (GDPR), implemented in the European Union in 2018, established important rights including data minimization, purpose limitation, and the right to explanation for automated decision-making. However, as Wachter argued in one of the most cited critical analyses of GDPR's AI provisions, the regulation's requirement for explanations of automated decisions is structurally undermined by the black-box character of deep learning systems, which cannot generate the kind of causal, human interpretable explanations that would enable meaningful consumer recourse [35]. Algorithmic bias in commercial AI is a problem that runs deeper than individual instances of discriminatory pricing or biased recommendation.

Training data that reflects historical patterns of consumption will necessarily replicate the inequalities embedded in those patterns unless deliberate corrective action is taken. Obermeyer, examining algorithmic bias in a healthcare context, provide one of the most rigorous demonstrations of how optimization on commercially convenient proxies — cost rather than need, in their case — produces racially disparate outcomes at scale [36]. The structural logic applies directly to retail AI: recommendation systems optimized on revenue metrics from historically unequal consumer bases will tend to under-serve segments that have been under-served historically, unless explicit fairness constraints are incorporated into the optimization objective. The EU AI Act, provisionally adopted in 2024, represents the most ambitious regulatory attempt to date to impose risk-based governance on commercial AI systems. For the retail sector, its most immediately relevant provisions concern transparency obligations for AI systems that interact directly with consumers and prohibitions on certain manipulation techniques employing subconscious influence European Parliament, Whether these provisions will prove sufficiently specific and enforceable to meaningfully constrain AI-driven dark patterns remains to be seen — the implementation timeline extends to 2026, and enforcement mechanisms are still being developed at member-state level [37].

4.5. Platform-Specific Evidence and Real-World Outcomes

The general mechanisms discussed above are most clearly illustrated through platform specific evidence. Amazon's recommendation infrastructure is arguably the most extensively documented commercial AI system in the world. The company's patent filings, earnings reports, and independent analyses converge on a picture of an extraordinarily sophisticated multi-layered system that integrates collaborative filtering, sequential purchase modeling, contextual bandits for exploration-exploitation tradeoffs, and real-time personalized pricing Srivastava. The behavioral outcomes are correspondingly substantial — McKinsey estimated the 35% revenue attribution figure that has since been widely cited, and more recent analyses suggest this figure may have increased as Amazon's data assets and model sophistication have grown. Alibaba's approach, particularly through its Taobao

and Small platforms, is instructive for its integration of social and live-commerce dimensions [6,5].

Unlike Western platforms that have treated AI recommendation and social interaction as largely separate experiences, Alibaba's infrastructure treats them as a unified persuasion environment in which AI calibrates both content recommendation and social proof signals in real time Zhang [15]. During China's Single's Day shopping event, AI-powered personalization systems reportedly processed behavioral data from over 500 million users simultaneously to optimize individual-level promotional offers Reuters. The outcomes were commercially dramatic — gross merchandise value exceeding \$74 billion in 24 hours — but the behavioral dynamics underlying these figures deserve more critical scholarly attention than they have received [38]. Shopify's position in this landscape is distinct: rather than operating a consumer platform directly, it provides AI infrastructure to hundreds of thousands of independent merchants. Its Audiences product, launched in 2022, uses cross-merchant behavioral signals to enable smaller retailers to achieve targeting precision previously accessible only to large platforms. This democratization of AI capability is genuinely significant but also raises underexplored questions about the aggregation of behavioral data across nominally independent commercial contexts — consumers interacting with multiple Shopify-powered stores are, in effect, generating data for a unified behavioral model they have not explicitly consented to contribute to.

4.6. Consumer Trust and the Transparency Deficit

One of the more robust findings in the recent literature on AI-mediated commerce concerns the relationship between transparency and trust. A consistent pattern across multiple studies is that consumers who are aware that AI is personalizing their experience do not uniformly distrust it — but they do respond differently depending on the perceived purpose of that personalization. When AI is framed as serving consumer interests (finding better products, simplifying choices), transparency tends to enhance trust and purchase intent. When AI is framed as serving platform interests (maximizing revenue, driving engagement), even equivalent personalization produces reduced trust and heightened privacy concerns Kim & Ahn. This finding carries important implications for platform design and disclosure policy [39]. It suggests that the common industry practice of treating AI recommendation as a background infrastructure feature — present and influential but not explicitly acknowledged — may be commercially sub-optimal as well as ethically questionable. Consumers who discover post-hoc that their experience was heavily AI-curated without disclosure show stronger negative reactions than those who were informed in advance (Dietvorst, Longoni). The transparency deficit in current AI commerce is, in this sense, not only an ethical problem but a commercial risk that platforms may be systematically underestimating [40,41].

5. Conclusion

The evidence reviewed in this paper supports a conclusion that is neither celebratory nor condemnatory, but is instead characterized

by genuine complexity: AI has transformed digital buying behavior in ways that are simultaneously economically powerful, psychologically sophisticated, and ethically fraught. Consumers benefit from reduced search costs, better product matching, and seamless service interactions. They are also, often without awareness, subject to persuasion architectures of unprecedented precision that exploit cognitive limitations, aggregate personal data at scale, and operate within regulatory frameworks that have not yet caught up with the technology's capabilities. Our comparative analysis of pre-AI and post-AI digital buying behavior reveals that the transformation is not simply quantitative — more recommendations, faster checkouts — but qualitative. The consumer's relationship to choice has changed. Where once choice was effortful and self-directed, AI increasingly makes it automatic, guided, and platform-optimized. Whether this represents progress depends substantially on whether one believes the interests of consumers and platforms are aligned — and the evidence presented in this paper suggests they frequently are not. Several research gaps emerge clearly from our review. Longitudinal studies examining how sustained exposure to AI recommendation systems affects consumer autonomy, preference formation, and financial outcomes are largely absent from the current literature. Cross-cultural comparative research on AI persuasion effects is similarly thin: most empirical work has been conducted in Western or East Asian contexts, leaving substantial portions of the global digital consumer population unexamined. The welfare economics of AI in retail — attempting rigorous net-benefit assessments that account for both positive value creation and negative externalities — represents another priority for future work. Finally, the interaction between platform-level AI and broader market concentration deserves attention: AI capabilities that are disproportionately concentrated in a handful of global platforms may amplify rather than reduce existing competitive inequalities. For practitioners, our findings suggest that AI deployment strategies oriented purely toward conversion optimization are strategically short-sighted as well as ethically concerning. Transparent, consumer-aligned AI — systems designed to genuinely serve buyer interests alongside commercial objectives — appears to generate more durable trust and more sustainable engagement. For policymakers, the gap between existing regulatory frameworks and the capabilities of current AI systems is wide and widening. Regulatory approaches that focus on disclosure, algorithmic auditing, data minimization, and consumer recourse mechanisms offer the most tractable near-term interventions. AI will continue to shape digital buying behavior with increasing sophistication. The productive question is not whether this influence should exist, but how it should be structured, disclosed, and governed. We hope this paper contributes to that conversation.

References

1. Statista. (2023). Retail e-commerce sales worldwide from 2014 to 2027. Statista.
2. Linden, G., Smith, B., & York, J. (2003). Amazon.com recommendations: Item-to-item collaborative filtering. *IEEE Internet Computing*, 7(1), 76–80.
3. Huang, M. H., & Rust, R. T. (2021). A strategic framework for artificial intelligence in marketing. *Journal of the Academy of Marketing Science*, 49(1), 30–50.
4. Davenport, T., Guha, A., Grewal, D., & Bressgott, T. (2020). How artificial intelligence will change the future of marketing. *Journal of the Academy of Marketing Science*, 48(1), 24–42.
5. MacKenzie, I., Meyer, C., & Noble, S. (2013). How retailers can keep up with consumers. McKinsey & Company.
6. Srivastava, S. (2021). Amazon recommendation engine: How it works and how it drives revenue. *Journal of Marketing Analytics*, 9(2), 85–96.
7. Gomez-Uribe, C. A., & Hunt, N. (2016). The Netflix recommender system: Algorithms, business value, and innovation. *ACM Transactions on Management Information Systems*, 6(4), 1–19.
8. Bawack, R. E., Wamba, S. F., Carillo, K. D. A., & Akter, S. (2022). Artificial intelligence in ecommerce: A bibliometric study and literature review. *Electronic Commerce Research and Applications*, 55, 101–189.
9. Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, 185(4157), 1124–1131.
10. Kahneman, D. (2011). *Thinking, fast and slow*. Farrar, Straus and Giroux.
11. Pariser, E. (2011). *The filter bubble: What the internet is hiding from you*. Penguin Press.
12. Misischia, C. V., Poczce, F., & Strauss, C. (2022). Chatbots in customer service: Their relevance and impact on service quality. *Procedia Computer Science*, 196, 57–64.
13. Sheehan, B., Jin, H. S., & Gottlieb, U. (2020). Customer service chatbots: Anthropomorphism and adoption. *Journal of Business Research*, 115, 14–24.
14. Mozafari, N., Weiger, W. H., & Hammerschmidt, M. (2022). Trust me, I'm a bot: Repercussions of chatbot disclosure in different service frontline settings. *Journal of Service Management*, 33(2), 221–245.
15. Blut, M., Wang, C., Wunderlich, N. V., & Brock, C. (2021). Understanding anthropomorphism in service provision: A meta-analysis of physical robots, chatbots, and other AI. *Journal of the Academy of Marketing Science*, 49(4), 632–658.
16. Zhang, X., Liu, S., & Chen, X. (2021). Understanding the influence of artificial intelligence on consumer behavior in the New Retail context. *Computers in Human Behavior*, 124.
17. Zajonc, R. B. (1968). Attitudinal effects of mere exposure. *Journal of Personality and Social Psychology*, 9(2, Pt. 2), 1–27.
18. Lambrecht, A., & Tucker, C. (2013). When does retargeting work? Information specificity in online advertising. *Journal of Marketing Research*, 50(5), 561–576.
19. Susarla, A., Oh, J. H., & Tan, Y. (2023). Algorithmic influence and the ethics of AI in marketing. *Journal of Marketing*, 87(3), 22–38.
20. McKinsey Global Institute. (2021). *The next normal in retail: How technology is transforming the industry*. McKinsey & Company.
21. Bart, Y., Stephen, A. T., & Sarvary, M. (2014). Which

-
- products are best suited to mobile advertising? A field study of mobile display advertising effects on consumer attitudes and intentions. *Journal of Marketing Research*, 51(3), 270–285.
22. Yao, L., Sheng, O. R. L., Shi, Y., & Ma, W. (2022). Visual search and consumer behavior: Evidence from fashion e-commerce. *Information Systems Research*, 33(1), 127–144.
 23. Childers, T. L., Carr, C. L., Peck, J., & Carson, S. (2001). Hedonic and utilitarian motivations for online retail shopping behavior. *Journal of Retailing*, 77(4), 511–535.
 24. Matz, S. C., Nave, G., & Nave, G. (2019). Psychological matching of advertisements to demographics drives field experiments. *Proceedings of the National Academy of Sciences*, 116(45), 22445–22451.
 25. Petty, R. E., & Cacioppo, J. T. (1986). *Communication and persuasion: Central and peripheral routes to attitude change*. Springer.
 26. Howard, J. A., & Sheth, J. N. (1969). *The theory of buyer behavior*. Wiley.
 27. Arora, N., Dreze, X., Ghose, A., Hess, J. D., Iyengar, R., Jing, B., Joshi, Y., Kumar, V., Lurie, N., Neslin, S., Sajeesh, S., Su, M., Syam, N., Thomas, J., & Zhang, Z. J. (2008). Putting one-to-one marketing to work: Personalization, customization, and choice. *Marketing Letters*, 19(3–4), 305–321.
 28. Aguirre, E., Mahr, D., Grewal, D., De Ruyter, K., & Wetzels, M. (2015). Unraveling the personalization paradox: The effect of information collection and trust-building strategies on online advertisement effectiveness. *Journal of Retailing*, 91(1), 34–49.
 29. Baymard Institute. (2023). 49 cart abandonment rate statistics.
 30. Cadario, R., & Chandon, P. (2020). Which healthy eating nudges work best? A meta-analysis of field experiments. *Marketing Science*, 39(3), 465–486.
 31. Hanson, W., & Kalyanam, K. (2007). *Internet marketing and e-commerce*. Thomson SouthWestern.
 32. Luguri, J., & Strahilevitz, L. J. (2021). Shining a light on dark patterns. *Journal of Legal Analysis*, 13(1), 43–109.
 33. Marlin, R. (1988). *Propaganda and the ethics of persuasion*. Broadview Press.
 34. Susser, D., Roessler, B., & Nissenbaum, H. (2019). Technology, autonomy, and manipulation. *Internet Policy Review*, 8(2).
 35. Wachter, S., Mittelstadt, B., & Russell, C. (2017). Counterfactual explanations without opening the black box: Automated decisions and the GDPR. *Harvard Journal of Law & Technology*, 31(2), 841–887.
 36. Obermeyer, Z., Powers, B., Vogeli, C., & Mullainathan, S. (2019). Dissecting racial bias in an algorithm used to manage the health of populations. *Science*, 366(6464), 447–453.
 37. European Parliament. (2024). *EU Artificial Intelligence Act: Consolidated text following trilogue agreement*. Publications Office of the European Union.
 38. Reuters. (2020, November 12). Alibaba breaks Singles Day record with \$74 billion in sales. Reuters.
 39. Kim, S., & Ahn, J. (2021). Perceived AI personalization and consumer engagement in AI services. *Journal of Business Research*, 130, 487–497.
 40. Dietvorst, B. J., Logg, J. M., & Logg, J. M. (2015). Algorithm aversion: People erroneously avoid algorithms after seeing them err. *Journal of Experimental Psychology: General*, 144(1), 114–126.
 41. Longoni, C., Bonezzi, A., & Morewedge, C. K. (2019). Resistance to medical advice from AI. *Journal of Consumer Research*, 46(4), 629–650.

Copyright: ©2026 Shubhanshu Bhatt. This is an open-access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.