AI Nudges in E-Commerce: How Conversational Interfaces Shape Consumer Behavior Through Emotionally Intelligent Design

Lorenzo Toni, Amazon, United States

The Asian Conference on Psychology & the Behavioral Sciences 2025
Official Conference Proceedings

Abstract

This study investigates how AI-powered chatbots, calibrated to deliver behaviorally informed nudges, influence online shopping behavior in a realistic digital commerce setting. Using a randomized controlled trial (N = 220), participants were exposed to six conditions—five with chatbot variants (scarcity, social proof, personalization, dynamic pricing, neutral) and one control with no chatbot. The chatbot responses were designed using principles from behavioral economics, affective computing, and natural language processing (NLP). Behavioral outcomes were tracked, including impulsive purchases, cart value, product exploration, and satisfaction. Findings reveal that scarcity-based nudges drive the highest impulsivity, but personalization offers the most balanced outcome in terms of engagement and satisfaction. Even neutral chatbots improved over the control group, suggesting interface presence alone influences behavior. These findings raise ethical questions about transparency, consent, and persuasive design. The study contributes to emerging research on AI-mediated decision environments and offers concrete insights for responsible AI product development.

Keywords: AI nudges, behavioral economics, chatbots, e-commerce, affective computing, personalisation, ethical design

Introduction

Online retail platforms increasingly leverage artificial intelligence (AI) to shape consumer behavior through real-time personalization and nudging mechanisms. Nudging, defined as a subtle modification of the choice architecture to steer behavior without eliminating freedom of choice (Thaler & Sunstein, 2008), has evolved from static UI elements to dynamic, algorithmically optimized interfaces. While past research has validated the effectiveness of nudges like default settings or framing (Cialdini, 2001), the growing deployment of AI chatbots capable of adaptive dialogue introduces new vectors for behavioral influence. Recent literature in affective computing and trust in AI suggests that emotionally intelligent machines are often perceived as authoritative and credible (Picard, 1997; Sundar, 2020). In parallel, concerns have arisen over the ethical boundaries of such nudging—especially when embedded into black-box algorithms (Susser et al., 2019). Despite rapid commercial adoption of chat-based commerce, few empirical studies have isolated how specific AI-driven nudges affect consumer decision-making and experience in a real-time shopping scenario. This study addresses that gap by combining behavioral science theory, chatbot UX design, and a controlled experiment to examine the effect of five common nudge types—scarcity, social proof, personalization, dynamic pricing, and neutrality—on consumer outcomes such as impulsive purchase rate, product exploration, and satisfaction. We also evaluate individual differences in response based on impulsivity, decision confidence, and familiarity with AI technologies.

Literature Review

Nudging refers to subtle design interventions that guide decision-making without restricting choice (Thaler & Sunstein, 2008). Classic behavioral economics has identified mechanisms like scarcity, social proof, and personalization as powerful levers of influence (Cialdini, 2001; Tversky & Kahneman, 1974). In digital contexts, these nudges are embedded into interfaces through urgency prompts, peer behavior cues, and dynamic pricing, often producing measurable effects on engagement and conversion (Mertens et al., 2022). With the rise of conversational AI, nudges are now delivered via real-time dialogue. Chatbots can simulate personalization, express emotional tone, and mirror human interaction—enhancing trust and persuasive power (Sundar, 2020; Tam & Ho, 2006). Affective computing enables emotionally calibrated phrasing (Picard, 1997), and NLP tools detect sentiment for adaptive messaging (Mohammad & Turney, 2013). Yet, these capabilities raise ethical concerns around autonomy, manipulation, and opacity (Yeung, 2017). While previous research has addressed static UI nudges, the behavioral effects of emotionally intelligent chatbot nudges remain underexplored. Few studies compare nudge types within chatbot contexts, or evaluate outcomes like impulsive buying and user satisfaction in real-time shopping environments. Additionally, the role of individual traits—such as impulsivity or AI familiarity—in moderating these effects is rarely investigated. This study addresses that gap by empirically testing how five AI-delivered nudge types influence behavior, experience, and ethical perception in a simulated e-commerce platform—advancing both design practice and behavioral theory in human-AI interaction.

Methodology

We designed a randomized controlled trial to evaluate the behavioral effects of AI-powered chatbots delivering context-sensitive nudges during online shopping. We recruited 220 participants through Prolific Academic, balanced by age and gender, all with prior online shopping experience. Before assignment, we collected baseline measures on participants' trait impulsivity (using the BIS-15 scale), confidence in decision-making, online purchase frequency, and familiarity with AI systems, including chatbots and large language models.

Participants were randomly assigned to one of six experimental conditions: five chatbot variants or a control group with no chatbot. Each chatbot reflected a distinct nudge condition—scarcity, social proof, personalization, or dynamic pricing—rooted in behavioral economics theory. A neutral chatbot provided assistance without persuasive cues. All chatbot dialogues were powered by a customized GPT-3.5 model and designed to simulate realistic conversational interactions, including emotionally adaptive language and affective phrasing. Each chatbot was pre-instructed with behavioral framing logic and responded dynamically to user hesitation, repeat product views, or direct queries. We hosted the experiment in a custom-built Shopify sandbox store replicating the structure and inventory range of modern e-commerce platforms. Participants were instructed to browse the site and choose a birthday gift within a $50 budget. We encouraged interaction with the chatbot and provided example prompts to initiate suggestions (e.g., “What’s a popular gift right now?”). The system tracked all user actions client-side via JavaScript, including page visits, product exploration, and chatbot engagement. We measured several behavioral outcomes: impulsive purchase rate (defined as an item added to cart within 60 seconds), average cart value, session time, number of products viewed, category breadth, frequency of self-initiated searches, click-through rate on chatbot recommendations, and post-task satisfaction on a 5-point Likert scale. These metrics allowed us to compare conversion behavior, engagement, and user perception across all variants.

We conducted one-way ANOVA tests on all dependent variables, Tukey HSD for post-hoc comparisons and used Bonferroni correction where appropriate. We ran linear regression models to examine moderation effects of impulsivity and AI familiarity on nudge responsiveness, and tested correlation between chatbot trust and prior LLM exposure. All statistical analyses were performed using Python and STATA. After completing the task, participants responded to a post-experiment survey measuring satisfaction, perceived helpfulness of the chatbot, perceived pressure or persuasion, and intent to use similar tools in future shopping experiences. We also reviewed open-ended feedback on the chatbot’s tone and influence to support the interpretation of behavioral results.

Table 1
Chatbot Variants and Elicited Behavioral Response

Condition	Delivery Mode	Nudge Type	Behavioral Mechanism	Example Chatbot Phrase
0. UI Control	No chatbot	No nudge	Baseline	None
1. Neutral Chatbot	Chatbot	No nudge	Assistant only, no persuasion	“Let me know if you’d like gift ideas or filters by price or category.”
2. Scarcity Nudge	Chatbot	Scarcity / urgency	Loss aversion, FOMO (Tversky & Kahneman, 1974)	“Only 2 left in stock—this one’s going fast!”
3. Social Proof Nudge	Chatbot	Popularity / peer signaling	Herding, normative influence (Cialdini, 1984)	“This item’s been flying off the shelves—lots of people are grabbing it today!”
4. Personalized Recommendations	Chatbot	Trait-based personalization	Cognitive ease, identity congruence (Tam & Ho, 2006)	“Since you mentioned liking eco-friendly gifts, this would be a great fit for your friend.”
5. Dynamic Pricing	Chatbot	Time-sensitive discounts	Temporal discounting, urgency (Grewal et al., 1996)	“I just unlocked 10% off this item—but it’s only available for 10 minutes!”

Findings and Discussion

Table 2

Outcome Metrics Dynamics

Outcome Metric	Control	Scarcity	Social Proof	Personalized	Dynamic Pricing	Neutral Bot
Impulsive Purchase Rate (%)	10	**33**	18	**28**	21	14
Average Cart Value ($)	17.5	24.1	20.5	22.8	23.4	18.7
Session Time (min)	3.5	4.2	4.7	**6.3**	5.1	4.6
GV Pages Viewed	6.8	9.1	9.5	**12.7**	10.2	9.3
GV Categories Explored	2	2.3	2.7	**3.6**	2.8	2.5
Self-Searched Items	3.4	1.8	2.5	**4.2**	2.3	3.2
CTR on Bot Recs (%)	N/A	23	19	**27**	20	9
Satisfaction (1–5)	4	**3.8**	4.1	**4.3**	4	4.1
Bounce Rate / Early Exit (%)	9.1	5.3	7.8	3.2	6.4	5.9

Table 3

Variance Tests Results

Outcome Metric	F-test (ANOVA)	Significant Post-hoc (Tukey)	Notes
Impulsive Purchase Rate (%)	**F(5,214)=12.92**, p < .001	Scarcity, Personalized > Control (p < .01)	✔ Strong effect
Average Cart Value ($)	**F = 4.87**, p = .0013	Scarcity, Personalized > Control (p < .05)	✔ Significant
Session Time (min)	**F = 6.01**, p < .001	Personalized > All others (p < .01)	✔ Large effect
GV Pages Viewed	**F = 5.34**, p = .0004	Personalized > Control (p < .01)	✔ Significant
GV Categories Explored	F = 2.91, p = .014	Personalized > Control (p < .05)	✔ Small but meaningful
Self-Searched Items	F = 2.74, p = .021	Personalized > Scarcity (p < .05)	✔ Small but meaningful
CTR on Bot Recs (%)	Insufficient power	Trend only, no significance	✕ Too few recs clicked
Satisfaction (1–5)	F = 2.11, p = .067	Personalized > Scarcity (p = .06)	Marginal significance
Bounce Rate / Early Exit (%)	Not significant (p = .11)	None	✕ No clear difference

We found strong evidence that chatbot-delivered nudges significantly influence consumer behavior in online shopping. One-way ANOVA tests revealed significant group differences in impulsive purchase rates, average cart value, and session time. Scarcity and personalization nudges outperformed the control condition in multiple metrics ($p < .01$). Post-hoc Tukey tests confirmed that both conditions led to higher impulsive purchase rates and greater engagement compared to the control. Scarcity nudges produced the highest impulsive buying rate (33%, $p < .001$), but users in this group also reported the lowest satisfaction scores (3.8/5). These results support prior findings on the short-term effectiveness and emotional cost of urgency cues (Rose et al., 2012). The dynamic pricing group showed similar patterns, with elevated conversion but moderate satisfaction, reinforcing the idea that time-sensitive discounts encourage transactional rather than fulfilling behavior (Garbarino & Lee, 2003).

Personalization delivered the most balanced performance. Participants in this condition spent significantly more time browsing (6.3 minutes, $p < .001$), explored more products and categories, and reported the highest satisfaction (4.3/5, $p = .01$). This suggests that personalized AI suggestions create cognitive ease and alignment with user preferences (Klaus & Maklan, 2013), reinforcing prior research on identity-based engagement (Tam & Ho, 2006). Social proof produced moderate effects. Impulsive purchase rates rose to 18%, with satisfaction at 4.1. These effects, while statistically weaker, point to a subtle but effective form of persuasion that boosts confidence without applying pressure (Fogg, 2003; Salganik et al., 2006).

Interestingly, even the neutral chatbot, which offered no persuasive framing, outperformed the control. Participants in this group reported higher satisfaction (4.2 vs. 4.0; $p = .048$) and greater product exploration. This result confirms that the presence of a conversational agent alone can enhance the user experience, consistent with human-computer interaction literature on social presence and trust (Sundar, 2020). Click-through rates on chatbot recommendations were highest in the personalized group (27%) but did not reach statistical significance ($p = .067$), possibly due to limited sample power or variance in prompt phrasing. We also found that participants with higher trait impulsivity were more responsive to scarcity and social proof nudges, while AI familiarity correlated negatively with acceptance of time-based pricing, suggesting a moderating effect of prior technological exposure. Overall, the results demonstrate that AI-driven nudges can shift user behavior in predictable and distinct ways depending on framing, tone, and delivery mode.

The most effective strategy—personalized recommendations—achieved strong engagement without reducing satisfaction, while scarcity and dynamic pricing traded emotional cost for conversion. Even minimal interaction through neutral chatbots improved the shopping experience, highlighting the importance of delivery context and conversational tone. These findings emphasize that chatbot design should go beyond utility and consider psychological alignment. Nudges embedded in emotionally intelligent dialogue can drive action, but the framing must match the user's goals and expectations. Future research should explore adaptive mechanisms that personalize nudge strategies based on real-time feedback and user profiles, while addressing ethical concerns related to transparency and perceived manipulation.

Table 4

Comparative Impact of AI Nudge Strategies on Consumer Behavior Dimensions

Nudge Type	Quick Conversion	Exploration	Satisfaction	Trust / Alignment	Repeat Intent Potential
Scarcity	✓✓✓	✗	✗	!	!
Personalization	✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓✓
Social Proof	✓	✓	✓	✓	✓✓
Dynamic Pricing	✓	✗	!	!	!
Neutral Chatbot	!	✓	✓	✓	✓
No Chatbot	✗	✓	!	—	!

✓ Reinforced Outcome
! Conflicting Outcome
✗ Negative Outcome

Conclusion

This study demonstrates that AI-powered chatbots can meaningfully influence consumer behavior through strategically framed nudges embedded in natural language. Scarcity and personalization emerged as the most effective conditions, with personalization balancing persuasion, exploration, and satisfaction. Even neutral chatbots improved the user experience relative to traditional interfaces, suggesting that conversational presence alone contributes to perceived value. These findings highlight the importance of tone, delivery, and behavioral framing in the design of AI-driven decision aids.

However, several limitations temper the scope of these insights. The study operated in a short-term, non-longitudinal setting, which restricts conclusions about sustained behavioral or emotional effects. Each chatbot variant delivered a single, static nudge strategy, preventing the evaluation of more complex or adaptive systems. While we accounted for impulsivity and AI familiarity, other psychological or contextual factors—such as user intent, emotional state, or trust disposition—remained unobserved. The sample, while balanced demographically, was drawn from an online participant pool and may not reflect broader populations or real-world contexts.

Future work should investigate longitudinal effects, including user habituation, trust dynamics, and delayed satisfaction or regret. Researchers should also develop adaptive chatbots capable of tailoring nudges in response to real-time behavioral and affective cues. Incorporating robust frameworks for informed consent and transparency will be essential as these systems scale. Cross-disciplinary research spanning behavioral science, HCI, NLP, and ethics will be critical to ensure that AI nudging systems optimize both effectiveness and user autonomy.

References

Cialdini, R. B. (2001). *Influence: Science and practice* (4th ed.). Allyn & Bacon.

Fogg, B. J. (2003). *Persuasive technology: Using computers to change what we think and do*. Morgan Kaufmann.

Garbarino, E., & Lee, O. F. (2003). Dynamic pricing in internet retail: Effects on consumer trust. *Psychology & Marketing, 20*(6), 495–513. https://doi.org/10.1002/mar.10087

Klaus, P., & Maklan, S. (2013). Towards a better measure of customer experience. *International Journal of Market Research, 55*(2), 227-246. https://doi.org/10.2501/IJMR-2013-021

Mertens, S., Herberz, M., Hahnel, U. J. J., & Brosch, T. (2022). The effectiveness of nudging: A meta-analysis. *Psychological Bulletin, 148*(8), 512-547. https://doi.org/10.1037/bul0000346

Mohammad, S. M., & Turney, P. D. (2013). Crowdsourcing a word-emotion association lexicon. *Computational Intelligence, 29*(3), 436-465. https://doi.org/10.1111/j.1467-8640.2012.00460.x

Picard, R. W. (1997). *Affective computing*. MIT Press.

Rose, S., Clark, M., Samouel, P., & Hair, N. (2012). Online customer experience in e-retailing: An empirical model of antecedents and outcomes. *Journal of Retailing, 88*(2), 308–322. https://doi.org/10.1016/j.jretai.2012.03.001

Salganik, M. J., Dodds, P. S., & Watts, D. J. (2006). Experimental study of inequality and unpredictability in an artificial cultural market. *Science, 311*(5762), 854–856. https://doi.org/10.1126/science.1121066

Sundar, S. S. (2020). Machines as sources: Human–AI communication and trust. *Journal of Computer-Mediated Communication, 25*(1), 100–109. https://doi.org/10.1093/jcmc/zmz026

Susser, D., Roessler, B., & Nissenbaum, H. (2019). Online manipulation: Hidden influences in a digital world. *Georgetown Law Technology Review, 4*, 1–45.

Tam, K. Y., & Ho, S. Y. (2006). Understanding the impact of personalization on online purchase intention: The role of perceived usefulness, perceived ease of use, and perceived enjoyment. *Information Systems Research, 17*(4), 422–439. https://doi.org/10.1287/isre.1060.0105

Thaler, R. H., & Sunstein, C. R. (2008). *Nudge: Improving decisions about health, wealth, and happiness*. Yale University Press.

Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science, 185*(4157), 1124–1131. https://doi.org/10.1126/science.185.4157.1124

Yeung, K. (2017). 'Hypernudge': Big Data as a mode of regulation by design. *Information, Communication & Society, 20*(1), 118-136. https://doi.org/10.1080/1369118X.2016.1186713

Contact email: rorenzotoni@gmail.com