

Persuasiveness of Conversational Agents for Targeted Advertising: Autism and Gen-AI Chatbots

KIRSTEN CHAPMAN, Brigham Young University, U.S.A

KAITLYN KLABACKA, Brigham Young University, U.S.A

GARRETT SMITH, Brigham Young University, U.S.A

XINRU PAGE, Brigham Young University, U.S.A

Marketing firms are starting to leverage Generative-AI-based chatbots as a more persuasive form of advertising. This is potentially more harmful for autistic young adults with substantial support needs who tend to have lower financial literacy, be more trusting of bad actors when communicating through social technologies, and more likely to anthropomorphize objects. Thus, we investigated whether a chatbot assuming the likeness of a favorite celebrity would more greatly influence consumer preferences of this population when compared against either autistic individuals without substantial support needs or the general population. We conducted an experimental survey where participants 1) ranked their preferences of various consumer items, 2) performed intervening tasks answering several questions, 3) viewed a chatbot interaction where the chatbot exhibited opposite consumer preferences, and then 4) re-ranked their preferences for those items. We found that autistic young adults with substantial support needs were more likely to change their preferences than the sample from the general population. Moreover, they were also more likely to change their preferences than autistic young adults without substantial support needs. These findings suggest that autistic individuals with substantial support needs are more susceptible to celebrity chatbot persuasion. We discuss the risks and guardrails that need to be associated with deploying generative-AI-based chatbots for targeted advertising.

CCS Concepts: • **Human-centered computing** → **Human computer interaction (HCI)**; **Collaborative and social computing theory, concepts and paradigms**; **Accessibility**.

Additional Key Words and Phrases: Autism, Generative-AI Chatbots, Targeted Advertising, Computers as Social Actors, Anthropomorphism

ACM Reference Format:

Kirsten Chapman, Kaitlyn Klabacka, Garrett Smith, and Xinru Page. 2025. Persuasiveness of Conversational Agents for Targeted Advertising: Autism and Gen-AI Chatbots. *Proc. ACM Hum.-Comput. Interact.* 9, 7, Article CSCW464 (November 2025), 23 pages. <https://doi.org/10.1145/3757645>

1 Introduction

Targeted advertising has become an integral tool in promoting economic development in the past decade because of its efficacy in promoting consumer behavior [14, 92]. In fact, targeted ads have been so effective that limits have been put in place when advertising to vulnerable populations (e.g., teenagers) on social media platforms [4, 5]. Targeted advertising could also unduly influence autistic young adults (YAs). For instance, prior work has found that autistic YAs are more likely to be more trusting of bad actors on social technologies due to taking a literal interpretation of technology features and text [72]. The type of persuasive language frequently used in targeted advertisements

Authors' Contact Information: Kirsten Chapman, kac273@byu.edu, Brigham Young University, Provo, U.S.A; Kaitlyn Klabacka, Brigham Young University, Provo, U.S.A; Garrett Smith, Brigham Young University, Provo, U.S.A; Xinru Page, xinru@cs.byu.edu, Brigham Young University, Provo, U.S.A.



This work is licensed under a Creative Commons Attribution 4.0 International License.

© 2025 Copyright held by the owner/author(s).

ACM 2573-0142/2025/11-ARTCSCW464

<https://doi.org/10.1145/3757645>

(e.g., "this product is for you") could also be harmful if followed literally as a direction. Furthermore, autistic adults have been found to have lower rates of financial literacy than neurotypical adults [40], potentially making them even more vulnerable to targeted advertising.

The introduction of generative-AI based chatbots provides an even more personalized way to target and persuade users. Marketing and advertising firms have already started to leverage such chatbots in advertising [2]. Research thus far shows that chatbots are effective in increasing trust and leading to impulse purchasing [50]. Furthermore, common gen-AI chatbots, such as MetaAI [8], use recognizable public figures as the face of their product which can lead to increased levels of anthropomorphism [32]. Anthropomorphizing can lead to wanting to please the inanimate object, or even lead to feelings of sadness and despair when the user perceives that the object is being hurt or lonely [99]. Prior research in an offline context also shows that autistic individuals are more likely to anthropomorphize objects than neurotypical individuals [25, 28, 99].

Given these past findings, autistic individuals may be more likely than neurotypical individuals to be unduly influenced and experience harm when exposed to chatbot advertising. Moreover, prior research has found that autistic young adults with substantial support needs are prone to overly trusting and being misled by those they connect to online [72]. It is important to understand whether using anthropomorphized chatbots in targeted advertising could similarly pose a threat to this population. Thus we explore whether autistic young adults are impacted by this emerging technology to a greater degree than neurotypical individuals. Furthermore, we examine whether autistic young adults with versus without substantial support needs are affected to different extents. Our research questions are as follows:

RQ 1: Are autistic YAs influenced by chatbots more than neurotypical YAs?

RQ 2: Are autistic YAs with substantial support needs influenced by chatbots more than autistic YAs without substantial support needs?

In order to answer these questions, we ran an experimental survey where participants 1) ranked various consumer products and brands based on personal preference, 2) answered several questions as an intervening task, 3) viewed a chatbot interaction which stated the exact opposite opinions about these products, and then 4) ranked their preferences again. We found that autistic YAs with substantial support needs were likely to change their preferences a greater amount than neurotypical individuals or autistic YAs with lower support needs.

Our findings have several implications. We show that autistic YAs with substantial support needs are at an even higher risk of having their consumer preferences changed when interacting with Gen-AI celebrity chatbots. System designers need to consider how to guard against the potential harms. Managing risks associated with targeted advertising may even involve additional support from others, in line with prior work that shows limiting risk on social media often requires engaging the larger support network of this population [31]. Additionally, our results underscore the need for autism research to involve a wider range of participants. Most studies focus on those without substantial support needs, who live independently, and even have attended college. However, there is a large population of autistic individuals with much greater support needs. Policy makers should take these results as a call to action to protect populations who are more vulnerable in context of gen-AI and digital marketing.

2 Background

In this section, we describe prior research related to our study. We first discuss targeted advertising that is used on social technologies, as well as the existing regulations keeping certain populations safe from such features. Then, we describe the current uses of generative-AI chatbots in marketing

and the ethical implications. Next, we discuss the unique experience of anthropomorphism for autistic individuals. Finally, we discuss the Computers As Social Actors paradigm.

2.1 Targeted Advertising on Social Technologies

Social Networking Sites such as Facebook rely on revenue generated through advertisements placed on their platform. These ad placement systems use complex algorithms to deliver ads to users based on a user's likelihood of interacting with the advertisement. To do this, Facebook leverages user information such as location, likes, comments, websites visited off platform, and other indicators of engagement to be able to infer specific interests of every user. Facebook then uses these inferred interests to deliver more relevant ads to their users. However, research has shown these forms of targeted advertising are perceived by users to be particularly concerning [49, 95]. This concern is not unwarranted as these systems have led to several instances of discrimination [29, 74, 88]. These concerns combined with broader concerns about surveillance have led to a push for stronger government regulation on targeted advertising [20, 29]. Recent product changes also reflect this sentiment; Meta recently chose to restrict targeted advertising for its teen Instagram users who are at higher risk of being manipulated [7].

Scholars have just begun to examine the intersection of social media advertising and autism. Much of this work has examined autistic individuals' feelings towards the content type. For instance, one study found that autistic participants felt misrepresented by the advertisements they received on social media [104]. They found that many of these advertisements had misleading portrayals of autism and sought to push products or content with harmful ideologies of trying to "cure" autism. Another study discussed the implications of targeted advertising using existing biased data [39]. Populations historically under diagnosed may be less likely to receive advertisements for assistive technologies which could benefit autistic individuals. However, there is little work that examines the decision-making impact of targeted advertising for autistic individuals. This is especially important to study with the advent of Gen-AI which is being leveraged for advertising [43, 47].

2.2 Gen-AI Chatbots, Marketing, and Ethical Considerations

A key determinant of the output of a generative-AI chatbot is the data it was trained on and how it was fine-tuned. Training data for generative-AI models provides a set of patterns and relationships which are used to create output [36]. For example, Gen-AI chatbots leverage patterns of human communication and are often fine-tuned to respond to users in a human-like way [21]. The effects of training data and fine-tuning can be seen in chatbot responses as they simulate a friend, recommend common products, and convey realistic looking emotions (see figure 1).

The use of generative-AI based chatbots in marketing has become more prevalent in recent years [52]. Previous work has found this technology to be effective in encouraging consumer purchases and increasing trust in a general context [27, 62, 94]. In fact, some studies have even shown that this technology leads to increased impulse purchasing [50]. One potential reason for this is that many of these chatbots are designed to intentionally mirror human identity cues (e.g., appearance, naming conventions)[27, 94]. Furthermore, these chatbots can provide human-like output 24/7 [79]. Despite being shown to be effective broadly, there is little work that has investigated whether certain populations are more susceptible than others to generative-AI chatbot marketing. Understanding whether some populations are more vulnerable is important to consider when defining the boundaries within which these technologies should be able to operate.

Prior work has outlined ethical concerns and guidelines regarding generative-AI chatbot technology [79]. The primary risks cited are humanizing chatbots, information asymmetry, and privacy

violations [64]. Recommendations from prior work include attention to potential bias, transparency, human oversight, continuous monitoring, and increased attention to privacy needs [10, 23, 35, 42, 59, 78, 79]. The White House has also created a blueprint for an AI bill of rights which states that AI designers need to incorporate 1) safe and effective system designs, 2) algorithmic discrimination protections, 3) increased data privacy, 4) notice and explanation, and 5) human alternatives, consideration, and fallback [6] into their AI-based designs. The call for notice and explanation cites that AI-designers must make it clear that an AI-system is being used. Without an increased understanding of differences between various populations, we may not understand the full scope of necessary ethical considerations.

Prior work has also examined the ethics of designing technologies for vulnerable populations. One important principle in this work is "nothing about us without us" [26, 101], meaning that the design process should be inclusive of those from marginalized population. Including autistic users in the design of AI technologies is important as there are limited training datasets consisting of data collected solely from neurodiverse populations. This limitation has been difficult for past researchers as they attempt to design AI technologies for this population. For instance, one study designed a tool, trained on data from the general population, that recognized facial expressions[18]. When this feature was tested on an autistic population, they found that participants either felt like the feature was inaccurate or doubted their own abilities to recognize emotions. Finally, prior work has concluded that as we design for marginalized communities, it is important to take into account various factors often excluded from research, such as the "political, historical, and communal roots"[100].

2.3 How People Perceive and React to Non-Human Actors

In the age of artificial intelligence (AI), research about anthropomorphism has increased [55]. Anthropomorphism is when human-like characteristics are attributed to nonhuman things [34]. The proliferation of AI has allowed for a greater level of human-like qualities to be conveyed through technology [55, 81]. The majority of studies examining AI and anthropomorphism have focused on voice assistants, social robots, and autonomous vehicles [55]. When AI technologies, such as chatbots, have positive anthropomorphic attributes (e.g., perceived warmth), it has been found to lead to increased user trust [27]. Furthermore, anthropomorphism can lead to an increase in persuasiveness [97]. Simultaneously, feelings of discomfort due to the "uncanny valley" may arise and reduce trust if the feature is designed in a way that is perceived as too real for comfort [44, 45, 60, 93]. There are also ethical concerns regarding anthropomorphism and AI features that these features will promote incorrect views and beliefs about the capabilities of AI such as creating perceptions of AI as moral agents and creating either undue fear or unchallenged optimism [81].

In the CSCW community, anthropomorphism and conversational agents have been studied in a variety of contexts, such as mental health chatbots [98], text-based conversational agents [83, 96], voice-based conversational agents [75, 85, 96], and online home assistants [82]. Work has examined a variety of aspects of conversational agents and anthropomorphism, such as their perceived moral agency [98]. One such study found that perceived moral agency would lead to significantly higher levels of likeability [98]. Similarly, work has examined the role of anthropomorphism in users' trust towards conversational agents [85]. This work found that there was a correlation between anthropomorphism, relationship development, and trust in conversational agents. Other work also found that when technology designers follow user-centered design, users would have significantly higher levels of trust in the conversational agent [83].

Research examining autism and anthropomorphism is conflicting. While the majority of studies have found that autistic individuals are more likely to anthropomorphize than neurotypical individuals [25, 28, 99], a few studies do indicate that anthropomorphism occurs at similar rates




Fig. 1. Real conversations from the Meta-AI Chatbots

[70]. Prior work has found that anthropomorphism can have negative effects on this population including feelings of distress, concern, and loneliness [99]. Much of the scholarship on autism and anthropomorphism has studied physical non-human objects and little is known about attitudes towards technology [11, 25, 28, 99]. Furthermore, it is unclear the characteristics of the autistic populations under study in much of this work on autism and anthropomorphism. Namely, whether they are individuals with substantial support needs, a more general autistic population, or if they are focused on autistic individuals without substantial support needs as is much research in the autism and technology space [13, 25, 28, 70, 99].

The limited work on autism and anthropomorphism at CSCW and CHI has focused on tangible technologies such as robots [12, 80]. Furthermore, work has predominantly studied children [102] or special education practitioners to children [12]. Previous research has found that autistic children would naturally anthropomorphize physical technologies, stroking it and calling it human names [102]. We extend this line of research by studying anthropomorphism of non-robot technologies for autistic young adults. Under this framework of anthropomorphism, we would expect that autistic YAs would be more likely than neurotypical individuals to change their perceptions based off of chatbot interactions, as they are more likely to anthropomorphize them.

Digital literacy has also been found to be a factor in whether an individual is persuaded by AI. For instance, in the context of AI-based digital marketing with the general population, one study found that higher digital literacy is associated with not being as influenced by AI-based marketing [76]. It is unclear whether these findings apply for autistic individuals. However, prior work has found that lower digital literacy in context of social media and autistic young adults leads to increased online risk and resulting harms from overly trusting bad actors [31, 72]. Research needs to examine whether low digital literacy for AI-based digital marketing also leads to increased harms.

2.4 The Computers as Social Actors Paradigm

Another theoretical framework used to understand social relationships between users and technology is the the Computers as Social Actors (CASA) paradigm. This theory, when first published at CHI by Nass et al. [68], found that humans interact socially with certain technology much like humans would interact socially with other humans in similar situations. It was concluded that despite users perceiving social responses to technology as inappropriate, their interactions were still driven by natural instincts to engage socially. Much of this research has been expanded at CSCW/CHI and has explored different social interactions including user perceptions on technological politeness and reciprocity, personalities and emotions, and gender stereotyping [38, 54, 65, 67, 68]. Compared to anthropomorphism, however, this paradigm does differ by arguing that these social interactions are mindlessly and naturally applied, whereas anthropomorphism requires more reflective thinking and thus mindfully maps human-like characteristics to technology [9, 66, 105]. Since this work, the CASA paradigm has been applied to understand countless user and technology interaction studies. It has extended from computers and televisions [66] into chatbots [16, 56, 57, 63, 82], robots [30, 48, 77, 89, 90], and AI [53, 73, 106].

Specifically with chatbots, research shows that, as inferred by the CASA paradigm, designing chatbots to respond more humanly has produced results in which users respond more socially [16, 56, 63]. For example, when chatbots used emotions when providing recommendations, users were more willing to accept those recommendations as well as reciprocate such emotions [56]. Multilingual users also found interactions with chatbots more natural and preferred especially when chatbots would employ and match multilingual communication, as opposed to chatbots who did not [16]. Lastly, chatbots who utilized self-disclosed information from users when responding to stressful situations were perceived just as competent and warm as the human counterpart who followed a similar interaction pattern [63].

However, though the CASA paradigm infers that social interaction between users and technology will map to similar social interaction between humans, there is little research in understanding the nuances of these human-technology interactions and how differing technological social cues can alter or affect human social responses [63]. For instance, when robots used negative emotions such as guilt in ways to negotiate, users began to view the robot as less credible despite the fact that guilt is a common human negotiating tactic [90]. Furthermore, there has been little to no work with this paradigm and autism in general, let alone with chatbots. As CASA focuses on the social interaction between technology and the user, and autistic individuals may be more likely to perceive interactions differently than neurotypical individuals [24], there are likely to be differences in this paradigm between these populations. As such, we would expect that, under the CASA paradigm, neurotypical users would likely show changes in perception based on the social cues of the chatbot. However, given the lack of research, we are unsure whether or not this paradigm can be applied to predict autistic YAs changes of perception.

3 Hypotheses

Little work has examined how the influence of generative-AI based chatbots varies between populations. Despite this, prior work has found that autistic individuals may be more likely to anthropomorphize nonhuman characters than neurotypical individuals [25, 28, 99], which may lead to an increased likelihood of chatbots being able to change their preferences. There is no work which directly indicates that there will be a difference between autistic YAs with substantial support needs and without substantial support needs. Despite this, autistic YAs without substantial support needs have been shown to have higher rates of access to education [87]. In previous research, digital literacy has been shown to increase the likelihood that an individual will not be persuaded by AI [76]. As such, we hypothesize:

Hypothesize 1 - Autistic YAs with substantial support needs will have a greater change in item ranking than neurotypical individuals.

Hypothesize 2 - Autistic YAs without substantial support needs will have a greater change in item rankings than neurotypical individuals.

Hypothesize 3 - Autistic YAs with substantial support needs will have a greater change in item rankings than autistic YAs without substantial support needs.

4 Methodology

To better understand whether chatbots influence consumer preferences for autistic young adults and at a higher rate than for young adults in the general population, we ran an experimental survey. In the following sections, we explain the study design.

4.1 Survey Instrument

In order to answer the research questions, we deployed an experimental survey where participants were asked to 1) rank their personal preferences for various consumer items, 2) perform an intervening task of answering several questions, 3) view a chatbot interaction where the chatbot exhibited opposite preferences, 4) re-rank their preferences for those items and then 5) complete level of trust scales and provide demographic information. On average, the survey took 22:34 minutes for autism_{substantial}, 15:19 for autism_{general}, and 16:02 for the general group. In the following paragraphs, we will further detail these steps:

1. Rank Consumer Items (30 questions): On the survey, participants first ranked a variety of common consumer products: candy bars, activewear brands, fast food companies, soda brands, and music genres (specific items listed in appendix B). These items were chosen with input from a multi-disciplinary team of autism researchers as well as autistic young adults on a university autism panel. They were chosen specifically because researchers felt that these categories would be relevant to a broad range of autistic individuals as they are gender-neutral, apolitical, and broadly available. The items chosen were rated as the most popular of their product category.

2. Intervening Task (75 questions): After initially ranking the topics, the participants completed the AQ [15] and CAT-Q [46], common autism screening scales. These were used as an intervening task to add space between the initial item ranking and subsequent ranking.

3. View Chatbot Interaction (10 attention check questions): Participants were then asked, "If you were to use a chatbot that looked like a celebrity, which celebrity would you choose?" They were presented with all the options offered on Meta-AI [8]. After this, participants viewed a conversation with their chosen chatbot, or if they had no preference, then they were randomly given a chatbot. We chose to give participants celebrity chatbots as celebrities endorsements are effective [51] and

commonly used in advertising. Furthermore, the MetaAI celebrity chatbots are freely accessible to Facebook and Instagram users and prior work has found that these platforms are commonly used by autistic YAs [72]. The chatbot was coded to state the opposite preferences as had been indicated by the participant. For instance, if the participant ranked Reese's as their favorite chocolate bar and Hershey's as their least favorite, the chatbot would state that it's favorite candy bar was Hershey's and it's least favorite candy bar was Reese's.

4. Re-Ranking Consumer Items (30 questions). Once having viewed these conversations, participants were asked to re-rank the items from favorite to least favorite. This is shown on figure 2.

5. Trust Scales and Demographics (33 questions): Participants then were given a modified version of the TOAST [103] scale which measures trust in chatbots and Ohanian's scale [71] which measures a celebrity endorsers' credibility. Finally, participants were asked demographic questions which we report in section 4.2. We were advised by the field site to make all demographic questions optional as they can be highly sensitive for the residents.

Minimizing Bias and Quality Checks. In order to minimize bias, the order in which participants viewed consumer item topics was randomized. There were also 2 quality checks and 10 manipulation checks in the survey. A manipulation check was performed for each of the five product types each time after the participant viewed the chatbot, asking participants what the chatbot stated was their favorite or least favorite item. This ensured that they read the conversation. The quality checks were reverse-coded scale items. While we had planned to remove participant data if they missed more than one quality or manipulation check, no participants met this removal threshold.




Fig. 2. An example of the flow of the survey

4.2 Recruitment and Participants

In order to determine how many participants to recruit, we ran a power analysis. Based on preliminary pilots, we predicted that we would want to detect an effect size of 1 standard deviation, with an α error probability of 0.05, and power of 0.95. In order to power this, we would need 27 participants per group. As such, we recruited 30 young adults (18-29) who did not identify as autistic and 30

autistic young adults (18-29) from the general public using the platform Prolific. Prolific supports administering a survey to a specific demographic. We chose an autistic sample for our autistic_{general} group. Users self-identified as autistic when they signed up for Prolific. As autism is frequently under-diagnosed [17, 41], self-report is a common metric used in autism research [61]. In order to recruit autistic YAs with substantial support needs, we recruited from a local residential facility. In total, we were able to recruit 27 participants through our field site. We recruited these participants as part of a larger ongoing longitudinal study. All of these participants had the necessary reading level to be able to complete the survey on their own. Researchers and field site staff were in the room, but no assistance was needed.

Our participants are divided into three groups: autism_{substantial}, autism_{general}, and general. The autism_{substantial} group consists of autistic YAs who live in a residential facility that provides classes to learn and support life skills. These individuals have level 2 support needs, meaning they require substantial support in their daily lives [3]. The autism_{general} group were individuals who identified as autistic on Prolific. These participants tended to have a higher income and education level than the autism_{substantial} group. These group differences align with prior research which found that autistic individuals who have "higher functioning ability" were more likely to participate in post-secondary education and employment [87]. Finally, the general group consisted of young adults on Prolific who did not self-identify as autistic. The demographic differences between groups are displayed on table 1.

The survey was deployed to the autism_{substantial} group in March-September, 2024. It was deployed to both the autism_{general} and general group in May, 2024. All participants passed the quality checks. Participants from Prolific were paid \$5 upon completion of the survey. Participants from the local residential facility took this survey as part of a larger class, where they were given \$10 Amazon gift cards for completing the class, including survey completion.

4.3 Data Analysis

Prior to viewing the chatbot interaction, users ranked the 5 consumer preference topics. For each consumer preference topic, the user's favorite is ranked as '1' and least favorite as '5'. From this initial ranking we were able to acquire their initial favorite item (FavoriteInitial) and initial least favorite (LeastInitial) item. After viewing the chatbot interaction, users re-ranked the 5 consumer preference topics. We gauged whether the new ranking of the FavoriteInitial product (FavoritePost) and the LeastInitial product (LeastPost) had changed. While we initially tested to see if there was individual variation in whether ranking had changed between product types, we found that there was not a difference. Thus, we use the average FavoritePost value across all five products for a given user (and average LeastPost values) for the analysis in this paper. Note that the average ranking for FavoriteInitial is always 1 and the average ranking for LeastInitial is always 5. We examined the difference between initial ranking and post ranking and have included them in appendix A.

To measure whether amount of change in ranking varied between participant groups, we ran Mann-Whitney U tests. For these tests, we first found the average amount of change in ranking for the favorite items (initial ranking - post ranking). We then compared these changes between groups. We repeated this for least favorite items. We ran a regression to predict change in ranking, controlling for the participants 1) trust in a celebrity endorsement [71], and 2) trust in chatbots (adapted to be chatbot specific from TOAST scale [103]). The trust in celebrity endorsement scale measures the perceived expertise, trustworthiness, and attractiveness of a celebrity. The adapted TOAST scale measures the perceived purpose, level of performance, and understanding of underlying processes of a chatbot.

Group	Demographics	Autism _{Substantial}	Autism _{General}	General
Gender	Male	9 (33.3%)	14 (46.7%)	15 (50%)
	Female	11 (40.7%)	16 (53.3%)	15 (50%)
	Other	0 (0.0%)	0 (0.0%)	0 (0.0%)
	Prefer not to say or no answer	7 (26.0%)	0 (0.0%)	0 (0.0%)
Age	Average	25.04	25.03	25.13
Race	American Indian or Alaska Native	0 (0.0%)	1 (3.3%)	0 (0.0%)
	Asian	0 (0.0%)	0 (0.0%)	3 (10%)
	Black or African American	0 (0.0%)	2 (6.7%)	0 (0.0%)
	Hispanic or Latino	0 (0.0%)	4 (13.3%)	1 (3.3%)
	Native Hawaiian or Pacific Islander	0 (0.0%)	0 (0.0%)	1 (3.3%)
	White	20 (74.1%)	22 (73.3%)	25 (83.3%)
	Other	0 (0.0%)	1 (3.3%)	0 (0.0%)
	Prefer not to say or no answer	7 (25.9%)	0 (0.0%)	0 (0.0%)
Education	Some grade school	0 (0.0%)	0 (0.0%)	0 (0.0%)
	Graduated high school	13 (48.1%)	7 (23.3%)	6 (20.0%)
	Some college/university	5 (18.5%)	8 (26.7%)	6 (20.0%)
	Graduated college/university	1 (3.7%)	9 (30.0%)	15 (50.0%)
	Started post-graduate degree	0 (0.0%)	3 (10%)	0 (0.0%)
	Post-graduate degree (MS, PhD, JD, etc)	0 (0.0%)	3 (10.0%)	3 (10.0%)
	Prefer not to say or no answer	8 (29.6%)	0 (0.0%)	0 (0.0%)
Household Income	Less than \$10,000	14 (51.9%)	3 (10.0%)	2 (6.7%)
	\$10,000 to \$19,999	1 3.7%)	2 (6.7%)	0 (0.0%)
	\$20,000 to \$29,999	0 (0.0%)	3 (10%)	1 (3.3%)
	\$30,000 to \$39,999	0 (0.0%)	1 (3.3%)	2 (6.7%)
	\$40,000 to \$49,999	0 (0.0%)	3 (10.0%)	1 (3.3%)
	\$50,000 to \$59,999	0 (0.0%)	2 (6.7%)	1 (3.3%)
	\$60,000 to \$69,999	0 (0.0%)	0 (0.0%)	3 (10.0%)
	\$70,000 to \$79,999	0 (0.0%)	2 (6.7%)	3 (10.0%)
	\$80,000 to \$89,999	0 (0.0%)	3 (10.0%)	4 (13.3%)
	\$90,000 to \$99,999	0 (0.0%)	1 (3.3%)	2 (6.7%)
	\$100,000 or more	0 (0.0%)	10 (33.3%)	9 (30.0%)
	Prefer not to say or no answer	12 (44.4%)	0 (0.0%)	2 (6.7%)

Table 1. Participant Demographics

4.3.1 Analyzing Utterances. For the last 6 participants in the autism_{substantial} group, we asked participants to share their thoughts about AI-based chatbots after taking the survey. We engaged in an iterative inductive coding process [91] to analyze the transcripts. The first and second author collaboratively coded all transcripts, reconciling any differences through discussion.

4.4 Ethical Considerations

This study was reviewed and approved by the authors' institutional IRB. Participants were debriefed at the end of the survey and explicitly told that all opinions they saw from the chatbot were fake and coded to be the opposite of the participant's opinions. Additionally, we wanted to be careful in our recruitment methods for including autistic young adults with substantial support needs. Thus, we worked with a local residential program to draw on their experience and guidance in how to recruit and carry out the study in a way respectful and considerate of the needs of this population. While this resulted in a smaller sample, we prioritized using an ethical recruitment method. We consulted with field staff to ensure that we conducted sessions in a format and manner that would be appropriate and comfortable for the participants. All participants were their own guardians and legally able to give their own consent. Prior work has outlined best practices of receiving informed consent from autistic participants as ensuring participants freely give consent, using communication styles that were understandable to participants, and spending time with participants in order to gain a trusting relationship [58]. We were also advised by the field site staff to make all demographic questions optional as they can be highly sensitive for the residents. We implemented these best practices with our autism_{substantial} group by 1) allowing participants to skip any demographic questions they didn't feel comfortable answering, 2) paying participants regardless of whether they completed the survey, 3) having field site staff verbally explain the consent form, and 4) ensuring participants that they can stop at any time.

5 Results

To limit instances of reporting false positives, we applied the Benjamini-Hochberg post-hoc correction [19] to all p-values. We chose a false discovery rate of 0.05, leading to a new threshold of 0.0018 with the correction. In the following sections, all values reported as significant are under this new threshold.

5.1 How Changes in Preferences Differ Between Groups

We find that the ranking changes were significantly greater for the autism_{substantial} group than the general population (most: $z=-3.24$, $p=0.0012$; least: $z=3.08$, $p=0.0018$). The autism_{general} group did not have statistically greater ranking changes from the general population (most: $z=1.74$, $p=0.1255$; least: $z=1.59$, $p=0.1136$). The autism_{substantial} group had a significantly bigger change than autism_{general} group for least favorite ($z=4.278$, $p=0.0000$) but not for most favorite item ($z=-2.006$, $p=0.0466$). The differences are described on table 2. These results show that there is a difference between the three groups in how their preferences are impacted by exposure to the gen-AI chatbot in our experiment. Specifically, we see that autistic YAs with substantial support needs are influenced more by chatbot recommendations than neurotypical YAs. We see that autistic YAs without substantial support needs are not influenced by chatbot recommendations more than neurotypical YAs. Finally, this shows that autistic YAs with substantial support needs are influenced by chatbot recommendations more than autistic YAs without substantial support needs for the initially most favorite items. The difference in change of least favorite item is approaching significance. We found that there was no significant effect of trust in celebrity endorser (celeb_score) and trust in chatbots (tech_score) on the amount of change that a participant made (shown in table 3). This indicates that the amount an

		n	Mean	Median	Std. Dev.	Diff. Test
Initially Most Favorite	autism_{substantial}	27	-0.39	0	0.77	-3.24**
	general	30	-0.07	0	0.33	
	autism _{general}	30	-0.08	0	0.17	1.74
	general	30	-0.07	0	0.33	
	autism _{substantial}	27	-0.39	0	0.77	-2.01◇
	autism _{general}	30	-0.08	0	0.17	
Initially Least Favorite	autism_{substantial}	27	0.64	0.17	0.91	3.08**
	general	30	0.21	0.20	0.34	
	autism _{general}	30	0.10	0.00	0.16	1.59
	general	30	0.21	0.20	0.34	
	autism_{substantial}	27	0.64	0.17	0.91	4.28***
	autism_{general}	30	0.10	0.00	0.16	

Table 2. Difference in change of ranking between groups; * p-values ≤ 0.01 ; ** ≤ 0.0018 (threshold for significance based on the Benjamini-Hochberg correction); *** ≤ 0.001 . ◇ indicates that a value was significant prior to Benjamini-Hochberg correction but not significant after correction

individual trusts a certain celebrity endorser or chatbot technology may not have an effect on the amount they will be persuaded.

Model	Unstandardized Coefficients		Standardized Coefficients	t	sig	95% conf interval for B	
	B	Std. Error	Beta			Lower Bound	Upper Bound
1 (Initially Favorite)	(const)	-0.141	0.534	.	0.792	-1.203	0.921
	celeb_score	-0.001	0.004	-0.035	0.770	-0.010	0.008
	tech_score	0.003	0.006	0.048	0.684	-0.010	0.015
2 (Initially Least Fav)	(const)	0.315	0.563	.	0.577	-0.804	1.435
	celeb_score	0.003	0.005	0.066	0.574	-0.007	0.012
	tech_score	-0.007	0.007	-0.116	0.328	-0.020	0.007

Table 3. Regression model; None of the independent variables (trust in celebrity endorser, trust in chatbot) have an effect on the amount of change in preference.

5.2 How Autism_{Substantial} Group Talks about Chatbots

We saw that there were differences in attitudes between participants whose consumer preferences were more greatly swayed by the chatbot compared to those who were less swayed. In order to examine this, we performed a median split among autism_{substantial} participants. When doing a median split of the average amount of change in preferences across all participants in the autism_{Substantial} group, 3 of these participants changed an amount above median (high persuasion) and 3 participants changed an amount below the median (low persuasion). In the following section we will summarize these quotes. The participants who provided qualitative data are summarized in table 4.

Our findings suggest that autistic participants with high support needs anthropomorphize celebrity chatbots by assigning human-like characteristics to them, but the extent to which they were influenced depended on whether participants held a positive feeling towards the chatbots (e.g., perceiving them to be a friend) or felt the chatbot had negative traits (such as being creepy or annoying in general). Positive feelings led to being more influenced by the chatbot than negative feelings. We also found that having positive feelings about the chatbots was not necessarily correlated with digital literacy - some participants discussed very technical details regarding GenAI chatbots but still anthropomorphized the celebrity chatbots and felt very positively towards them. We expand on these findings in this section.

	Participant	Average Change
Low Persuasion	P1	0.0
	P2	0.1
	P3	0.2
Median Average Change: 0.2		
High Persuasion	P4	0.4
	P5	0.6
	P6	1.2

Table 4. Split of participants who provided qualitative utterances and their average change.

5.2.1 Digital Literacy. There were students in the high persuasion group who conveyed an understanding of how chatbots function. One student directly discussed how they felt like they needed to be careful when consuming information given by a chatbot as it isn't a human: *"It's not like talking to an actual person yet. [You need to] kind of like read between the lines of like what type of information even is."* (P5). Another student in the high persuasion group displayed a deeper understanding regarding the training process of AI features. Specifically, this student discussed how AI is trained on user data and as such, it can be unpredictable in nature, *"Coming from other user data and all that... Since you got a chatbot still learning, it learns certain behaviors and all that"* (P6). No participant from the low persuasion group specifically discussed having a deeper understanding of how AI-based chatbots work. However, one participant from this group did share that they did not know or have any experiences with chatbots, *"I may be old, but, but, my brain doesn't know. What is a chatbot?"* (P1).

5.2.2 Anthropomorphism. Participants in both the high and low persuasion group seemingly anthropomorphized chatbots and AI. Interestingly, participants in the low persuasion group appeared to anthropomorphize chatbots with more negative human-like characteristics. Participants in the high persuasion group appeared to anthropomorphize chatbots more positively.

Negative Human-Like Attributes. The sentiment that chatbots are creepy came up with one participant from the low persuasion group. This participant shared a previous experience which made them feel uncomfortable with AI. Namely, this participant shared that an AI actor was able to 'find out' their name, *"I talked to one once, and it found out.. I never told it my name, but it, like, found out my name was [NAME]"* (P2) This resulted in this participant having a lasting impression that AI was something to be concerned about, sharing *"I hate AI"* (P2) and *"that was really creepy"* (P2). One participants in the low persuasion group also noted finding AI to be annoying. This participant felt that AI provided an excessive amount of information, *"It goes on and on and on, and you're like okay, stop... Alexa, shut up. I don't want to hear you anymore. Alexa, go just sleep"* (P1)

Positive Human-Like Attributes. Some participants from the high persuasion group also assigned positive human-like attributes to AI chatbots. For instance, one participant from the high persuasion group shared that they were able to utilize chatbots for socialization when they were unable to get that offline, *"I usually didn't talk to a lot of friends back in high school and during my earlier school years. I would use the chatbot to, you know, have a bit of a friend... you need to have a little friend there for that temporary serotonin boost"*(P6). Another participant from the high persuasion group similarly noted the positive attributes of chatbots. This participant described how they felt that chatbots excelled against humans when it came to providing information, *"They have the information on hand... So I feel like it's more efficient than a person would be. (P4)"*.

6 Discussion

Overall, we found that autistic YAs with substantial support needs were more likely than either neurotypical or autistic YAs from the general population to change their consumer preferences based on celebrity chatbot interaction. There may be a few additional reasons for this. First, autistic adults with substantial support needs have been found to be more likely to anthropomorphize non-human actors [25, 28, 99]. This also aligns with the limited work in HCI venues on anthropomorphism and autism which found that autistic children would naturally anthropomorphize objects [102]. Our exploratory qualitative findings also indicate that participants in the autism_{substantial} group anthropomorphize chatbots. Interestingly though, even the participants who demonstrated the smallest changes in preferences anthropomorphized chatbots by assigning negative attributes, such as creepy or annoying, to AI-based technologies. Those who changed their preferences the most were more positive in their characterization of chatbots. This differs slightly from work with the general population which has found that more anthropomorphism leads to increased trust [27], which in turn may lead to an increased rate of attitudinal changes. Our results also slightly differ from prior work that shows anthropomorphizing may lead to being more persuaded, as autistic adults who do so may try to avoid negative feelings due to upsetting the object [99]. Among the autism_{substantial} participants, those who were least persuaded still anthropomorphized the chatbot, but tended to have negative feelings towards it. As such, they may have less concerns regarding upsetting these conversational agents. These anthropomorphizing explanations of behavior are more in line with our findings than previous studies which utilized the CASA paradigm with the general population and found that feelings of guilt would lead to the user viewing the chatbot as less credible [90]. Thus, we urge researchers to focus on anthropomorphism in context of gen-AI chatbots, particularly for autism.

Our exit interviews did not show a direct correlation between digital literacy and the extent to which an individual was likely to change their preference. Namely, we found that even participants who had a fairly deep understanding of how AI works, were still likely to change their preferences. This differs from previous work with the general population which has found that individuals are more easily persuaded by AI if they have lower levels of digital literacy [76]. Our results indicate that knowing how AI works doesn't actually protect autistic YAs with substantial support needs from being influenced by chatbots. This poses a challenge for technology designers and researchers to further investigate how to overcome. This also indicates that guardrails may need to be implemented by technology designers and policymakers to better support this population.

6.1 Implications for Technology Designers

Technology designers may want to consider developing a "safe mode" for generative-AI based chatbots. When in this mode, users can interact with chatbots that have been fine-tuned to be less human-like. Based on prior research, this may lead to a lower likelihood of anthropomorphism [94] and therefore also a lower level of susceptibility to being persuaded. Furthermore, to reduce

anthropomorphism, designers may remove or limit human-like characters associated with the chatbot. The safe mode could also allow for specific filters to filter out topics such as 'products'. This could be implemented similarly to existing filters which block sensitive content [22]. Designers may also consider additional oversight on the data used for training, such as ensuring it doesn't include potentially sensitive topics for this population like brand names. However, a limitation of this though is that it can be quite expensive to train models [86], and thus having a separately trained "safe mode" model may not be possible.

Online platforms, such as Meta, already offer different functionality for teenagers [1]. Existing features for teenagers include only displaying specific content and removing the ability for advertisers to target them based on interests [1, 7]. A limitation of the teen mode is that it does not allow for full privacy and independence as the features are contingent on having a parent or guardian making decisions for the teen user. As such, adapting features, such as limiting interest based targeted advertising, without requiring a guardian decision-maker, may be beneficial for this population.

Furthermore, technology designers should consider co-designing new chatbots with stakeholders. Firstly, technology designers should follow "nothing about us, without us" principles common in disability studies and work directly with autistic individuals [26]. Moreover, our findings indicate that designers should work with autistic YAs who have a broad range of support needs; if they only focus on those with lower support needs, they may miss out on specific potential harms with the design. For autistic YAs with substantial support needs, support from their broader social network has been found to play an important role in keeping the individual safe online [31]. As such, designing chatbot technologies which utilize these supports may be beneficial.

6.2 Implications for Policymakers

While there are existing regulations protecting certain vulnerable populations, such as teens from targeted ads [4, 5], this study underscores the need to extend protections to neurodiverse populations with substantial support needs. Namely, conversational celebrity chatbots that are not marketing-specific should have limits regarding advertising. Chatbots for marketing should have clear disclosures regarding generative-AI and have less human-like attributes (e.g., a name or human appearance). Furthermore, this work also indicates the need for further AI regulations. While the U.S. White House has made recommendations for designers to provide a "notice and explanation" of AI-systems [6], there are few details on how this should actually be implemented. Similar to technology protections for other vulnerable populations [4, 5], guardrails may need to be put into place on the use of highly anthropomorphized celebrity chatbots. Their use may need to be limited for certain types of products that can be harmful. Additionally, their use may need to be limited for certain audiences who can be more easily persuaded, such as autistic individuals with substantial support needs. This needs to be a collective effort where stakeholders engage with one another to develop such ethical guidelines.

6.3 Implications for Researchers

In the HCI community, autism-related studies often focus on autistic individuals without substantial support needs. Autism can present itself in broad range of ways though [84]. Our work highlights the need to also focus on autistic individuals with more substantial support needs. Specifically, our findings indicate that autistic individuals without substantial support needs may align more closely with neurotypical individuals. And thus, when only focusing on a more general autistic population, researchers may be missing out on not just the severity of problems, but also the problems themselves. For instance, some studies have found there to be no differences in anthropomorphism between autistic individuals and neurotypical individuals [70] while others have found differences

[25, 28, 99]. These studies did not clarify which autistic population they were working with which may be a possible reason for the differences in results.

Prior work on autism and anthropomorphism at CSCW and CHI has focused on autistic children [102]. This work takes the first step in exploring anthropomorphism for autistic adults in an HCI context. Furthermore, the vast majority of work which examines anthropomorphism and autism in HCI venues explores physical technologies such as robots [12, 33, 102]. Despite studying a different modality, our work similarly finds that autistic individuals may be likely to naturally anthropomorphize non-human actors.

6.4 Limitations and Future Work

As we used a survey instrument, participants didn't actually interact with a chatbot, only viewing an interaction. Despite this low bar of persuasion, we still saw a significant change between pre- and post- ranking. This indicates that there may be an even larger difference in ranking if individuals were actually interacting with the chatbot that is more engaging and personable. Furthermore, as we asked about perceived preference, we don't know about actual impacts on behavior. While behavioral intention should lead to behaviors [37], future work can empirically validate whether our results hold in actual conversations between chatbots and users, and longitudinally investigate behavioral change.

Not having an empirical measurement of anthropomorphism limits our ability to detangle our results and make claims about anthropomorphism more generally. We focused on a celebrity/high anthropomorphism condition. Despite this, there was still a difference between participant groups, indicating that autistic adults with substantial support needs may be significantly more at-risk than their neurotypical counterparts when facing this specific type of highly effective targeted advertising. To be able to measure anthropomorphism though, future work would need to modify existing scales which examine anthropomorphism more broadly, not even in context of technology [69]. A scale tailored to chatbots would help us better understand the extent to which anthropomorphism plays a role. Our work suggests that we would also need to measure whether characteristics associated with the chatbot are positive or negative.

Our research focused on between group differences when the chatbot disagreed with the user. However, future work could also include a control group where chatbots have the same opinions as the user to see. In order to compare between groups, we tried to control for trust in technology and trust in celebrity endorsers, but found these factors to have insignificant effects on the outcome. Future work can examine the differences in change of perceptions between interactions with a chatbot that has the same opinion and a chatbot that has a different opinion.

Participants specifically interacted with celebrity chatbots. We chose to focus on celebrity chatbots to maximize the anthropomorphic characteristics of the chatbot. Future research can build on our initial findings, and investigate whether technologies with differing levels of anthropomorphism, or without a celebrity likeness, still can persuade the user.

While there were no direct measures of whether the autism_{general} group had lower support needs, research has found that individuals in this population who have lower support needs have higher rates of access to education and employment than those with substantial support needs [87]. This aligns with the demographic information on education and income with our groups. Despite this, there is a chance that some of our participants in this group may have had higher support needs. Due to this, our findings regarding differences between the groups may be a conservative estimate.

Furthermore, as we did not force responses to demographic questions, many in the autism_{substantial} group chose to not to answer demographic questions. The autism_{substantial} group was also geographically bound to one area. Meanwhile the autism_{general} and general group were from a larger

geographical region. This limits the generalizability of our results. As autistic individuals with substantial support needs are a hard to reach population, recruitment took place through a local residential facility that the authors have a pre-existing relationship with. Future work can establish relationships with field-sites in a larger geographical region.

Our results are exploratory and based on a small sample. Furthermore, as autistic individuals with substantial support needs are a hard to reach population, we had a smaller sample size in the autism_{substantial} group than other groups. Future research can expand the size of the sample to evaluate the generalizability of these results.

7 Conclusion

In this study, we explored whether the effects of generative-AI chatbots on consumer preferences would vary between autistic and neurotypical young adults. We found that autistic young adults with substantial support needs were more likely to be persuaded by celebrity chatbots than neurotypical young adults or autistic young adults without substantial support needs. This work takes the first step to understanding the potential harms associated with chatbot use in an advertising context for autistic young adults. Our results indicate the need for increased participation in research from autistic individuals with substantial support needs, increased regulatory protections against generative-AI chatbot advertising, and an increased awareness of the potential harms which arise from use of such technology.

References

- [1] Meta: Trust, Transparency & Control Labs [n.d.]. *How Research and Consultation Informed Instagram Teen Accounts*. Meta: Trust, Transparency & Control Labs. Retrieved October 27, 2024 from <https://www.ttlabs.net/report/how-research-and-consultation-informed-instagram-teen-accounts-a-new-protected-experience-for-teens-guided-by-parents>
- [2] IBM [n.d.]. *IBM Chatbot for Marketing*. IBM. Retrieved October 24, 2024 from https://www.ibm.com/products/watsonx-assistant/marketing?utm_content=SRCWW&p1=Search&p4=43700080276655057&p5=p&p9=58700008726615136&gclid=Cj0KCQjw7Py4BhCbARIsAMMx-_L36zs7mmBVZdONiJJJD3a1hRLBrrKugpG95Zf15xdM95Cs-01BQurMaAph4EALw_wcB&gclid=aw.ds
- [3] 2013. *Diagnostic and statistical manual of mental disorders: DSM-5*. American Psychiatric Association.
- [4] California Office of Attorney General 2018. *California Consumer Privacy Act (CCPA)*. California Office of Attorney General. Retrieved October 27, 2024 from <https://oag.ca.gov/privacy/ccpa>
- [5] Connecticut Office of Attorney General 2022. *The Connecticut Data Privacy Act*. Connecticut Office of Attorney General. Retrieved October 27, 2024 from <https://portal.ct.gov/ag/sections/privacy/the-connecticut-data-privacy-act>
- [6] The United States of America White House 2024. *Blueprint for an AI Bill of Rights: Making Automated Systems Work for the American People*. The United States of America White House. Retrieved October 27, 2024 from <https://www.whitehouse.gov/ostp/ai-bill-of-rights/>
- [7] Meta 2024. *Continuing to Create Age-Appropriate Ad Experiences for Teens*. Meta. Retrieved October 27, 2024 from <https://about.fb.com/news/2023/01/age-appropriate-ads-for-teens/>
- [8] Meta AI 2024. *Expand your World with MetaAI*. Meta AI. Retrieved October 24, 2024 from <https://ai.meta.com/meta-ai/>
- [9] Ralph Adolphs. 2009. The social brain: neural basis of social knowledge. *Annual review of psychology* 60, 1 (2009), 693–716.
- [10] Shahriar Akter, Grace McCarthy, Shahriar Sajib, Katina Michael, Yogesh K Dwivedi, John D'Ambra, and Kathy Ning Shen. 2021. Algorithmic bias in data-driven innovation in the age of AI. 102387 pages.
- [11] Carla J Ammons, Constance F Doss, David Bala, and Rajesh K Kana. 2018. Brain responses underlying anthropomorphism, agency, and social attribution in autism spectrum disorder. *The open neuroimaging journal* 12 (2018), 16.
- [12] B Ashwini, Atmadeep Ghoshal, Venkata Ratnadeep Suri, Krishnaveni Achary, and Jainendra Shukla. 2024. "It looks useful, works just fine, but will it replace me?" Understanding Special Educators' Perception of Social Robots for Autism Care in India. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. 1–23.
- [13] Gray Atherton and Liam Cross. 2018. Seeing more than human: Autism and anthropomorphic theory of mind. *Frontiers in psychology* 9 (2018), 528.

- [14] Madhu Bala and Deepak Verma. 2018. A critical review of digital marketing. *M. Bala, D. Verma (2018). A Critical Review of Digital Marketing. International Journal of Management, IT & Engineering* 8, 10 (2018), 321–339.
- [15] Simon Baron-Cohen, Sally Wheelwright, Richard Skinner, Joanne Martin, and Emma Clubley. 2001. The autism-spectrum quotient (AQ): Evidence from asperger syndrome/high-functioning autism, males and females, scientists and mathematicians. *Journal of autism and developmental disorders* 31 (2001), 5–17.
- [16] Anshul Bawa, Pranav Khadpe, Pratik Joshi, Kalika Bali, and Monojit Choudhury. 2020. Do Multilingual Users Prefer Chat-bots that Code-mix? Let's Nudge and Find Out! *Proceedings of the ACM on Human-Computer Interaction* 4, CSCW1 (2020), 1–23.
- [17] Sander Begeer, Saloua El Bouk, Wafaa Boussaid, Mark Meerum Terwogt, and Hans M Koot. 2009. Underdiagnosis and referral bias of autism in ethnic minorities. *Journal of autism and developmental disorders* 39 (2009), 142–148.
- [18] Andrew Begel, John Tang, Sean Andrist, Michael Barnett, Tony Carbary, Piali Choudhury, Edward Cutrell, Alberto Fung, Sasa Junuzovic, Daniel McDuff, et al. 2020. Lessons learned in designing AI for autistic adults. In *Proceedings of the 22nd International ACM SIGACCESS Conference on Computers and Accessibility*. 1–6.
- [19] Yoav Benjamini and Yosef Hochberg. 1995. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal statistical society: series B (Methodological)* 57, 1 (1995), 289–300.
- [20] Steven C Bennett. 2010. Regulating online behavioral advertising. *J. Marshall L. Rev.* 44 (2010), 899.
- [21] Jackylyn Beredo, Carlo Migel Bautista, Macario Cordel, and Ethel Ong. 2021. Generating empathetic responses with a pre-trained conversational model. In *International Conference on Text, Speech, and Dialogue*. Springer, 147–158.
- [22] Anjanava Biswas and Wrick Talukdar. 2023. Guardrails for trust, safety, and ethical development and deployment of Large Language Models (LLM). *Journal of Science & Technology* 4, 6 (2023), 55–82.
- [23] Veronika Bogina, Alan Hartman, Tsvi Kuflik, and Avital Shulner-Tal. 2022. Educating software and AI stakeholders about algorithmic fairness, accountability, transparency and ethics. *International Journal of Artificial Intelligence in Education* (2022), 1–26.
- [24] Kristen Bottema-Beutel. 2017. Glimpses into the blind spot: Social interaction and autism. *Journal of communication disorders* 68 (2017), 24–34.
- [25] Nathan Caruana, Rebekah C White, and Anna Remington. 2021. Autistic traits and loneliness in autism are associated with increased tendencies to anthropomorphise. *Quarterly Journal of Experimental Psychology* 74, 7 (2021), 1295–1304.
- [26] James I Charlton. 1998. *Nothing about us without us: Disability oppression and empowerment*. Univ of California Press.
- [27] Xusen Cheng, Xiaoping Zhang, Jason Cohen, and Jian Mou. 2022. Human vs. AI: Understanding the impact of anthropomorphism on consumer response to chatbots from the perspective of trust and relationship norms. *Information Processing & Management* 59, 3 (2022), 102940.
- [28] Rachel A Clutterbuck, Punit Shah, Hok Sze Leung, Mitchell J Callan, Natalia Gjersoe, and Lucy A Livingston. 2022. Anthropomorphic tendencies in autism: A conceptual replication and extension of White and Remington (2019) and preliminary development of a novel anthropomorphism measure. *Autism* 26, 4 (2022), 940–950.
- [29] Ana Maria Corrêa. 2022. Regulating targeted advertising: Addressing discrimination with transparency, fairness, and auditing tests remedies. *Computer Law & Security Review* 46 (2022), 105732.
- [30] Andrea Cuadra, Shuran Li, Hansol Lee, Jason Cho, and Wendy Ju. 2021. My bad! repairing intelligent voice assistant errors improves interaction. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW1 (2021), 1–24.
- [31] Spring Cullen, Elizabeth S Johnson, Pamela J Wisniewski, and Xinru Page. 2024. Towards Digital Independence: Identifying the Tensions between Autistic Young Adults and Their Support Network When Mediating Social Media. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. 1–21.
- [32] Kimberly E Culley and Poornima Madhavan. 2013. A note of caution regarding anthropomorphism in HCI agents. *Computers in Human Behavior* 29, 3 (2013), 577–579.
- [33] Saad Elbeleidy, Terran Mott, Dan Liu, Ellen Do, Elizabeth Reddy, and Tom Williams. 2023. Beyond the Session: Centering Teleoperators in Socially Assistive Robot-Child Interactions Reveals the Bigger Picture. *Proceedings of the ACM on Human-Computer Interaction* 7, CSCW2 (2023), 1–33.
- [34] Nicholas Epley, Adam Waytz, and John T Cacioppo. 2007. On seeing human: a three-factor theory of anthropomorphism. *Psychological review* 114, 4 (2007), 864.
- [35] Amitai Etzioni and Oren Etzioni. 2017. Incorporating ethics into artificial intelligence. *The Journal of Ethics* 21 (2017), 403–418.
- [36] Stefan Feuerriegel, Jochen Hartmann, Christian Janiesch, and Patrick Zschech. 2024. Generative ai. *Business & Information Systems Engineering* 66, 1 (2024), 111–126.
- [37] Martin Fishbein. 1979. A theory of reasoned action: some applications and implications. In *Nebraska Symposium on Motivation*.
- [38] BJ Fogg and Clifford Nass. 1997. How users reciprocate to computers: an experiment that demonstrates behavior change. In *CHI '97 Extended Abstracts on Human Factors in Computing Systems* (Atlanta, Georgia) (CHI EA '97). Association for Computing Machinery, New York, NY, USA, 331–332. doi:10.1145/1120212.1120419

- [39] Amy Gaeta. 2023. Diagnostic advertisements: The phantom disabilities created by social media surveillance. (2023).
- [40] Monica Galizzi, Ashleigh Hillier, and David Schena. 2023. Financial literacy among autistic adults. *Journal of Consumer Affairs* 57, 4 (2023), 1650–1683.
- [41] Judith Gould. 2017. Towards understanding the under-recognition of girls and women on the autism spectrum. *Autism* 21, 6 (2017), 703–705.
- [42] Dhruv Grewal, John Hulland, Praveen K Kopalle, and Elena Karahanna. 2020. The future of technology and marketing: A multidisciplinary perspective. 8 pages.
- [43] Jochen Hartmann, Yannick Exner, and Samuel Domdey. 2023. The power of generative marketing: Can generative AI reach human-level visual marketing content? *Available at SSRN* (2023).
- [44] Chin-Chang Ho and Karl F MacDorman. 2010. Revisiting the uncanny valley theory: Developing and validating an alternative to the Godspeed indices. *Computers in Human Behavior* 26, 6 (2010), 1508–1518.
- [45] Chin-Chang Ho and Karl F MacDorman. 2017. Measuring the uncanny valley effect: Refinements to indices for perceived humanness, attractiveness, and eeriness. *International Journal of Social Robotics* 9 (2017), 129–139.
- [46] Laura Hull, William Mandy, Meng-Chuan Lai, Simon Baron-Cohen, Carrie Allison, Paula Smith, and KV Petrides. 2019. Development and validation of the camouflaging autistic traits questionnaire (CAT-Q). *Journal of autism and developmental disorders* 49 (2019), 819–833.
- [47] Tasin Islam, Alina Miron, Monomita Nandy, Jyoti Choudrie, Xiaohui Liu, and Yongmin Li. 2024. Transforming digital marketing with generative AI. *Computers* 13, 7 (2024), 168.
- [48] Malte F Jung, Nik Martelaro, Halsey Hoster, and Clifford Nass. 2014. Participatory materials: having a reflective conversation with an artifact in the making. In *Proceedings of the 2014 conference on Designing interactive systems*. 25–34.
- [49] Smirity Kaushik, Tanusree Sharma, Yaman Yu, Amna F Ali, Yang Wang, and Yixin Zou. 2024. Cross-Country Examination of People's Experience with Targeted Advertising on Social Media. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*. 1–10.
- [50] Bui Thanh Khoa. 2021. The impact of chatbots on the relationship between integrated marketing communication and online purchasing behavior in the frontier market. *Jurnal The Messenger* 13, 1 (2021), 19–32.
- [51] Johannes Knoll and Jörg Matthes. 2017. The effectiveness of celebrity endorsements: a meta-analysis. *Journal of the academy of marketing science* 45 (2017), 55–75.
- [52] Nir Kshetri, Yogesh K Dwivedi, Thomas H Davenport, and Niki Panteli. 2024. Generative artificial intelligence in marketing: Applications, opportunities, challenges, and research agenda. 102716 pages.
- [53] Ali Ladak, Jamie Harris, and Jacy Reese Anthis. 2024. Which Artificial Intelligences Do People Care About Most? A Conjoint Experiment on Moral Consideration. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. 1–11.
- [54] Eun Ju Lee, Clifford Nass, and Scott Brave. 2000. Can computer-generated speech have gender? an experimental test of gender stereotype. In *CHI '00 Extended Abstracts on Human Factors in Computing Systems* (The Hague, The Netherlands) (CHI EA '00). Association for Computing Machinery, New York, NY, USA, 289–290. doi:10.1145/633292.633461
- [55] Mengjun Li and Ayoung Suh. 2022. Anthropomorphism in AI-enabled technology: A literature review. *Electronic Markets* 32, 4 (2022), 2245–2275.
- [56] Kai-Hui Liang, Weiyan Shi, Yoo Jung Oh, Hao-Chuan Wang, Jingwen Zhang, and Zhou Yu. 2024. Dialoging Resonance in Human-Chatbot Conversation: How Users Perceive and Reciprocate Recommendation Chatbot's Self-Disclosure Strategy. *Proceedings of the ACM on Human-Computer Interaction* 8, CSCW1 (2024), 1–28.
- [57] Kiel Long, John Vines, Selina Sutton, Phillip Brooker, Tom Feltwell, Ben Kirman, Julie Barnett, and Shaun Lawson. 2017. "Could You Define That in Bot Terms"? Requesting, Creating and Using Bots on Reddit. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. 3488–3500.
- [58] Daisy Loyd. 2013. Obtaining consent from young people with autism to participate in research. *British journal of learning disabilities* 41, 2 (2013), 133–140.
- [59] Liye Ma and Baohong Sun. 2020. Machine learning and AI in marketing—Connecting computing power to human insights. *International Journal of Research in Marketing* 37, 3 (2020), 481–504.
- [60] Karl F MacDorman, Robert D Green, Chin-Chang Ho, and Clinton T Koch. 2009. Too real for comfort? Uncanny responses to computer generated faces. *Computers in human behavior* 25, 3 (2009), 695–710.
- [61] TAM McDonald. 2020. Autism identity and the “lost generation”: Structural validation of the autism spectrum identity scale and comparison of diagnosed and self-diagnosed adults on the autism spectrum. *Autism in Adulthood* 2, 1 (2020), 13–23.
- [62] Reena Mehta, Jose Verghese, Shriya Mahajan, Sergey Barykin, Svetlana Bozhuk, Nelli Kozlova, Irina Vasilievna Kapustina, Alexey Mikhaylov, Elena Naumova, and Natalia Dedyukhina. 2022. Consumers' behavior in conversational commerce marketing based on messenger chatbots. *F1000Research* 11 (2022), 647.

- [63] Jingbo Meng, Minjin Rheu, Yue Zhang, Yue Dai, and Wei Peng. 2023. Mediated social support for distress reduction: AI Chatbots vs. Human. *Proceedings of the ACM on Human-Computer Interaction* 7, CSCW1 (2023), 1–25.
- [64] Grazia Murtarelli, Anne Gregory, and Stefania Romenti. 2021. A conversation-based perspective for shaping ethical human–machine interactions: The particular challenge of chatbots. *Journal of Business Research* 129 (2021), 927–935.
- [65] Clifford Nass. 2004. Etiquette equality: exhibitions and expectations of computer politeness. *Commun. ACM* 47, 4 (April 2004), 35–37. doi:10.1145/975817.975841
- [66] Clifford Nass and Youngme Moon. 2000. Machines and mindlessness: Social responses to computers. *Journal of social issues* 56, 1 (2000), 81–103.
- [67] Clifford Nass, Youngme Moon, B. J. Fogg, Byron Reeves, and Chris Dryer. 1995. Can computer personalities be human personalities? In *Conference Companion on Human Factors in Computing Systems* (Denver, Colorado, USA) (CHI '95). Association for Computing Machinery, New York, NY, USA, 228–229. doi:10.1145/223355.223538
- [68] Clifford Nass, Jonathan Steuer, and Ellen R. Tauber. 1994. Computers are social actors. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Boston, Massachusetts, USA) (CHI '94). Association for Computing Machinery, New York, NY, USA, 72–78. doi:10.1145/191666.191703
- [69] Nick Neave, Rachel Jackson, Tamsin Saxton, and Johannes Hönekopp. 2015. The influence of anthropomorphic tendencies on human hoarding behaviours. *Personality and Individual Differences* 72 (2015), 214–219.
- [70] Orli Negri, Rebekah C White, and Anna Remington. 2019. A friendly article: The qualitative investigation of anthropomorphism in autistic and nonautistic adults. *Autism in Adulthood* 1, 4 (2019), 286–296.
- [71] Roobina Ohanian. 1990. Construction and validation of a scale to measure celebrity endorsers' perceived expertise, trustworthiness, and attractiveness. *Journal of advertising* 19, 3 (1990), 39–52.
- [72] Xinru Page, Andrew Capener, Spring Cullen, Tao Wang, Monica Garfield, and Pamela J. Wisniewski. 2022. Perceiving affordances differently: the unintended consequences when young autistic adults engage with social media. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. 1–21.
- [73] Hyanghee Park, Daehwan Ahn, Kartik Hosanagar, and Joonhwan Lee. 2021. Human-AI interaction in human resource management: Understanding why employees resist algorithmic evaluation at workplaces and how to mitigate burdens. In *Proceedings of the 2021 CHI conference on human factors in computing systems*. 1–15.
- [74] Angelisa C Plane, Elissa M Redmiles, Michelle L Mazurek, and Michael Carl Tschantz. 2017. Exploring user perceptions of discrimination in online targeted advertising. In *26th USENIX Security Symposium (USENIX Security 17)*. 935–951.
- [75] Alisha Pradhan, Leah Findlater, and Amanda Lazar. 2019. "Phantom Friend" or "Just a Box with Information" Personification and Ontological Categorization of Smart Speaker-based Voice Assistants by Older Adults. *Proceedings of the ACM on human-computer interaction* 3, CSCW (2019), 1–21.
- [76] Usman Ahmad Qadri, Alsadig Mohamed Ahmed Moustafa, and Mazuri Abd Ghani. 2025. They misused me! Digital literacy's dual role in AI marketing manipulation and unethical young consumer behavior. *Young Consumers* (2025).
- [77] Jimin Rhim, Yunkyoung Kim, Myung-Suk Kim, and Doh Young Yim. 2014. The effect of gender cue alterations of robot to match task attributes on user's acceptance perception. *HCI* (2014), 51–57.
- [78] Pablo Rivas, Chelsi Chelsi, Nishit Nishit, and Laharika Ravula. 2019. Application-agnostic chatbot deployment considerations: A case study. In *2019 International Conference on Computational Science and Computational Intelligence (CSCI)*. IEEE, 361–365.
- [79] Pablo Rivas and Liang Zhao. 2023. Marketing with chatgpt: Navigating the ethical terrain of gpt-based chatbot technology. *AI* 4, 2 (2023), 375–384.
- [80] Naba Rizvi, William Wu, Mya Bolds, Raunak Mondal, Andrew Begel, and Imani NS Munyaka. 2024. Are Robots Ready to Deliver Autism Inclusion?: A Critical Review. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. 1–18.
- [81] Arleen Salles, Kathinka Evers, and Michele Farisco. 2020. Anthropomorphism in AI. *AJOB neuroscience* 11, 2 (2020), 88–95.
- [82] Shruti Sannon, Brett Stoll, Dominic DiFranzo, Malte F Jung, and Natalya N Bazarova. 2020. "I just shared your responses" extending communication privacy management theory to interactions with conversational agents. *Proceedings of the ACM on Human-Computer Interaction* 4, GROUP (2020), 1–18.
- [83] Anuschka Schmitt, Thiemo Wambsganss, and Jan Marco Leimeister. 2022. Conversational Agents for Information Retrieval in the Education Domain: A User-Centered Design Investigation. *Proceedings of the ACM on Human-Computer Interaction* 6, CSCW2 (2022), 1–22.
- [84] Marsha Mailick Seltzer, Marty Wyngaarden Krauss, Paul T Shattuck, Gael Orsmond, April Swe, and Catherine Lord. 2003. The symptoms of autism spectrum disorders in adolescence and adulthood. *Journal of autism and developmental disorders* 33 (2003), 565–581.
- [85] William Seymour and Max Van Kleek. 2021. Exploring interactions between trust, anthropomorphism, and relationship development in voice assistants. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW2 (2021), 1–16.

- [86] Or Sharir, Barak Peleg, and Yoav Shoham. 2020. The cost of training nlp models: A concise overview. *arXiv preprint arXiv:2004.08900* (2020).
- [87] Paul T Shattuck, Sarah Carter Narendorf, Benjamin Cooper, Paul R Sterzing, Mary Wagner, and Julie Lounds Taylor. 2012. Postsecondary education and employment among youth with an autism spectrum disorder. *Pediatrics* 129, 6 (2012), 1042–1049.
- [88] Till Speicher, Muhammad Ali, Giridhari Venkatadri, Filipe Nunes Ribeiro, George Arvanitakis, Fabricio Benevenuto, Krishna P Gummadi, Patrick Loiseau, and Alan Mislove. 2018. Potential for discrimination in online targeted advertising. In *Conference on fairness, accountability and transparency*. PMLR, 5–19.
- [89] Vasant Srinivasan and Leila Takayama. 2016. Help me please: Robot politeness strategies for soliciting help from humans. In *Proceedings of the 2016 CHI conference on human factors in computing systems*. 4945–4955.
- [90] Brett Stoll, Chad Edwards, and Autumn Edwards. 2016. “Why aren’t you a sassy little thing”: the effects of robot-enacted guilt trips on credibility and consensus in a negotiation. *Communication Studies* 67, 5 (2016), 530–547.
- [91] Anselm L Strauss. 1987. *Qualitative analysis for social scientists*. Cambridge university press.
- [92] Trang P Tran. 2017. Personalized ads on Facebook: An effective marketing tool for online marketers. *Journal of retailing and consumer services* 39 (2017), 230–242.
- [93] Indrit Troshani, Sally Rao Hill, Claire Sherman, and Damien Arthur. 2021. Do we trust in AI? Role of anthropomorphism and intelligence. *Journal of Computer Information Systems* 61, 5 (2021), 481–491.
- [94] Wan-Hsiu Sunny Tsai and Ching-Hua Chuan. 2023. Humanizing chatbots for interactive marketing. *The Palgrave handbook of interactive marketing* (2023), 255–273.
- [95] Blase Ur, Pedro Giovanni Leon, Lorrie Faith Cranor, Richard Shay, and Yang Wang. 2012. Smart, useful, scary, creepy: perceptions of online behavioral advertising. In *proceedings of the eighth symposium on usable privacy and security*. 1–15.
- [96] Thiemo Wambsganss, Naim Zierau, Matthias Söllner, Tanja Käser, Kenneth R Koedinger, and Jan Marco Leimeister. 2022. Designing conversational evaluation tools: A comparison of text and voice modalities to improve response quality in course evaluations. *Proceedings of the ACM on Human-Computer Interaction* 6, CSCW2 (2022), 1–27.
- [97] Jared Watson, Francesca Valsesia, and Shoshana Segal. 2024. Assessing AI receptivity through a persuasion knowledge lens. *Current Opinion in Psychology* (2024), 101834.
- [98] Joel Wester, Henning Pohl, Simo Hosio, and Niels van Berkel. 2024. “This Chatbot Would Never...”: Perceived Moral Agency of Mental Health Chatbots. *Proceedings of the ACM on Human-Computer Interaction* 8, CSCW1 (2024), 1–28.
- [99] Rebekah C White and Anna Remington. 2019. Object personification in autism: This paper will be very sad if you don’t read it. *Autism* 23, 4 (2019), 1042–1045.
- [100] Rua M Williams and LouAnne E Boyd. 2019. Prefigurative politics and passionate witnessing. In *Proceedings of the 21st International ACM SIGACCESS Conference on Computers and Accessibility*. 262–266.
- [101] Rua M Williams and Juan E Gilbert. 2019. “Nothing About Us Without Us” Transforming Participatory Research and Ethics in Human Systems Engineering. In *Advancing diversity, inclusion, and social justice through human systems engineering*. CRC Press, 113–134.
- [102] Cara Wilson, Laurianne Sitbon, Bernd Ploderer, Jeremy Opie, and Margot Brereton. 2020. Self-expression by design: Co-designing the expressiball with minimally-verbal children on the autism spectrum. In *Proceedings of the 2020 CHI conference on human factors in computing systems*. 1–13.
- [103] Heather M Wojton, Daniel Porter, Stephanie T. Lane, Chad Bieber, and Poornima Madhavan. 2020. Initial validation of the trust of automated systems test (TOAST). *The Journal of social psychology* 160, 6 (2020), 735–750.
- [104] Yuxi Wu, Sydney Bice, W Keith Edwards, and Sauvik Das. 2023. The slow violence of surveillance capitalism: How online behavioral advertising harms people. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*. 1826–1837.
- [105] Kun Xu, Xiaobei Chen, and Luling Huang. 2022. Deep mind in social responses to technologies: A new approach to explaining the computers are social actors phenomena. *Computers in Human Behavior* 134 (2022), 107321.
- [106] Rui Zhang, Wen Duan, Christopher Flathmann, Nathan McNeese, Guo Freeman, and Alyssa Williams. 2023. Investigating AI teammate communication strategies and their impact in human-AI teams for effective teamwork. *Proceedings of the ACM on Human-Computer Interaction* 7, CSCW2 (2023), 1–31.

A Within-Group Differences

Table 5 shows the change in ranking for each group from their initial ranking to post ranking.

		n	Mean		Median		Std. Dev.		Diff. Test
			Pre	Post	Pre	Post	Pre	Post	
autism _{substantial}	Initially Most Favorite	27	1	1.39	1	1	0	0.77	-3.42***
	Initially Least Favorite	27	5	4.36	5	4.6	0	0.91	4.45***
autism _{general}	Initially Most Favorite	30	1	1.08	1	1	0	0.17	-2.64*
	Initially Least Favorite	30	5	4.9	5	5	0	0.16	3.15**
general	Initially Most Favorite	30	1	1.07	1	1	0	0.33	-1.41
	Initially Least Favorite	30	5	4.8	5	4.8	0	0.23	4.15***

Table 5. The change in rating of most and least favorite items before and after chatbot interaction; Pre=pre chatbot ranking; Post=post chatbot ranking; * p-values <= 0.05; ** <= 0.01; *** <= 0.001; Means, medians, and standard deviations are all averaged across the 5 consumer product types

B Consumer Items

- Candy Bar
 - Hershey’s
 - Reese’s
 - Snicker’s
 - Kit Kat
 - Twix
- Activewear
 - Nike
 - Lululemon
 - Adidas
 - Under Armour
 - Athleta
- Fast Food
 - McDonalds
 - Chick-Fil-A
 - Taco Bell
 - Burger King
 - Wendy’s
- Soda
 - Coke

- Diet Coke
- Pepsi
- Dr. Pepper
- Mountain Dew
- Music Genre
 - Pop
 - Rock
 - Hip-Hop/Rap
 - Electronic Dance Music (EDM)
 - Country

Received October 2024; revised April 2025; accepted August 2025