



Raise the Curtains: The Effect of Awareness About Targeting on Consumer Attitudes and Purchase Intentions

Sonam Samat, Alessandro Acquisti, and Linda Babcock, *Carnegie Mellon University*

<https://www.usenix.org/conference/soups2017/technical-sessions/presentation/samat-awareness>

**This paper is included in the Proceedings of the
Thirteenth Symposium on Usable Privacy and Security (SOUPS 2017).**

July 12–14, 2017 • Santa Clara, CA, USA

ISBN 978-1-931971-39-3

**Open access to the Proceedings of the
Thirteenth Symposium
on Usable Privacy and Security
is sponsored by USENIX.**

Raise the Curtains: The Effect of Awareness About Targeting on Consumer Attitudes and Purchase Intentions

Sonam Samat

Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA 15232
sonamsam@andrew.cmu.edu

Alessandro Acquisti

Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA 15232
acquisti@andrew.cmu.edu

Linda Babcock

Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA 15232
lb2k@andrew.cmu.edu

ABSTRACT

We investigate the effect of awareness about targeting on users' attitudes towards a targeted ad and behavioral intentions towards the advertised product. Specifically, we study the effect of a notice that makes individuals aware that a particular advertisement has been targeted to them on their attitudes about the product and intentions to purchase the product. We find that, among individuals who have negative opinions about the practice of targeted advertising, awareness about targeting significantly worsens attitudes towards the targeted product and reduces the likelihood of purchasing the targeted product. Among individuals who have positive and neutral opinions about targeted advertising, awareness about targeting does not impact attitudes or purchase intentions towards the targeted product. We develop a scale to measure opinions about targeted ads and find that a substantial proportion (at least 33%) of our participants have negative opinions about targeted ads. This suggests that the self-regulated advertising industry is not incentivized to follow recommendations from the U.S. Federal Trade Commission to make consumers aware about their targeted advertising practices.

1. INTRODUCTION

Behavioral or targeted advertising is defined as “the practice of tracking an individual’s online activities in order to deliver advertising tailored to the individual’s interests” [16]. These online activities include webpages visited and content viewed, search queries, and online purchases. Many of the steps involved in performing targeted advertising (data collection and aggregation, data mining, bidding for ad spaces, etc.) are hidden from consumers. Consumers are typically only asked for overall consent for data collection when they register with an online service. After this initial consent, there are rarely any reminders about the continuous practice of online data collection. Consumers are also not made aware of the aggregation of their data from multiple sources, the mining of their data to select ads shown to them, or the real-time bidding process used to sell ad spaces. There is no easy way for consumers to distinguish between targeted ads and non-targeted ads, or to figure out what information was used in the

targeting process. Therefore, as far as consumers are concerned, targeted advertising happens “behind the scenes.” We investigate the effect of “raising the curtains,” by making individuals aware that a particular ad is targeted to them, on their attitudes and purchase intentions towards the advertised product.

The industry favors the use of targeted advertising because it generates higher click-through rates [14] and higher sales [2] in comparison to non-targeted ads. But consumer surveys about perceptions of targeted advertising suggest that, by and large, people do not like being tracked and do not wish to receive behaviorally targeted advertisements [29, 34, 39]. These concerns are not unfounded given the growing privacy risks associated with large-scale data collection and use, such as the use of consumer data for price discrimination [28, 40] and in revealing embarrassing personal details [20]. In order to address these privacy risks, the U.S. Federal Trade Commission (FTC) has laid out a number of recommendations for best practices, one of which focuses on being transparent about how consumer information is collected and used, so that consumers are not kept in the dark and they can make informed decisions about their online activities [15, 16, 17].

The ad industry has made some efforts to achieve this goal. In 2011, the Digital Advertising Alliance (DAA), which is a coalition of advertising, media, marketing and technology companies, developed a set of icons (Figure 1) that may be displayed on targeted advertisements delivered by its members [9, 11]. The goal of these icons is to communicate how behavioral targeting works and provide consumers with avenues for opting out. However, Leon et al.’s 2012 work finds that people severely misunderstand these icons: 53% of their participants incorrectly believed that more ads would pop up if they clicked on the icon and 45% incorrectly believed that the accompanying ‘AdChoices’ tagline was intended to sell advertising space [22].



Figure 1. Icons used by ad industry on targeted ads.

We are interested in understanding how awareness about targeted advertising impacts consumers’ attitudes and purchase intentions towards the advertised product. Previous research provides contradictory evidence about this effect. Research from the recommendations systems literature suggests that providing explanations for how recommendations are selected increases users’ trust in the recommendation system and their likelihood to

Copyright is held by the author/owner. Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee.

Symposium on Usable Privacy and Security (SOUPS) 2017, July 12 – 14, 2017, Santa Clara, California.

use the system in the future [8, 37]. Awareness about targeting could work in a similar way if it helps users understand how advertisements are selected for them and if it increases users' purchase intentions towards the advertised product. On the other hand, highly personalized advertisements can be perceived as intrusive or creepy, prompting individuals to avoid the advertisement [12, 26, 41]. Awareness about targeting may elicit feelings of intrusiveness because it reminds people about the continuous tracking of their data, and may therefore decrease purchase intentions towards the advertised product. By understanding how awareness about targeting impacts consumers' attitudes and purchase intentions, we aim to inform policy makers about the incentives of a self-regulated ad industry in implementing the FTC's recommendation about being transparent when targeting.

To this end, we conducted a series of studies to understand the impact of awareness about targeting on consumers' attitudes and purchase intentions towards the advertised product. We implemented awareness through the use of a text message shown along with the ad that indicated that the displayed ad had been selected for the user based on some information about the user. We hypothesize that the effect of awareness about targeting on attitudes and purchase intentions depends on individuals' overall opinions about the practice of targeted advertising. To test this hypothesis, we built and validated a scale that measures individuals' overall opinions about targeted advertising, as existing literature has not established a scale to measure this construct. We find that, for the participants who have negative opinions about targeted ads, awareness about targeting worsens attitudes and reduces purchase intentions for the advertised product. We also find that a substantial portion of our participants have negative opinions about targeted ads (53% in our first study and 33% in our third study). This suggests that a self-regulated ad industry is not incentivized to use notices that make consumers aware of their targeting practices, as it can significantly reduce their revenues from advertising.

2. RELATED WORK

2.1.1 Recommendation Systems

Recommendation systems (or recommendation agents) like Netflix or Amazon share several similarities with the practice of targeted advertising. Recommendation systems also collect data about consumers' preferences and behaviors, and use this data to recommend products to consumers. Research from this field shows that explaining how a recommendation was selected generally has a positive effect on users' attitudes. For instance, Chen and Pu's 2005 work shows that explanations can be an effective way to increase users' trust in the recommendation system [8]. In their 2002 work, Sinha and Swearingen asked participants to report how much they like the recommendations made by music recommender systems as well as how transparent they perceive these systems to be. They found that users like the recommendations made by systems they perceive to be transparent more than those they perceive to be non-transparent [37]. Bilgic and Mooney's 2005 work further shows that some explanations can even improve users' own accuracy in predicting how much they will like the recommended item [4]. In addition, 86% of the participants in a survey conducted by Herlocker et al. in 1999 said they would like to see explanations for the choices made by the recommendation system used in their experiment [19]. Johnson and Johnson's 1993 work attempted to explain this positive effect by highlighting that explanations provide an association between antecedent and consequent, that is, a link between cause and effect [21]. Notices

that make consumers aware of the fact that an advertisement is targeted to them may also provide an explanation for why the particular advertisement was selected for them and, similar to recommendation systems, could improve consumers' attitudes about the advertised product.

There are some key differences between recommendation systems and behavioral targeting. Users can actively subscribe to recommendation systems, whereas they are typically subjected to targeted ads without their explicit consent. Therefore, the extent to which individuals expect (or even desire) personalized recommendations may be different for recommendation systems and behavioral targeting. While users expect recommendations for movies on Netflix and consumer goods on Amazon, they may not expect information from their Google search queries to be used in targeted advertisements on the New York Times website. As a result, the positive effect of explanations on users' intentions may be restricted to recommendations made within a specific platform and may not carry forward to behavioral targeting of advertisements that we study in this paper.

2.1.2 Targeted Advertising

Awareness about targeting might positively influence attitudes towards the targeted product due to the placebo effect, by which simply telling someone a product has been selected for them can positively influence their opinions about the product. Cosley et al.'s 2003 work shows that users' ratings of a recommendation can be manipulated simply by showing them what the recommendation system predicts their rating of the product will be, irrespective of whether the prediction is accurate or not [10]. In another example, a recent study conducted by OkCupid.com (an online dating platform) shows that the effect of simply telling people that a particular individual is a 90% match for them, when truly the individual is only a 30% match, is just as strong as when the individual is actually a 90% match [35]. In other words, the mere myth of compatibility works just as well as actual compatibility. This may happen if consumers trust that sophisticated algorithms with large amounts of data will make the best selections for them. In our context, a placebo effect may lead to improved attitudes and higher purchase intentions toward the advertised product if awareness notices convey the fact that an ad has been *selected* or *recommended* for the user.

Additionally, consumers may react to what a targeted ad signals about themselves. For instance, upon viewing a targeted ad for an environmentally friendly product, a consumer might believe that the advertiser thinks of her as someone who is environmentally friendly (i.e., the targeted ad can act as an implied social label). She may then adjust her own perceptions about being environmentally friendly and in turn also change her behavior towards environmentally friendly products. In their 2016 work, Summers et al. demonstrate that targeted advertisements can act as social labels causing consumers to adjust their own self-perceptions, and even behaviors, to match the implied labels. While they do not directly test if making individuals aware of targeting impacts purchase intentions, they explore whether awareness notices act as implied social labels and cause adjustments to self-perceptions, which in turn influence purchase intentions towards the advertised product. [38]. We borrow from the study design used by Summers et al. in their 2016 paper and extend their work by testing the direct impact of awareness notices on attitudes and purchase intentions towards the advertised product. More importantly, we test whether this

impact varies with the *a priori* opinions that individuals may have about targeted advertising.

Consumer surveys on perceptions toward behavioral targeting help provide support for the opposite hypothesis that awareness about targeting can negatively influence attitudes and purchase intentions. In a nationally representative survey conducted by Turov et al. in 2009, 66% of the participants claimed that they did not want behaviorally targeted advertisements [39]. In 2010, McDonald and Cranor found that the majority of their participants (55%) also did not wish to receive advertisements tailored to them [29]. A 2012 survey conducted by the Pew Research Center found that 68% of adult Internet users say that they are not okay with targeted advertising because they “don’t like having [their] online behavior tracked and analyzed” [34]. These data suggest that attitudes and purchase intentions towards the targeted ad may be negatively affected if awareness notices inherently remind consumers about the continuous tracking of their personal information.

In their 2012 work, Leon et al. show that users struggle to use the existing tools for opting out of targeted advertising, including tools that block access to advertising websites, tools that set cookies indicating a user’s preference to opt out of targeted advertising, and privacy tools that are built directly into web browsers [23]. If users experience lack of control over their data then awareness notices indicating which ads are specifically targeted to them may cause users to retaliate by specifically avoiding the products shown in these targeted ads. This retaliation is called ‘reactance’ [6], and has been shown to occur when ads are perceived to be highly intrusive [12] and when messages are perceived to be highly personalized [41].

3. Hypotheses

As highlighted in the previous section, existing literature does not provide a clear indication of how awareness about targeting may impact individuals’ purchase intentions. We believe that consumers’ opinions about the practice of targeted advertising will play a moderating role in determining how awareness about targeting impacts attitudes and purchase intentions towards the targeted product. Specifically, for consumers who have a positive opinion about targeted advertising in general, awareness about targeting will increase purchase intentions towards the advertised product relative to no awareness about targeting. On the other hand, awareness about targeting will decrease purchase intentions relative to no awareness about targeting for consumers who have a negative opinion about targeted advertising. These hypotheses are developed based on three factors that determine attitudes towards the product: 1) how much the person likes that product in the absence of targeting, 2) how useful the person finds having a product recommended to her based on her personal information (in other words, how much the person likes targeted advertising), and 3) how invasive the person finds targeted advertising (in other words, how much the person dislikes targeted advertising). When the individual is not aware that an advertisement is targeted, the second and third factors are not activated and only the first factor determines the individual’s attitudes towards the product. However, when the individual is made aware of targeting, the second and third factors are activated and a combination of all three factors determines her attitudes towards the product. Holding the first factor constant, the relative strength of the second and third factors determines the influence of awareness about targeting on purchase intentions.

In order to test this moderating role, we need to effectively measure opinions about targeting practices. Existing literature has not established a scale to measure opinions about targeting practices. Several researchers have used ad-hoc measures for this construct, but none have attempted to build and validate a scale to measure it. For instance, Leon et al.’s 2012 work included four questions towards the end of their survey to measure participants’ opinions towards behavioral advertising. These questions were: “How comfortable are you with behavioral advertising?”, “In general, I find behavioral advertising useful”, “In general, I like behavioral advertising” and “Behavioral advertising is privacy invasive” [22]. In 2016, Melicher et al. conducted in-person interviews with 35 participants asking them questions to capture their opinions about online tracking. They identified four groups of opinions (generally negative, generally neutral, mixed, and conditionally positive) but they did not attempt to build a scale to measure opinions in a closed-ended format [30]. In 2010, McDonald and Cranor interviewed 14 participants, and used their responses to create two closed-ended measures: one for feelings towards current targeting practices and another for reasons to accept or reject targeted advertising. They later used these measures in an online survey but they did not perform any formal validation analysis for their measures [29].

4. OUR CONTRIBUTION

To test our hypotheses, we conducted three studies. The first was an exploratory study to measure participants’ opinions about targeted advertising using open-ended questions. The goal of this study was to capture the different kinds of opinions that participants express about targeting practices in their own words, and use them to build a scale that can measure opinions about targeting in future studies. The second study was a validation study in which we validated the closed-ended scale created from the responses received in the exploratory study. We also shortened the scale so participants can complete it quickly. We then tested convergent and divergent validity, as well as the test-retest validity of our scale. Finally, in the third study, we evaluated the effect of awareness about targeting on attitudes and purchase intentions towards the advertised product. In this study, we tested the hypothesis that opinions about targeting moderate the effect of awareness of targeting on attitudes and purchase intentions.

We recruited participants for all studies from Amazon Mechanical Turk. The platform allows researchers to approve or reject participants’ payment based on their performance. Therefore, each participant has an approval rating (the percentage of his or her previously completed surveys or tasks that have been approved). Following the recommendations established by previous researchers, we implemented a minimum requirement of a 95% approval rating during our recruitment process [33]. All our studies were approved by the Institutional Review Board at Carnegie Mellon University and participants in all studies consented to participate in this research.

4.1 Exploratory Study

4.1.1 Methods

We believe consumers’ opinions about targeting practices will play a moderating role in the effect of awareness on purchase intentions. The primary goal of this study was to identify the different types of opinions (for example: positive, negative, conditionally positive, etc.) that individuals have towards targeted advertising practices, and build a scale to measure these opinions reliably and quickly. From this study, we also gained insights into the prevalence of

different opinions, the reasons behind these opinions, and the factors that correlate with these opinions.

We conducted an online survey on Amazon’s Mechanical Turk platform. The survey lasted 15–20 minutes and participants were paid \$1 as compensation. Participants were shown a hypothetical scenario about a targeted advertising experience. Targeting can be done in several different ways and individuals’ opinions may vary across different types of targeting. In this study, we measured opinions about three types of targeting: 1) Remarketing: where the targeted ad is for a specific product that the individual has looked at before, 2) Interest-based targeting: where the targeted ad is for a product that is similar to other products that the individual has looked at before, and 3) Contextual targeting: where the targeted ad is related to the content of the website where the ad is shown. By random assignment, each participant in our study was shown one of three hypothetical scenarios involving one of the three types of ads (see Appendix A.1). Once participants read the scenario, they were asked open-ended questions to capture their opinions, and the reasons behind those opinions, about the type of targeting practice in their scenario. Specifically, we asked participants the following questions: Q1: “How would you feel if you experienced this scenario?” and Q2: “Please tell us why you would feel this way if you experienced this scenario.” Finally, we measured individual-level factors such as the IUIPC scale for privacy concerns [27], previous experiences with online privacy invasions involving targeted ads and in general (self-developed), previous actions taken to avoid targeted ads [22], perceived control over this type of targeted advertising (self-developed), level of Internet usage [24], current interest in buying the advertised product (self-developed), the Domain-Specific Risk-Taking (DOSPERT) scale [5], Internet usage levels, and demographics. The self-developed measures are reported in Appendix A.2.

4.1.2 Results

One hundred eighteen participants (Mean Age = 31.4; 58% Male) completed this study. Two independent coders read and coded participants’ open-ended responses about their opinions. After they completed the first round of coding independently, they met and discussed their codes to consolidate them and form the final set of codes. Then, they coded the open-ended responses once again using this final set of codes. The inter-rater reliability for the final codes was satisfactory for both questions (Cohen’s kappa=0.71 for Q1 and Cohen’s kappa=0.77 for Q2). The cases where the coders differed in their final codes (25% in Q1 and 19% in Q2) were resolved by the first author, by reading the participants’ responses and selecting the code that seemed a better fit from among the two codes assigned by the coders.

4.1.2.1 Categories of Opinions

Eight different categories of opinions about targeted advertising emerged from participants’ open-ended responses to Q1. Figure 2 shows the distribution of these opinions across all three conditions. As can be seen from this figure, a large proportion of our participants (37%) reported that they would feel neutral if they experienced the targeted advertising scenario shown to them. Some responses in this category were “I’d feel fine, that’s pretty normal,”

“I would have neutral feelings,” and “I would not feel any special way.” A small proportion of our participants (6%) reported feeling positively about the targeted advertising scenario shown to them, and an even smaller proportion (4%) reported feeling mixed emotions. Some responses from the positive category were “[I] would feel interested” and “I would feel excited to have this new [recommendation] and I wouldn’t really care that my browsing history was tracked,” and a response we observed in the mixed category was “I would have mixed emotions. I’d feel a bit weird about the tracking thing, but I’d also be interested in checking out the shoes, probably.”




Figure 2. Distribution of participants’ open-ended responses to Question 1.¹

We observed a wide range of negative opinions towards these types of targeted ads. The spectrum of responses ranged from “slightly annoyed” to “violated,” with three categories in between these two extremes. The total proportion of participants that reported negative feelings is 53%. Some examples of responses from the modal ‘Dislike’ category were “I would feel like my personal information is not safe. It would make me feel uncomfortable,” “I would feel watched and unhappy,” and “I would feel like my privacy has been invaded.”




Figure 3. Distribution of participants’ open-ended responses to Question 1, split by targeting type condition.

Next, we looked at the differences between the three conditions (contextual, interest-based, and remarketed). For simplicity, we combined the five categories of negative responses. As Figure 3 shows, the contextual ad condition had more neutral responses and

¹ Seven participants’ responses to Q1 were coded as unusable either because the response did not make any sense or because the

response did not answer the question in anyway. The percentages and graphs shown here are for the remaining 111 participants.

fewer negative responses than the interest-based and remarketed conditions. Further statistical tests confirmed that the distribution of responses (using the consolidated 4 categories) did not significantly differ from each other in the interest-based and remarketed conditions ($\chi^2(3) = 2.66, p=0.44$). However, the distribution of responses in the contextual condition was significantly different from the interest-based ($\chi^2(3) = 8.86, p=0.031$) and remarketing conditions ($\chi^2(3) = 9.40, p=0.024$).

4.1.2.2 Reasons Behind Opinions

Nine different categories emerged from participants' open-ended responses to Q2 about why they would feel the way they reported feeling in Q1. As shown in Figure 4, the most common reason we observed (provided by 29% of our participants) was "because it happens all the time," suggesting that individuals have become used to seeing targeted ads. Some responses we observed in this category were "I have had it happen before many times, so it has become normal," "I would feel this way because it is a scenario that I have experienced in the past, and currently experience," and "I would feel this way because it is not the first time this has happened. It's common practice in my opinion." Not surprisingly, most of these participants (87%) reported feeling neutral in Q1.




Figure 4. Distribution of participants' open-ended responses to Question 2.²

The second most common reason we observed (provided by 23% of our participants) was "because it feels like an invasion of my privacy." Many participants in this category used the exact words "because it feels like an invasion of my privacy" as part of their response, while some others reported "This would make me [feel] that [it's] not safe for me to look at anything online because it [feels] as if someone is watching me" and "they made me feel as if [I] was being watched." Almost all participants in this category (96%) reported feeling one of the negative opinions in Q1.

The third most common category we observed (provided by 12% of our participants) was "because I feel they are exploiting me for money" which includes responses where participants specifically said they feel they are being taken advantage of, or being manipulated for money, in such scenarios. Interestingly, 62% of participants who reported this reason were in the contextual targeting condition. One participant in the contextual targeting condition reported, "The advertiser pays them, so the newspaper feels obligated to place the ad somewhere that readers can feel

[encouraged] to buy this particular shoe. They are trying to lead unsuspecting readers to this particular store. It's consumer manipulation." Other categories we observed were "because I find such ads useful" or "because I understand why they are needed" (9%), "because it makes me suspicious" (9%), "because I did not provide permission" (7%, labeled 'lack of consent' in Figure 16), "because I don't care about targeting" (6%), "because I do not know how to turn it off" (4%, labeled 'lack of control' in Figure 17), and "because I don't notice ads" (2%).

4.1.2.3 Correlations with Opinions

Next, we explored how perceived control over targeted advertising and previous experiences with online privacy invasions correlate with opinions. Before doing this analysis, we noted that the biggest category of responses is negative and that the other categories are relatively small, suggesting that consolidating categories might be necessary. Accordingly, we created two consolidated categories for 'negative' vs. 'non-negative' opinions, combining the neutral, mixed, and positive categories into one. We ran logistic regressions with the consolidated two-category dummy variable for negative opinions as our dependent variable (coded as '0' for non-negative opinions and '1' for negative opinions) and the individual-level factors measured in our survey as the independent variables. This analysis suggests that, for the interest-based and remarketing conditions, the level of perceived control over targeting significantly predicts participants' opinions about targeting (odds ratio=0.30, $p=0.002$). Participants in the interest-based and remarketing conditions who feel more control over this type of targeting are less likely to have negative opinions about it. In the contextual targeting condition, a previous unpleasant experience with targeted ads marginally predicts opinions, such that participants who report having a previous unpleasant experience with targeted ads are more likely to have negative opinions about targeting (odds ratio = 7.7, marginally significant $p = 0.054$). The full results of our logistics regressions are shown in Appendix B.1. We validate these relationships with a larger sample size in Validation Study A.

4.1.3 Discussion

Our results suggest that many participants (53%) feel negatively about targeted advertising, especially interest-based and remarketed advertising. It may be that the collection and use of browsing history (which happens in the interest-based and remarketing scenarios but not in the contextual targeting condition) plays a key role in explaining how participants feel about targeting practices.

In this study, we also learn that a substantial number of participants (37%) have neutral opinions about targeting practices. Given this result, we re-evaluate our original hypothesis about the moderating role of opinions about targeting in the effect of awareness on purchase intentions. What would be the effect of awareness about targeting on purchase intentions for individuals who feel neutral about targeting practices? As 68% of participants who reported having neutral opinions also reported feeling this way because targeting "happens all the time," we believe there will be no effect of awareness notices on purchase intentions for these individuals.

² Seven participants' responses to Q2 were coded as unusable either because the response did not make any sense or because the

response did not answer the question in any way. The percentages and graphs shown here are for the remaining 111 participants.

These individuals expect such targeting to occur anyway, so they will obtain no new information from awareness messages. We find that a small proportion of individuals reported positive or mixed opinions about targeted ads. While the effect of awareness about targeting may be different for these individuals, it is difficult to test such an effect if very few people report having such opinions.

In the next studies, we narrow our focus to develop a scale to measure opinions about one type of targeted advertising because different types of targeting may require different measurement models. We choose to focus on interest-based targeting because we believe it is more difficult for individuals to identify interest-based ads in comparison to remarketed and contextually targeted ads, since remarketed ads advertise the same product that an individual has looked at before and contextual ads advertise a product related to the content of the page where the ad is shown. In comparison, interest-based ads display products similar, though not identical, to ones that an individual has looked at before and as such the consumer may not be aware of the targeting.

4.2 Validation Study A

4.2.1 Methods

In order to measure opinions about the practice of targeted advertising without having to collect and code open-ended responses, we built a scale starting with the 14 most common responses provided by our participants in the first question in the Exploratory Study (Table 1). In this study, we recruited participants from Amazon's Mechanical Turk platform for a study advertised to collect individuals' opinions about a hypothetical online scenario.

In order to avoid selection bias, our Mechanical Turk post did not mention anything about targeted ads (or ads in general). The survey lasted 5 minutes and participants were paid \$0.30 as compensation. Participants were shown the interest-based targeting scenario used in the Exploratory Study. After participants read the scenario, they were told, "We are interested in understanding how you would feel about the type of targeted advertising described in the scenario on the previous page. Please indicate how strongly you agree or disagree with the following statements." For each statement, participants were asked to indicate their level of agreement on a 1–7 scale from 'Strongly Agree' to 'Strongly Disagree'. The order of the 14 statements was randomized.

We asked participants their perceived level of control over this type of targeted advertising (self-developed) and their previous experiences with online privacy invasions either involving targeted ads or in general (self-developed). Then, we asked questions about their current interest in purchasing shoes (the product used in our ad), how often they purchase shoes online and offline, their general Internet usage [25], actions they have taken to avoid being tracked online [22], their Internet privacy concerns [27], and demographics.

4.2.2 Results

Two hundred ninety-six participants (Mean Age = 34.3; 53% Male) completed this study. Our 14-item scale shows high reliability (Cronbach's alpha = 0.9265, well above the acceptable threshold of 0.70 [32]). As can be seen in Table 1, eliminating the two mixed category items results in a higher Cronbach's alpha value of 0.9554. Therefore, we decided to drop the two mixed items from our scale.

Table 1. Correlations, covariance, and Cronbach's alpha values for the 14-item scale. *Cronbach's alpha value only increases when the two mixed category items are deleted.

Item	Category	Item-test correlation	Item-rest correlation	Average inter-item covariance	Alpha if item is deleted
I would feel ok about this type of advertising	Neutral	0.90	0.88	1.49	0.9141
I would feel indifferent about this type of advertising	Neutral	0.63	0.56	1.60	0.9246
I would not be bothered by this type of advertising	Neutral	0.83	0.80	1.51	0.9168
I would think this is clever advertising	Positive	0.68	0.62	1.58	0.9228
I would be interested in products shown in this type of advertising	Positive	0.76	0.72	1.55	0.9195
I would feel this type of advertising can be helpful to me	Positive	0.81	0.78	1.54	0.9178
I would have mixed feelings (good and bad) about this type of advertising	Mixed	0.08	-0.03	1.83	0.9428*
This type of advertising would make me feel worried but I can also see the benefit to me	Mixed	0.14	0.05	1.80	0.9392*
I would be annoyed by this type of advertising	Negative	0.87	0.84	1.50	0.9154
I would feel creeped out by this type of advertising	Negative	0.86	0.83	1.48	0.9155
I would not like this type of advertising	Negative	0.88	0.86	1.49	0.9148
This type of advertising would make me feel like my privacy has been invaded	Negative	0.86	0.83	1.50	0.9157
I would feel upset about this type of advertising	Negative	0.83	0.79	1.51	0.9170
This type of advertising would make me feel violated and manipulated	Negative	0.87	0.85	1.49	0.9152
Overall				1.56	0.9265

Table 2. The 6 items in our final scale, along with their inter-item correlations. *Significant at the 0.001 alpha level.

	Item	Category	1	2	3	4	5	6
1	I would feel ok about this type of advertising	Neutral	1					
2	I would not be bothered by this type of advertising	Neutral	0.76*	1				
3	I would be interested in products shown in this type of advertising	Positive	0.70*	0.62*	1			
4	I would feel this type of advertising can be helpful to me	Positive	0.74*	0.65*	0.76*	1		
5	I would be annoyed by this type of advertising	Negative	-0.79*	-0.71*	-0.63*	-0.65*	1	
6	I would not like this type of advertising	Negative	-0.79*	-0.71*	-0.64*	-0.71*	0.78*	1

4.2.2.1 Scale Consolidation

Our next goal was to reduce the size of the scale so participants can complete it in a shorter period of time, while still maintaining the high reliability of the scale. We decided to reduce the scale to 6 items, and consequently computed the Cronbach's alpha value for all combinations of 6 items from the remaining 12 items. We identified the best combination as the one that equally represents the remaining three categories of opinions (neutral, positive, and negative) and has a high Cronbach's alpha value of 0.9362. Table 2 shows the 6 items that make our final scale. All pairwise correlations between the 6 items are statistically significant at the 0.001 alpha level. In addition to reducing the time burden on participants, this shorter scale also ensures that it is not biased towards any single category as it includes an equal number of items (two) from each category.

Table 3. Factor loadings from the exploratory factor analysis (Validation Study A) and confirmatory factor analysis.

Item	Factor Loading (EFA)	Completely Standardized Parameter (CFA)
I would feel ok about this type of advertising	0.9021	0.9378
I would not be bothered by this type of advertising	0.8115	0.8183
I would be interested in products shown in this type of advertising	0.7890	0.8513
I would feel this type of advertising can be helpful to me	0.8311	0.8957
I would be annoyed by this type of advertising	-0.8456	-0.8528
I would not like this type of advertising	-0.8640	-0.8406

4.2.2.2 Exploratory Factor Analysis

Exploratory factor analysis with our 6-item scale yielded a single factor with eigenvalue greater than one. As shown in Table 3, each of the 6 items is highly correlated with the single underlying factor. This suggests that our scale measures a single underlying construct: opinions towards targeted advertising. Positive and neutral items have positive factor loadings while the negative items have negative factor loadings. Therefore, our scale can be consolidated into a single value by reverse coding the negative items ($8 - \text{value}$) and then averaging the 6 items. Lower numbers of the consolidated value indicate more negative opinions towards targeted advertising. We use this consolidated value to validate the preliminary results obtained in the Exploratory Study. Just as in the Exploratory Study, we find that the level of perceived control over targeting significantly predicts participants' opinions about targeting ($\beta_{\text{control}} = 0.140, p=0.011$). We also find that having a previous unpleasant or uncomfortable experience with targeted ads significantly predicts participants' opinions about targeted ads ($\beta_{\text{prevtar}} = -1.16, p<0.001$). The regression coefficients are reported in Appendix B.2.

4.3 Validation Study B

4.3.1 Methods

The goal of this study was three-fold: 1) to confirm the single-factor model for our 6-item scale that emerged in Validation Study A, 2) to test the divergent validity of our 6-item scale, and 3) to evaluate the test-retest reliability of our 6-item scale. Towards the first goal, we recruited participants from Amazon Mechanical Turk in the same way as in Validation Study A (participants who completed Validation Study A were not allowed to take this survey). The survey lasted 10–15 minutes and participants were paid \$0.75 as compensation. Participants were shown the hypothetical scenario about interest-based targeting and asked to indicate their level of agreement with the 6 items (presented in random order) on a 1–7 scale from 'Strongly Agree' to 'Strongly Disagree'.

Towards the second goal for this study, we included three existing scales in our survey: 1) the IUIPC privacy concerns scale [27], which measures Internet users' information privacy concerns, 2) the Domain-Specific Risk-Taking (DOSPERT) scale [5], which measures individuals' attitudes towards engaging in risky behaviors, and 3) the General Decision Making Style (GDMS) scale [36], which measures decision-making styles. Recent work

Table 4. Scale descriptive statistics, reliabilities, and correlations with our 6-item scale. The three sub-scales of the IUIPC scale show significant (at the 0.001 alpha level) and moderate correlations with our 6-item scale.

Scale	Sub-scale	Mean	Standard deviation	Cronbach's alpha	Correlation with our 6-item scale
IUIPC	Control	5.68	1.08	0.7903	-0.36*
	Awareness	6.03	1.06	0.8654	-0.25*
	Collection	5.44	1.34	0.9095	-0.47*
DOSPERT	Financial	2.72	1.36	0.8577	0.07
	Health/Safety	2.83	1.22	0.7261	0.08
	Recreational	2.88	1.45	0.8392	0.01
	Ethical	2.24	1.10	0.8097	0.05
	Social	4.81	1.18	0.7779	-0.09
GDMS	Rational	3.97	0.66	0.8330	-0.07
	Avoidant	2.43	1.06	0.9227	0.01
	Dependent	3.07	0.84	0.8329	0.06
	Intuitive	3.09	0.87	0.8550	0.09
	Spontaneous	2.44	0.90	0.8620	0.11

has shown that these three established scales are good predictors of privacy preferences and behaviors [13]. Towards the third goal, all participants were invited to take a follow-up survey after a two-week gap, in which they were shown the same scenario and asked to indicate the extent to which they agreed with the same 6 items.

4.3.2 Results

Two hundred ninety-four participants (Mean Age = 33.8; 53% Male) completed this study. Again, our 6-item scale shows high reliability (Cronbach's alpha = 0.9385) and the elimination of no item yields a higher Cronbach's alpha (α -if-item-deleted_i < 0.9385, for all items).

4.3.2.1 Confirmatory Factor Analysis

Confirmatory factor analysis suggests that the data fit a single-factor model well, with a Bentler's comparative fit index (CFI) of 0.955, a Tucker-Lewis index (TLI) of 0.925, and a standardized root mean squared residual (SRMR) of 0.027. The confirmatory factor loadings are shown in Table 3 and are all significant at the 0.001 alpha level with t-statistics greater than 29.5.

4.3.2.2 Divergent Validity

Next, we tested the divergent validity of our 6-item scale. We compared our scale to three established scales: IUIPC, DOSPERT, and GDMS. Table 4 shows the mean, standard deviations, Cronbach's alpha coefficients, and correlations with our 6-item scale for each subscale of the three established scales. The IUIPC and DOSPERT scales use a 1–7 response scale while the GDMS scale uses a 1–5 response scale. As shown in Table 5, only the sub-

scales of the IUIPC scale are significantly correlated with our 6-item scale (at the 0.001 alpha level). Correlations of our 6-item scale with the sub-scales of the DOSPERT and GDMS scales are all not significant. The significant correlations between our 6-item scale and the three sub-scales of the IUIPC scale are negative and moderate in nature. It makes intuitive sense that peoples' opinions about targeted ads are correlated with their informational privacy concerns, but given the moderate size of these correlations, we can conclude that our measure is distinct from the construct measured by the IUIPC scale.

We conducted a second test of divergent validity by comparing the square root of the Average Variance Extracted (AVE) by our single factor with the correlations between our 6-item scale and the three established scales. The square root of the AVE by our single factor is 0.849, which is higher than all the correlations between our 6-item scale and the three established scales, thus establishing discriminant validity.

4.3.2.3 Test-Retest Reliability

Next, we evaluated the test-retest reliability of our 6-item scale. One hundred sixty-six of the original two hundred ninety-four participants completed our follow-up survey (response rate = 57%). The survey lasted 5 minutes and participants were paid \$0.30 as compensation. Our 6-item scale continues to show high internal reliability (Cronbach's alpha = 0.9383) and also shows high test-retest reliability ($r(164) = 0.75$, $p < 0.001$). Therefore, our 6-item scale is validated to reliably measure opinions about the practice of targeting advertising.

4.3.3 Discussion

In Validation Studies A and B, we validated our 6-item scale to measure opinions about targeted ads. Our scale shows high reliability, test-retest validity, and discriminant validity. Therefore, our 6-item scale can be used to reliably measure individuals' opinions about targeted ads. In order to use our scale, researchers should present participants with the hypothetical scenario about interest-based targeting shown in Appendix A.1. Then, participants should be told, "We are interested in understanding how you would feel about the type of targeted advertising described in the scenario on the previous page. Please indicate how strongly you agree or disagree with the following statements." The 6 items should be presented in random order and, for each item, participants should be asked to indicate their level of agreement on a 1–7 scale from 'Strongly Agree' to 'Strongly Disagree'.

4.4 Evaluation Study

4.4.1 Methods

The goal of this study was to evaluate how awareness about targeting impacts participants' purchase intentions. We borrowed the methodology used by Summers et al. in their 2016 work but adapted their study design to meet the goals of our study [38]. Specifically, we first ran a pilot study in which participants evaluated the perceived environmental friendliness of 32 different products belonging to 8 different product categories (such as light bulbs, laundry detergents, notebooks, etc.). These 32 products were the same ones used by Summers et al. in their 2016 work and are listed in Appendix A.3. The goal of this pilot study was to confirm that the Mechanical Turk population perceives at least some of these 32 products to be environmentally friendly. Participants were shown the 32 products in random order (one product at a time) and asked to indicate on a 1–7 scale (1 – Not at all, 7 – A lot) how environmentally friendly each product seems to them relative to other products in the same product category. Next, we recruited new participants from the same population for our main study. This study was conducted in two phases. In the first phase, we captured participants' opinions towards targeting practices using our validated 6-item scale from the Validation Study. By measuring these opinions *before* we showed participants our ad stimulus, we were able to collect an unbiased measure of our hypothesized moderator variable. In this phase, we also measured participants' tendency to express their value of environmental protection through the purchase of goods and services by asking participants to complete the Green Consumption Values scale [18]. We included a third unrelated scale in phase 1, the 'Dysfunctional Beliefs and Attitudes about Sleep' scale [31], as a decoy in order to ensure that participants are not able to make an obvious guess about the goal of our study or draw a direct connection between the two phases of our study. Finally, participants were asked to answer some demographic questions.

After a gap of about two weeks, we invited the participants who completed phase 1 to an ostensibly new study, which constituted the second phase of our main study. This second phase was almost identical to the study conducted by Summers et al. in their 2016 work [38]. Participants were told they would need to complete three different tasks in this study. The first task was a shopping task in which participants were asked to select one product from a set of four product options in several different product categories (the list of products is provided in Appendix A.3). We used this task to make the story of targeted advertising plausible. The second task

was an advertisement evaluation task in which participants were shown an advertisement for an acoustic speaker. Between conditions, we manipulated whether participants were made aware that the ad was targeted to them. We did this by telling participants in the 'Awareness' condition, "Our software will customize an advertisement for you based on your responses from the shopping task you completed earlier" before they were shown the ad. Participants in the 'No Awareness' condition were shown the same ad and were not given any information about targeting. In addition to the message about software matching, participants in the 'Awareness' condition were also shown the AdChoices icon on the acoustic speaker advertisement, the title of the page where the ad was shown was 'Targeted Advertisement', and the following text was shown above the ad: "The ad displayed below is customized for you based on your shopping choices earlier in today's session. Please take a moment to consider the advertisement below." In the 'No Awareness' condition the advertisement did not include the AdChoices icon, the title of the page was 'Advertisement', and the following text was shown above the ad: "Please take a moment to consider the advertisement below." All participants were shown the same advertisement, so it was not truly matched to any shopping choices. The advertisement shown to participants is provided in Appendix A.4.

As in Summers et al.'s 2016 work, we first asked participants how much they liked the advertisement on a 7-point scale (1 – Not at all, 7 – A lot) [38]. This question was included to confirm that participants' perceptions about the *advertisement* did not differ between conditions. Next, we asked participants how much they liked the product on a 7-point scale (1 – Not at all, 7 – A lot). This question captures attitudes towards the product. Finally, we measured purchase intentions towards the advertised product by using the same question as Summers et al. (2016). Specifically, we asked participants how likely they were to buy the acoustic speaker on a 7-point scale (1 – Very Unlikely, 4 – Undecided, 7 – Very Likely). On the next page, as a manipulation check, we asked participants the extent to which they agreed with the statement, "The advertisement shown to me was matched to my purchase choices from the earlier task in this study" on a 7-point scale (1 – Strongly Disagree, 7 – Strongly Agree).

Then participants continued to the final task in this study, which was almost identical to the final task used by Summers et al. in their 2016 work [38]. In this task, participants were informed that the researchers conducting this study have decided to partner with a different charity each month, to help our participants support these charities. They were informed that the selected charity this month is 'Rainforest Alliance' which "is a non-governmental organization (NGO) working to conserve biodiversity and ensure sustainable livelihoods by transforming land-use practices, business practices and consumer behavior." Participants were then informed that they would be entered in a lottery in which five participants would be randomly selected to win \$10, and they can choose to donate some or all of their winnings to the Rainforest Alliance if they are selected as one of the winners. They were also informed that the researchers would match any donation they chose to make towards the Rainforest Alliance. Then, we asked participants if they would like to donate to the Rainforest Alliance should they be chosen as a winner in the lottery. Those who responded affirmatively were asked for the exact amount that they wished to donate. This measure of donation behavior was used as an additional dependent variable to test whether the impact of awareness about targeting carries forward to subsequent decisions related to the factor

purportedly used in the targeting process (environmental friendliness). In other words, if participants feel that they were shown a targeted ad for an environmentally friendly product because the advertiser believes they are environmentally friendly, then they might change their behavior when deciding to subsequently donate to an environmentally friendly charity. Finally, participants were debriefed about the fact that the advertisement they saw was not truly targeted to them.

4.4.2 Results

4.4.2.1 Pilot Study

One hundred forty-four participants (Mean Age = 34.5; 60% Male) completed this study. The survey lasted 10 minutes and participants were paid \$0.50 as compensation. In order to determine whether a product is perceived to be environmentally friendly, we tested whether the mean value of reported environmental friendliness was significantly different from the midpoint of the scale, 4. In six of the eight product categories, at least one of the four products was perceived to be environmentally friendly. In the remaining two product categories (mouthwashes and digital cameras), none of the four products was perceived to be environmentally friendly. In order to improve the plausibility of our targeting scenario, we dropped these two product categories from the shopping task in the second phase. Therefore, the shopping task in the second phase included four product options in each of the following six product categories: light bulbs, laundry detergents, notebooks, air purifiers, dish scrubbers, and water bottles. The full results of this pilot, along with the means and standard deviations of the similar pilot study conducted by Summers et al. in their 2016 work, are presented in Appendix B.3.

4.4.2.2 Phase 1

Nine hundred ninety-two participants (Mean Age = 35.7; 53% Male) completed this study. The survey lasted 10 minutes and participants were paid \$0.50 as compensation. We created a consolidated value for opinions towards targeted ads by reverse coding the two negative items (8 – value) and then averaging the 6 scale items. Figure 5 shows the distribution of this consolidated value. Lower numbers indicate negative opinions towards targeted ads and higher numbers indicate neutral and positive opinions towards targeted ads. As is evident from Figure 5, many participants fell in the middle of the scale for our consolidated value of opinions about targeted ads. This may be because of participants' tendency to "pile on the midpoint" of response scales [1].




Figure 5. Distribution of opinions towards targeted ads using our 6-item scale.

4.4.2.3 Phase 2

Of the nine hundred ninety-two participants invited to participate in this study, six hundred ninety-seven (Mean Age = 36.7; 55% Male) completed Phase 2 (Response Rate = 70.3%). The survey lasted 10 minutes and participants were paid \$2 as compensation. The full results of all the models reported below in this section are presented in Appendix B.4. First, we analyze the responses to our manipulation check questions. Participants in the 'Awareness' condition believed that the advertisement was matched to their purchase choices more than those in the 'No Awareness' condition ($\text{Mean}_{\text{Awareness}} = 4.68$, $\text{Mean}_{\text{NoAwareness}} = 3.54$; $t = 8.78$; $p < 0.001$). This suggests that participants in the 'Awareness' condition believed that the advertisement was targeted. Participants' attitudes towards the advertisement, measured by how much they like the ad, were not significantly different between the two conditions ($\text{Mean}_{\text{Awareness}} = 4.46$, $\text{Mean}_{\text{NoAwareness}} = 4.55$; $t = 0.84$; $p = 0.40$). Next, we analyze the responses for our dependent variables, attitudes and purchase intentions. Overall, we find no significant effect of awareness about targeting on attitudes towards the product ($\text{Mean}_{\text{Awareness}} = 4.56$, $\text{Mean}_{\text{NoAwareness}} = 4.65$; $t = 0.63$; $p = 0.53$) or purchase intentions towards the product ($\text{Mean}_{\text{Awareness}} = 3.52$, $\text{Mean}_{\text{NoAwareness}} = 3.47$; $t = 0.37$; $p = 0.71$).

Next, we separately analyze how awareness about targeting impacts attitudes and purchase intentions towards the product among 1) participants with negative opinions about targeted advertising and 2) participants with neutral and positive opinions about targeted advertising. The former category comprises participants whose consolidated value on our measure for opinions about targeted advertising is 3 or lower ($N=228$), and the latter category comprises participants whose consolidated value on our measure for opinions about targeted advertising is 5 or greater ($N=176$). We estimate the following econometric models separately for these two categories of participants:

$$\text{LikeProduct}_i = \beta_0 + \beta_{\text{Awareness}} \text{Awareness} + \beta_{\text{Male}} \text{Male} + \beta_{\text{Age}} \text{Age} + \beta_{\text{Caucasian}} \text{Caucasian} + \beta_{\text{Education}} \text{Education} + \varepsilon_i$$

$$\text{BuyProduct}_i = \beta_0 + \beta_{\text{Awareness}} \text{Awareness} + \beta_{\text{Male}} \text{Male} + \beta_{\text{Age}} \text{Age} + \beta_{\text{Caucasian}} \text{Caucasian} + \beta_{\text{Education}} \text{Education} + \varepsilon_i$$

where 'LikeProduct' and 'BuyProduct' represent participants' attitudes towards the product and purchase intentions towards the product respectively; 'Awareness' represents a dummy variable that takes the value 1 for participants who are randomly assigned to the 'Awareness' condition and 0 for participants who are randomly assigned to the 'No Awareness' condition; 'Male' represents a dummy variable that takes the value 1 for male participants and 0 for female participants; 'Age' represents a continuous variable that represents the participant's age; 'Caucasian' is a dummy variable that takes the value 1 for Caucasian participants and 0 for all other participants; 'Education' is a categorical variable with seven categories ranging from "No high school" up to "Graduate degree"; and ' ε ' is the random error term.

Estimating the models above for participants with negative opinions about targeted ads showed that awareness about targeting has a significant negative impact on 'LikeProduct' ($\beta_{\text{Awareness}} = -0.45$; $p = 0.045$) and a marginally significant negative impact on 'BuyProduct' ($\beta_{\text{Awareness}} = -0.40$; $p = 0.08$) for these participants. This suggests that awareness about targeting worsens attitudes towards the targeted product and (marginally) reduces purchase intentions towards the targeted product among individuals with negative opinions about targeted ads. Estimating the models for

participants who have neutral and positive opinions about targeted ads showed no effect of awareness about targeting on ‘LikeProduct’ ($\beta_{\text{Awareness}} = -0.04$; $p=0.87$) and ‘BuyProduct’ ($\beta_{\text{Awareness}} = -0.02$; $p=0.96$). This suggests that awareness about targeting does not influence the attitudes and purchase intentions of individuals who have neutral and positive opinions towards targeted ads. The correlation between ‘LikeProduct’ and ‘BuyProduct’ is observed to be fairly high ($r=0.75$, $p<0.001$) so we created a combined dependent variable by averaging the values of these two dependent variables. This combined dependent variable, ‘CombinedProduct’, confirmed the results reported above with a significant effect of awareness about targeting ($\beta_{\text{Awareness}} = -0.42$; $p=0.048$) for the group of individuals who have negative opinions about targeted advertising, and no significant effect of awareness about targeting ($\beta_{\text{Awareness}} = -0.03$; $p=0.91$) for the group of individuals who have neutral and positive opinions about targeted advertising.

Next, we tested whether the size of the effect of awareness about targeting on attitudes and purchase intentions significantly varies between those who have negative opinions about targeted advertising and those who have neutral and positive opinions. We did this by introducing two additional variables to the basic econometric model above: 1) a dummy variable ‘Negative’ that takes the value 1 for individuals who have negative opinions about targeted advertising and 0 for those who have neutral and positive opinions about targeted advertising and 2) an interaction term between this dummy variable and the ‘Awareness’ dummy variable. To ensure that we are comparing the negative group to just the neutral and positive group, we did not include individuals whose consolidated values on our opinions measure fall between 3 and 5. The coefficient on the interaction term is not significant for ‘LikeProduct’ ($\beta_{\text{Awareness}*\text{Negative}} = -0.38$; $p=0.25$), ‘BuyProduct’ ($\beta_{\text{Awareness}*\text{Negative}} = -0.39$; $p=0.26$), and ‘CombinedProduct’ ($\beta_{\text{Awareness}*\text{Negative}} = -0.38$; $p=0.23$). This result suggests that the sizes of the effects of awareness about targeting on attitudes and purchase intentions do not differ significantly between individuals who have negative opinions about targeted advertising (‘LikeProducts’ Cohen’s $d = 0.23$; ‘BuyProducts’ Cohen’s $d = 0.19$) and individuals who have neutral and positive opinions (‘LikeProducts’ Cohen’s $d = 0.05$; ‘BuyProducts’ Cohen’s $d = 0.02$).

To explore alternative specifications of our moderation hypothesis, we also used a continuous measure of opinions about targeting (instead of splitting participants into groups). This allows us to include all participants in the analysis (even those whose consolidated opinions measure falls between 3 and 5) and treat participants with different values on the consolidated opinions measure differently (as opposed to grouping together everyone with values 3 or below and everyone with values 5 or above). While it is useful to treat the opinions measure as a continuous variable, doing so also introduces additional unexplained variance as we are now including participants who are not sure about their own opinions about targeted advertising (i.e., participants whose consolidated opinions measure falls between 3 and 5). We introduced the

consolidated value of our measure for opinions about targeting and the interaction of this measure with our ‘Awareness’ dummy variable to the basic model described above. The coefficient on the interaction term is not significant for ‘LikeProduct’ ($\beta_{\text{Awareness}*\text{Opinions}} = 0.09$; $p=0.27$), ‘BuyProduct’ ($\beta_{\text{Awareness}*\text{Opinions}} = 0.11$; $p=0.18$), and ‘CombinedProduct’ ($\beta_{\text{Awareness}*\text{Opinions}} = 0.10$; $p=0.19$), suggesting that the effect of awareness about targeting on attitudes and purchase intentions does not significantly vary when opinions about targeted advertising are varied on a continuous 1–7 scale.³

Finally, we analyze whether the donation behavior of our participants varied between the ‘Awareness’ and ‘No Awareness’ conditions. Participants reported whether or not they would choose to donate to the environmentally friendly charity if they win the lottery, and those who did choose to donate indicated the amount they would like to donate. As the former dependent variable is a dummy variable, we estimate the basic model as a probit. Overall, we find no significant difference in the likelihood to donate between the ‘Awareness’ and ‘No Awareness’ conditions ($\beta_{\text{Awareness}} = 0.03$; $p=0.79$). We also do not find any significant differences in the likelihood to donate between ‘Awareness’ and ‘No Awareness’ conditions when specifically looking at individuals who have negative opinions about targeted ads ($\beta_{\text{Awareness}} = 0.26$; $p=0.13$) and those who have neutral and positive opinions ($\beta_{\text{Awareness}} = -0.13$; $p=0.50$). We also analyzed the donation amounts using tobit models and find the same null results. Therefore, we do not find evidence that awareness about targeting impacts subsequent donation decisions made by individuals. Results from all models reported in this section are presented in Appendix B.4.

5. DISCUSSION

We investigated the effect of awareness about targeting on individuals’ attitudes and purchase intentions towards the advertised product. We find that, among individuals who have negative opinions about targeted ads, awareness about targeting worsens attitudes towards the advertised product and marginally reduces purchase intentions towards the advertised product. We find no effect of awareness on attitudes and purchase intentions towards the advertised product among individuals who have neutral and positive opinions about targeted ads. We also find that 53% of participants in our exploratory study (Exploratory Study) and 33% in our final study (Evaluation Study) reported having negative opinions about targeted advertising. Surveys conducted by previous researchers also suggest that a sizeable proportion of individuals have negative feelings about targeted advertising practices (66% of participants in [39] and 55% of participants in [29] said they do not want targeted advertisements). In addition to uncovering how awareness about targeting impacts attitudes and purchase intentions, we also built and validated a scale that can reliably measure individuals’ opinions about targeted advertising.

5.1 Implications

This research raises the important question of whether mere recommendations from the FTC to a self-regulated advertising industry about making consumers aware of targeting are enough to

³ We conducted the same analysis using the average of only the positive and negative items from our scale (i.e., without the neutral items) and obtained substantively similar results.

protect consumers' privacy. We find that awareness about targeting worsens attitudes towards the advertised product. Since a substantially large proportion of individuals have negative opinions about targeted advertising, our results suggest that the advertising industry is not incentivized to make consumers aware of targeted advertising, as such awareness could lead to lower effectiveness of advertisements. Because targeting practices such as interest-based advertising are not transparent, consumers may be unaware of how their information is being used to influence their purchase behaviors. We believe policy makers should consider introducing and enforcing regulations that require companies to make consumers aware of targeting practices. Another implication of our work is the ability for future researchers to capture attitudes towards interest-based targeted advertising in a reliable manner with our short 6-item scale.

5.2 Limitations

This work presents some limitations. First, we recruited participants from a single participant pool, Amazon Mechanical Turk, in all our studies. It is important to validate our results about targeted advertising with other participant pools. Although previous researchers have shown that Mechanical Turk workers are more demographically diverse than the typical convenience samples of American college students and that results using MTurk samples are similar to more traditional population pools [3, 7], this participant pool is likely to be more savvy about computers than the typical U.S. resident. A second limitation is that our scale validation and evaluation of awareness about targeting on attitudes is only conducted for the interest-based targeted advertising scenario. It is important to validate (and, if needed, modify) our findings for other types of targeting practices such as contextual and remarketed advertising. A third limitation of our work is highlighted by the difference in percentage of negative opinion participants between the Exploratory Study (53%) and the Evaluation Study (33%). This is likely due to differences in how we measured opinions in the Exploratory Study (with open-ended responses) and in the Evaluation Study (with our 6-point scale). A scale is useful in measuring opinions quickly without having to code open-ended responses, but it can introduce bias. A fourth limitation of our work is that we did not collect information about our participants' general offline and online shopping habits, which could have helped reduce statistical noise between conditions. Finally, our Evaluation Study does not account for users' possible overall dislike of ads. It may be interesting to account for that in future work by including a baseline condition with non-contextual, non-behavioral ads.

6. ACKNOWLEDGEMENTS

This research has been in part supported by the National Science Foundation under grants 1012763, 1514192, and 1327992, and by a Google Focused Research Award. We would like to thank all reviewers for their feedback and suggestions.

7. REFERENCES

- [1] Alreck, P. L., & Settle, R. B. (1985). *The Survey Research Handbook*, Homewood Ill: Irwin.
- [2] Beales, H. (2010). "The Value of Behavioral Targeting. Network Advertising Initiative," http://www.networkadvertising.org/pdfs/Beales_NAI_Study.pdf
- [3] Berinsky, A. J., Huber, G. A., & Lenz, G. S. (2012). "Evaluating online labor markets for experimental research: Amazon.com's Mechanical Turk," *Political Analysis*, 20(3), 351-68.
- [4] Bilgic, M., & Mooney, R. J. (2005). "Explaining Recommendations: Satisfaction vs. Promotion," *In Beyond Personalization Workshop*, IUI Vol. 5.
- [5] Blais, A.-R., & Weber, E. U. (2006). "A domain-specific risk-taking (DOSPERT) scale for adult populations," *Judgment and Decision Making*, 1, 33-47.
- [6] Brehm, J. W. (1966). *A Theory of Psychological Reactance*, New York: Academic Press.
- [7] Buhrmester, M., Kwang, T., & Gosling, S. D. (2011). "Amazon's Mechanical Turk: a new source of inexpensive, yet high-quality, data?" *Perspectives on Psychological Science*, 6(1), 3-5.
- [8] Chen, L., & Pu, P. (2005). "Trust Building in Recommender Agents," *In Proceedings of the Workshop on Web Personalization, Recommender Systems and Intelligent User Interfaces at the 2nd International Conference on E-Business and Telecommunication Networks*.
- [9] Clifford, S. (2010). "A Little 'i' to Teach About Online Privacy," *The New York Times*, Jan 26, http://www.nytimes.com/2010/01/27/business/media/27adco.html?_r=0
- [10] Cosley, D., Lam, S. K., Albert, I., Konstan, J. A., & Riedl J. (2003). "Is Seeing Believing?: How Recommender System Interfaces Affect Users' Opinions," *In ACM CHI, Vol. 1 of Recommender Systems and Social Computing*, 585-92.
- [11] Digital Advertising Alliance (2011). "Your Ad Choices," <http://www.youradchoices.com>
- [12] Edwards, S. M., Li, H., & Lee, J. (2002). "Forced exposure and psychological reactance: antecedents and consequences of the perceived intrusiveness of pop-up ads," *Journal of Advertising*, 31(3), 83-95.
- [13] Egelman, S., & Peer, E. (2015). "Predicting privacy and security attitudes," *ACM SIGCAS Computers and Society*, 45(1), 22-28.
- [14] Farhat, A., & Bailey, M. C. (2012). "How Effective Is Targeted Advertising?" *In Proceedings of the 21st International Conference on World Wide Web, ACM, Lyon, France*.
- [15] Federal Trade Commission (2000). "Fair Information Practices in the Electronic Marketplace," Federal Trade Commission Staff Report, Washington, DC.
- [16] Federal Trade Commission (2009). "Self-regulatory Principles for Online Behavioral Advertising," Federal Trade Commission Staff Report, Washington, DC.
- [17] Federal Trade Commission (2012). "Protecting Consumer Privacy in an Era of Rapid Change," Federal Trade Commission Staff Report, Washington, DC.
- [18] Haws, K., Winterich, K. P., & Reczek, R. W. (2013). "Seeing the world through GREEN-tinted glasses: Green consumption values and responses to environmentally

- friendly products,” *Journal of Consumer Psychology*, 24(3), 336-54.
- [19] Herlocker, J. L., Konstan, J. A., Borchers, A., & Riedl, J. (1999). “An Algorithmic Framework for Performing Collaborative Filtering,” *In Proceedings of the 22nd International Conference on Research and Development in Information Retrieval*, New York, 230-37.
- [20] Hill, K. (2012). “How Target Figured Out A Teen Girl Was Pregnant Before Her Father Did,” *Forbes*, Feb 12, <http://www.forbes.com/sites/kashmirhill/2012/02/16/how-target-figured-out-a-teen-girl-was-pregnant-before-her-father-did/#19b4300034c6>
- [21] Johnson, H., & Johnson, P. (1993). “Explanation Facilities and Interactive Systems,” *In Proceedings of Intelligent User Interfaces*, 159-66.
- [22] Leon, P. G., Cranshaw, J., Cranor, L. F., Graves, J., Hastak, M., Ur, B., & Xu, G. (2012a). “What Do Online Behavioral Advertising Privacy Disclosures Communicate to Users?” *In Proceedings of the 2012 ACM Workshop on Privacy in the Electronic Society*, 19-30.
- [23] Leon, P. G., Ur, B., Shay, R., Wang, Y., Balebako, R., and Cranor, L. F. (2012b). “Why Johnny Can’t Opt Out: A Usability Evaluation of Tools to Limit Online Behavioral Advertising,” *In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 589-98.
- [24] Leon, P. G., Ur, B., Wang, Y., Sleeper, M., Balebako, R., Shay, R., Bauer, L., Christodorescu, M., & Cranor, L. F. (2013). “What Matters to Users?: Factors That Affect Users’ Willingness to Share Information with Online Advertisers,” *In Proceedings of the Ninth Symposium on Usable Privacy and Security*, ACM, 7.
- [25] Leon, P. G., Rao, A., Schaub, F., Marsh, A., Cranor, L. F., & Sadeh, N. (2014). “Why People Are (Un)willing to Share Information with Online Advertisers,” *In Workshop on Privacy in Electronic Society*, <http://www.andrew.cmu.edu/user/pgl/wpes2014oba.pdf>
- [26] Malheiros, M., Jennett, C., Patel, S., Brostoff, S., & Sasse, M. A. (2012). “Too Close for Comfort: A Study of the Effectiveness and Acceptability of Rich-Media Personalized Advertising,” *In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ACM, 579-88.
- [27] Malhotra, Naresh K., Kim, Sung S., & Agarwal, J. (2004). “Internet users’ information privacy concerns (IUIPC): The construct, the scale, and a causal model,” *Information Systems Research*, 15(4), 336-55.
- [28] Mattioli, D. (2012). “On Orbitz, Mac Users Steered to Pricier Hotels,” *The Wall Street Journal*, Aug 23, <http://www.wsj.com/articles/SB10001424052702304458604577488822667325882>
- [29] McDonald, A., & Cranor, L. F. (2010). “Beliefs and Behaviors: Internet Users’ Understanding of Behavioral Advertising,” *In Telecommunications Policy Research Conference*, Arlington.
- [30] Melicher, W., Sharif, M., Tan, J., Bauer, L., Christodorescu, M., & Leon, P. G. (2016). “(Do Not) Track Me Sometimes: Users’ Contextual Preferences for Web Tracking,” *In Proceedings on Privacy Enhancing Technologies*, Vol. 2, 13554.
- [31] Morin, C. M., Vallieres, A., & Ivers, H. (2007). “Dysfunctional beliefs and attitudes about sleep (DBAS): validation of a brief version (DBAS-16),” *Sleep*, 30, 1547-54.
- [32] Nunnally, J. C. (1978). *Psychometric Theory*, 2nd ed., New York: McGraw-Hill.
- [33] Peer, E., Vosgerau, J., & Acquisti, A. (2014). “Reputation as a sufficient condition for data quality on Amazon Mechanical Turk,” *Behavior Research Methods*, 46(4), 1023-31.
- [34] Purcell, K., Brenner J., & Rainie, L. (2012). “Search Engine Use 2012,” PEW Research Center Technical Report, http://www.pewinternet.org/files/old-media/Files/Reports/2012/PIP_Search_Engine_Use_2012.pdf
- [35] Rudder, C. (2014). “We Experiment on Human Beings!” OkCupid blog, Jul 28, <http://blog.okcupid.com/index.php/we-experiment-on-human-beings/>
- [36] Scott, S. G., & Bruce, R. A. (1995). “Decision-making style: The development and assessment of a new measure,” *Educational and Psychological Measurement*, 55(5), 818-31.
- [37] Sinha, R., & Swearingen, K. (2002). “The Role of Transparency in Recommender Systems,” *In CHI Extended Abstracts on Human Factors in Computing Systems*, ACM, 830-31.
- [38] Summers, C. A., Smith, R. W., & Reczek, R. W. (2016). “An Audience of One: Behaviorally Targeted Ads as Implied Social Labels,” *Journal of Consumer Research*, ucw012.
- [39] Turow, J., King, J., Hoofnagle, C. J., Bleakley, A., & Hennessy, M. (2009). “Americans Reject Tailored Advertising and Three Activities That Enable It,” SSRN 1478214.
- [40] Valentino-Devries, J., Singer-Vine, J. & Soltani, A. (2012). “Websites Vary Prices, Deals Based on Users’ Information,” *The Wall Street Journal*, Dec 24, <http://www.wsj.com/articles/SB10001424127887323777204578189391813881534>
- [41] White, T. B., Zahay, D. L., Thorbjørnsen, H., & Shavitt, S. (2007). “Getting too personal: reactance to highly personalized email solicitations,” *Marketing Letters*, 19 (1), 39-50.

APPENDIX

Appendix A

This appendix contains the study materials.

A.1 Scenarios – Exploratory Study

A.1.1 Contextual Targeting Condition

Imagine that you are reading an article on the New York Times website about how to protect your shoes from being damaged by rain and snow. You notice that the advertisement next to the article is for a pair of shoes.

You realize that the advertisement has been specifically targeted to you based on the article that you're reading.

A.1.2 Interest-Based Targeting Condition

Imagine that you are looking to buy a new pair of shoes online. You look at a few different websites that sell shoes. You have not made your decision about which pair of shoes you want to purchase.

A few days later, you are reading an article on the New York Times website. You notice that the advertisement next to the article is for a pair of shoes. You had not seen this particular pair of shoes when you were browsing for shoes some days before.

You realize that your browsing history had been tracked and that it is being used to display an advertisement that has been specifically targeted to you.

A.1.3 Remarketing Condition

Imagine that you are looking to buy a new pair of shoes online. You look into a few different websites that sell shoes. You find a pair of shoes that you like, on a website that you've never visited before. You spend some time looking at this particular pair of shoes and then move on to doing something else. You haven't decided whether or not you want to purchase this particular pair of shoes.

A few days later, you are reading an article on the New York Times website. You notice that the advertisement next to the article is for the same pair of shoes that you were looking at the other day.

You realize that your browsing history had been tracked and that it was being used to display an advertisement that has been specifically targeted to you.

A.2 Self-developed Measures – Exploratory Study

A.2.1 Perceived Control

To what extent do you feel that you have control over this type of targeted advertising? In other words, to what extent do you feel that you can stop receiving this particular type of targeted ads, if you no longer want them?

Please answer this question with respect to the type of targeting described in the scenario on the previous page. [5 points scale from 'Not at all' to 'Very much']

<If the response is 4 or 5 on the previous question then the following question is shown>

Please tell us how you would control this type of targeted ads. In other words, how would you stop receiving this particular type of targeted ads, if you no longer want them? [Open-ended response]

A.2.2 Previous Experience with Online Privacy Invasions Related to Targeted Ads

Have you personally ever experienced an unpleasant or uncomfortable scenario involving targeted ads?

<If yes> Please describe the unpleasant or uncomfortable scenario that you experienced involving targeted ads.

A.2.3 Previous Experience with Online Privacy Invasions in General

Have you personally ever been the victim of what you felt was an improper invasion of your online privacy (irrespective of whether it involved targeted ads or not)? [Modified from Surveys by Louis Harris and Associates and Harris Interactive⁴]

<If yes> Please describe the online privacy invasion that you experienced.

A.2.4 Current Interest in Purchasing Shoes

How interested are you currently in purchasing a new pair of shoes, either online or offline? [5 points scale from 'Not at all' to 'Very much']

A.3 Products – Evaluation Study [38]

All 32 products shown below were tested in the pilot study. The four digital cameras and four mouthwashes were dropped from the list of products used in the shopping task in the second phase.

 GE 25-Watt Tiffany Stained Glass Light Bulb \$5.39	 GE Reveal 53-Watt Halogen Clear A19 General Purpose Light Bulbs 2-ct. \$7.99	 GE Energy Smart 60-Watt General Purpose Light Bulbs 2-ct. \$15.39	 GE Energy Smart 13-Watt Soft White Light Bulbs 2-pk. \$14.49
 Seventh Generation Natural Liquid Laundry Detergent \$17.39	 Tide Total Care HE Laundry Detergent - Renewing Rain \$19.99	 Arm & Hammer Liquid Laundry Detergent For Sensitive Skin \$17.99	 Mrs. Meyer's Clean Days Lavender Laundry Detergent \$19.99

⁴ Surveys by Louis Harris and Associates for Southern New England Telephone, September 1-11, 1983, and by Harris Interactive for Business Week, March 3-6, 2000, reported about by The Roper Center for Public Opinion Research:

<http://www.ropercenter.cornell.edu/public-perspective/ppscan/116/116012.pdf>

			
Nikon COOLPIX S3500 20MP Digital Camera with 7x Optical Zoom \$109.99	Polaroid 300 Instant Camera - Purple (PIC-300L) with 10 Pack \$79.99	Canon PowerShot SX-500 16MP Digital Camera with 30x Optical Zoom - Black \$199.99	PENTAX Optio WG-10 14MP Waterproof Digital Camera with 5x Optical Zoom \$179.99
			
Act Fluoride Rinse - Mint \$7.99	Colgate® Phos-Flur® Ortho Protect Rinse - Mint \$10.99	Listerine Total Care Fresh Mint \$9.29	Tom's of Maine Natural Cool Mountain Mint \$10.54
			
Moleskine Hard Cover Notebook - Orange \$18.95	Lang Deluxe Journal Morning Has Broken \$11.95	Greenroom Recycled Spiral Blank Journal \$5.99	Blank Journal Markings \$7.49
			
Holmes Eco-Friendly Air Purifier \$39.99	CleanAirBall Air Purifier \$49.99	Honeywell True HEPA Air Purifier \$119.09	Vornado AC300 Whole Room Air Purifier \$149.99
			
Scotch-Brite Natural Fiber Non-Scratch Scrub Sponge \$4.49	O-Cel-O No-Scratch Scrub & Wipe Pad \$6.19	KitchenAid Black Soap Dispensing Palm Brush \$8.29	WayClean Mesh Scrubber \$3.99
			
Rive Saboy Water Bottle \$12.99	Elio Pure Fizz \$12.99	Contigo Double Wall Water Bottle \$10.99	Aladdin Recycle & Recyclable Travel Mug \$10.99

A.4 Advertisement– Evaluation Study [38]

A.4.1 Awareness Condition



Acoustic Speaker by Houd
Green, energy-free speaker crafted from sustainably sourced Colombian wood.
\$85.00

A.4.2 No Awareness Condition



Acoustic Speaker by Houd
Green, energy-free speaker crafted from sustainably sourced Colombian wood.
\$85.00

Appendix B

This appendix contains the study results.

B.1 Logistic Regression – Exploratory Study

Odds Ratio:

	(1)	(2)	(3)
	Negative Opinions	Negative Opinions	Negative Opinions
	All conditions	Contextual targeting conditions	Interest-based targeting and Remarketing conditions
Perceived Control	0.472*** (0.103)	0.567 (0.229)	0.304*** (0.114)
Previous Unpleasant Targeted Ad Experience	4.040** (2.642)	7.708* (8.17)	2.402 (2.215)
Male	0.662 (0.320)	1.014 (0.843)	0.419 (0.296)
Age	0.958 (0.025)	0.976 (0.042)	0.937 (0.038)
Caucasian	0.370* (0.212)	0.379 (0.327)	0.181* (0.171)
Unemployed	0.635 (0.420)	0.719 (0.880)	0.312 (0.305)
Experience in IT	0.923 (0.576)	0.705 (0.770)	1.30 (1.26)

*p<0.10; **p<0.05; ***p<0.01; Standard errors in brackets

B.2 OLS Regression – Validation Study A

	(1)
	Consolidated value of our 6-item scale
Perceived Control	0.140** (0.055)
Previous Unpleasant Targeted Ad Experience	-1.163*** (0.167)
Male	0.058 (0.168)
Age	-0.013* (0.008)
Caucasian	-0.328* (0.196)
Unemployed	0.216 (0.202)
Experience in IT	-0.398* (0.219)
Constant	5.222*** (0.543)

*p<0.10; **p<0.05; ***p<0.01; Standard errors in brackets

B.3 Results – Evaluation Study (Pilot)

Product Category	Product	Mean	Std Dev	T test against scale midpoint 4	Environmentally friendly?	Summers et al. (2016) Mean, N = 45	Summers et al. (2016) Std Dev, N = 45
Light bulbs	GE Tiffany Stained Glass	2.63	1.47	$t(143) = -11.20$, $p < 0.0001$	No	2.74	1.42
	GE Reveal Halogen	3.21	1.57	$t(143) = -6.06$, $p < 0.0001$	No	3.47	1.35
	GE Energy Smart	5.37	1.45	$t(143) = 11.31$, $p < 0.0001$	Yes	5.28	1.41
	GE Energy Smart-Soft White	5.29	1.45	$t(143) = -10.67$, $p < 0.0001$	Yes	5.55	1.19
Laundry detergents	Seventh Generation Natural Liquid	5.42	1.34	$t(143) = 12.73$, $p < 0.0001$	Yes	5.36	1.17
	Tide Total Care	3.10	1.53	$t(143) = -7.02$, $p < 0.0001$	No	3.48	1.31
	Arm & Hammer Sensitive Skin	3.26	1.39	$t(143) = -6.34$, $p < 0.0001$	No	3.26	1.12
	Mrs. Meyer's Clean Lavender	4.12	1.73	$t(143) = 1.16$, $p = 0.25$	Can't tell	3.04	1.84
Digital cameras	Nikon COOLPIX S3500	2.85	1.38	$t(143) = -9.94$, $p < 0.0001$	No	2.74	1.48
	Polaroid 300 Instant	2.54	1.34	$t(143) = -13.09$, $p < 0.0001$	No	2.48	1.22
	Canon PowerShot Sx-500	2.71	1.43	$t(143) = -10.82$, $p < 0.0001$	No	2.85	1.62
	PENTAX Optio WG-10	2.93	1.43	$t(143) = -8.96$, $p < 0.0001$	No	3.00	1.74
Mouthwashes	Act Fluoride Rinse	2.97	1.37	$t(143) = -9.09$, $p < 0.0001$	No	3.28	1.46
	Colgate Phos-Flur Ortho Protect Rinse	3.07	1.37	$t(143) = -8.17$, $p < 0.0001$	No	2.89	1.35
	Listerine Total Care	3.10	1.45	$t(143) = -7.50$, $p < 0.0001$	No	3.33	1.38
	Tom's of Maine Natural Cool	4.49	1.56	$t(143) = 3.79$, $p < 0.001$	No	3.98	1.78
Notebook	Moleskin Hard Cover	3.06	1.34	$t(143) = -8.37$, $p < 0.0001$	No	2.78	1.60
	Lang Journal Deluxe	3.15	1.48	$t(143) = -6.92$, $p < 0.0001$	No	2.85	1.71

	Greenroom Recycled Spiral	5.01	1.70	$t(143) = 7.17,$ $p < 0.0001$	Yes	5.07	1.69
	Blank Journal Markings	3.16	1.38	$t(143) = -7.32,$ $p < 0.0001$	No	2.98	1.73
Air purifiers	Holmes Eco-Friendly	5.00	1.45	$t(143) = 8.29,$ $p < 0.0001$	Yes	4.39	1.71
	CleanAirBall	4.15	1.55	$t(143) = 1.13,$ $p = 0.26$	Can't tell	4.48	1.41
	Honeywell True HEPA	3.89	1.54	$t(143) = -0.97,$ $p = 0.33$	Can't tell	4.13	1.47
	Vornado AC300 Whole Room	3.92	1.43	$t(143) = -0.70,$ $p = 0.48$	Can't tell	3.96	1.58
Dish scrubbers	Scotch-Brite Natural Fiber	5.51	1.30	$t(143) = 13.97,$ $p < 0.0001$	Yes	5.07	1.39
	O-Cel-O No Scratch	3.17	1.32	$t(143) = -7.56,$ $p < 0.0001$	No	3.07	1.44
	KitchenAid Soap Dispensing Palm Brush	3.26	1.34	$t(143) = -6.60,$ $p < 0.0001$	No	3.48	1.41
	WayClean Mesh	3.33	1.28	$t(143) = -6.25,$ $p < 0.0001$	No	3.39	1.37
Water bottle	Rive Saboy	4.01	1.63	$t(143) = 0.05,$ $p = 0.96$	Can't tell	4.04	1.93
	Ello Pure Fizz	4.15	1.57	$t(143) = 1.12,$ $p = 0.27$	Can't tell	4.17	1.77
	Contigo Double Wall	3.91	1.62	$t(143) = -0.67,$ $p = 0.50$	Can't tell	4.07	1.82
	Aladdin Recycle Travel	5.69	1.26	$t(143) = 16.08,$ $p < 0.0001$	Yes	5.39	1.76

B.4 Results – Evaluation Study (Phase 2)

OLS regression coefficients among all participants (N=697):

	(1) LikeProduct	(2) BuyProduct	(3) CombinedProduct
Awareness	−0.084 (0.120)	0.042 (0.129)	−0.021 (0.117)
Male	−0.128 (0.121)	−0.160 (0.130)	−0.144 (0.117)
Age	0.003 (0.005)	0.003 (0.006)	0.003 (0.005)
Caucasian	0.243* (0.147)	0.220 (0.158)	0.231 (0.143)
Education	−0.108** (0.047)	−0.126** (0.050)	−0.117** (0.045)
Constant	4.989*** (0.330)	3.92*** (0.354)	4.456*** (0.320)

*p<0.10; **p<0.05; ***p<0.01; Standard errors in brackets

OLS regression coefficients among participants who have negative opinions about targeted ads (N=228):

	(1) LikeProduct	(2) BuyProduct	(3) CombinedProduct
Awareness	−0.448** (0.222)	−0.401* (0.229)	−0.424** (0.213)
Male	−0.048 (0.230)	0.056 (0.236)	0.004 (0.220)
Age	0.025** (0.010)	0.023** (0.010)	0.024*** (0.009)
Caucasian	−0.184 (0.269)	−0.042 (0.276)	−0.113 (0.258)
Education	−0.173** (0.087)	−0.209** (0.090)	−0.191** (0.084)
Constant	4.807*** (0.626)	3.697*** (0.644)	4.252*** (0.601)

*p<0.10; **p<0.05; ***p<0.01; Standard errors in brackets

OLS regression coefficients among participants who have neutral and positive opinions about targeted ads (N=176):

	(1) LikeProduct	(2) BuyProduct	(3) CombinedProduct
Awareness	−0.042 (0.243)	−0.015 (0.266)	−0.029 (0.242)
Male	−0.239 (0.244)	−0.477* (0.266)	−0.358 (0.243)
Age	−0.010 (0.010)	−0.0004 (0.011)	−0.005 (0.010)
Caucasian	0.231 (0.337)	0.357 (0.369)	0.294 (0.336)
Education	−0.061 (0.092)	−0.036 (0.100)	−0.048 (0.092)
Constant	5.422*** (0.659)	3.828*** (0.720)	4.625*** (0.656)

*p<0.10; **p<0.05; ***p<0.01; Standard errors in brackets

OLS regression coefficients for models including the interaction between ‘Awareness’ dummy and ‘Negative’ dummy among participants who have negative opinions or neutral and positive opinions (N=404):

	(1) LikeProduct	(2) BuyProduct	(3) CombinedProduct
Awareness	−0.034 (0.248)	0.007 (0.262)	−0.014 (0.242)
Male	−0.125 (0.168)	−0.164 (0.177)	−0.144 (0.163)
Age	0.009 (0.007)	0.013* (0.008)	0.011 (0.007)
Caucasian	0.005 (0.210)	0.121 (0.221)	0.063 (0.204)
Education	−0.120* (0.063)	−0.136** (0.067)	−0.128** (0.062)
Negative	−0.133 (0.236)	−0.171 (0.249)	−0.152 (0.230)
Awareness*Negative	−0.378 (0.329)	−0.391 (0.347)	−0.384 (0.320)
Constant	5.130*** (0.465)	3.889*** (0.490)	4.510*** (0.453)

*p<0.10; **p<0.05; ***p<0.01; Standard errors in brackets

OLS regression coefficients for models including the interaction between 'Awareness' dummy and continuous measure for opinions about targeted advertising (N=697):

	(1) LikeProduct	(2) BuyProduct	(3) CombinedProduct
Awareness	−0.419 (0.324)	−0.394 (0.347)	−0.407 (0.313)
Male	−0.115 (0.121)	−0.146 (0.129)	−0.131 (0.117)
Age	0.004 (0.005)	0.004 (0.006)	0.004 (0.005)
Caucasian	0.223 (0.147)	0.196 (0.157)	0.209 (0.142)
Education	−0.107** (0.047)	−0.127** (0.050)	−0.117** (0.045)
Opinions	0.052 (0.057)	0.054 (0.061)	0.053 (0.055)
Awareness*Opinions	0.087 (0.079)	0.113 (0.085)	0.100 (0.077)
Constant	4.780*** (0.403)	3.708*** (0.432)	4.244*** (0.390)

*p<0.10; **p<0.05; ***p<0.01; Standard errors in brackets

Probit results for likelihood to donate to the charity:

	(1) Donate All participants (N=697)	(2) Donate Only 'Negative' group (N=228)	(3) Donate Only 'Neutral & Positive' group (N=176)
Awareness	0.010 (0.038)	0.102 (0.067)	−0.053 (0.078)
Male	−0.047 (0.038)	−0.108 (0.069)	−0.065 (0.078)
Age	0.004** (0.002)	−0.005 (0.003)	0.008** (0.004)
Caucasian	0.029 (0.047)	−0.033 (0.081)	−0.009 (0.110)
Education	−0.019 (0.015)	−0.009 (0.026)	−0.042 (0.030)

*p<0.10; **p<0.05; ***p<0.01; Standard errors in brackets

Tobit results for amount of donations made to the charity (left censored at 0):

	(1) Donation Amount All participants (N=697)	(2) Donation Amount Only 'Negative' group (N=228)	(3) Donation Amount Only 'Neutral & Positive' group (N=176)
Awareness	−0.077 (0.405)	0.848 (0.694)	−0.858 (0.750)
Male	−0.488 (0.407)	−0.667 (0.716)	−0.558 (0.754)
Age	0.036** (0.017)	−0.040 (0.030)	0.057* (0.031)
Caucasian	0.455 (0.502)	−0.424 (0.832)	0.561 (1.047)
Education	−0.125 (0.157)	−0.246 (0.272)	−0.313 (0.284)
Constant	−0.383 (1.112)	3.413* (1.969)	0.574 (2.028)

*p<0.10; **p<0.05; ***p<0.01; Standard errors in brackets