

When and Why Do People Want Ad Targeting Explanations? Evidence from a Four-Week, Mixed-Methods Field Study

Hao-Ping (Hank) Lee[†], Jacob Logas[‡], Stephanie Yang[‡], Zhouyu Li^{*}, Natã Barbosa[§], Yang Wang[§], Sauvik Das[†]
*Carnegie Mellon University[†], Georgia Institute of Technology[‡],
 North Carolina State University^{*}, University of Illinois Urbana-Champaign[§]*
 haopingl@cs.cmu.edu, {logasja, syang479}@gatech.edu, zli85@ncsu.edu, {natamb2, yvw}@illinois.edu, sauvik@cmu.edu

Abstract—Many people are concerned about how their personal data is used for online behavioral advertising (OBA). Ad targeting explanations have been proposed as a way to reduce this concern by improving transparency. However, it is unclear *when* and *why* people might want ad targeting explanations. Without this insight, we run the risk of designing explanations that do not address real concerns. To bridge this gap, we conducted a four-week, mixed-methods field study with 60 participants to understand when and why people want targeting explanations for the ads they actually encountered while browsing the web. We found that users wanted explanations for around 30% of the 4,251 ads we asked them about during the study, and that subjective perceptions of how their personal data was collected and shared were highly correlated with when users wanted ad explanations. Often, users wanted these explanations to confirm or deny their own preconceptions about how their data was collected or the motives of advertisers. A key upshot of our work is that one-size-fits-all approaches to ad explanations are likely to fail at addressing people’s lived concerns about ad targeting; instead, more personalized explanations are needed.

1. Introduction

Online Behavioral Advertising (OBA) is “the practice of monitoring people’s online behavior and using the collected information to show people individually targeted advertisements” [1]. At its best, OBA is beneficial for all stakeholders: advertisers reach their target audience, end-users find products and services specific to their interests, and ad brokers profit by connecting the two. Indeed, digital ad spending in the U.S. alone is estimated to surpass 230 billion USD in 2022, and the majority of that spending will go towards OBA [2]. In reality, while OBA benefits advertisers and ad brokers, many end-users find these targeted advertisements to be “surprising” and “creepy” [3], and surveys from the Pew Research Center and Statista in 2019 suggest that the majority of U.S. adult Internet users view behavioral profiling and targeted advertisements negatively [4]. These negative user perceptions and the privacy concerns thereof stem partially from the fact that personal data that is processed to target end-users in OBA is often collected covertly and without informed consent [3], [5]. This lack of tracking and targeting transparency, in turn, can lead to other security concerns: e.g., the malicious use of ad targeting to spread misinformation, malware, and scams [6].

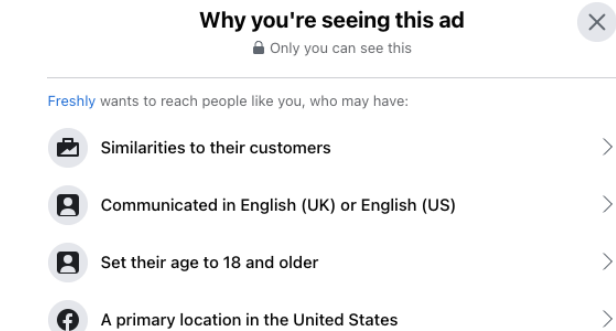


Figure 1. An ad targeting explanation on Facebook. Ad targeting explanations are commonly employed to help improve transparency in OBA, but little is known about when and why people might want such explanations.

Ad targeting explanations — layman descriptions of why a user is seeing a particular targeted ad — have been introduced to improve OBA transparency. For example, Google, Facebook and Twitter provide user-specific descriptions of inferred interests used for ad targeting on their platforms (e.g., Figure 1). However, prior work suggests that extant explanations are largely ineffective: users are either unaware that they exist and/or ignore them altogether [3] and find the explanations insufficient and unnerving [7]. Prior work has employed human-centered design methodologies to provide insights into *how* to craft such explanations to enhance OBA transparency [7], [8] and increase user satisfaction with individual explanations. However, we still have little insight into *when* and *why* end-users might want ad targeting explanations. Answering the *when* question is important because users might see many targeted advertisements and not all of them might warrant concern; thus, showing ad targeting explanations for every targeted ad will invariably lead to habituation effects as has been shown to occur in prior work on security & privacy warnings and notifications [9]. Answering the *why* question is important to craft explanations that contain information that users *want* to see for ads that actually elicit concern — for example, if a user wants to know how an ad broker learned they are in the market for shoes.

Here, we aim to answer two key research questions:

- RQ1** When do people want ad targeting explanations?
- RQ2** Why do people want ad targeting explanations?

We answer RQ1 — the *when* question — through a four-week field study with 60 participants. We developed a browser extension to keep track of the ads that participants would see in their everyday web browsing throughout

the study. We asked participants to complete at most six questionnaires daily, each corresponding to an online ad captured that day. The questionnaire asked participants to reflect on their feelings about the ad and if they wanted an explanation for why they were targeted with the ad. To answer RQ2 — the *why* question — we conducted 60-minute exit interviews with 36 of the 60 participants where we asked participants to expand on their reasoning for (not) wanting explanations. Note that not all online ads are targeted. In this study, we nevertheless presented all collected advertisements as “targeted ads.” Thus, all online ads that participants encountered were used as a probe to model when and why they might have wanted explanations.

We found that users wanted ad targeting explanations for approximately 30% of the 4,251 ads we brought to their attention. Their subjective perceptions of specific ads — e.g., idiosyncratic beliefs of how their data was collected and shared, motives of advertisers, and emotional reactions to ad content — were highly correlated with their desire for explanations. These subjective perceptions were more predictive of when users wanted explanations than objectively measurable properties of the ad, correlated user behavior, viewing context, and general user attitudes towards security & privacy.

We also uncovered key reasons why users (did not) want ad explanations, which related to their expectations (e.g., was the targeting accurate?), preconceptions (e.g., were their conversations heard?), perceptions (e.g., is the targeting creepy?), and attitudes (e.g., is the ad offensive?) toward an ad and being targeted with it. A key implication of our results is that existing approaches to crafting ad targeting explanations — often one-size-fits-all templates — are likely unsatisfying for users. Instead, *when* users want ad targeting explanations, they want to test their hypotheses of data collection and inference, investigate advertiser motivations, or situate an ad within a noticed trend.

To summarize, our work makes three main contributions:

- Using ecologically valid data collected from a four-week field study, we build quantitative models that help explain and predict *when* people (do not) want ad targeting explanations.
- Drawing from interviews with 36 participants, we construct a qualitative model of *why* people (do not) want ad targeting explanations.
- We synthesize design implications for ad targeting explanations that are rooted in these models.

2. Background and Related Work

Our work builds upon and extends a rich tradition of prior work that (i) models user perceptions and attitudes towards OBA, and (ii) designs and evaluates novel mechanisms to improve ad targeting transparency. Prior work in the former category was instrumental in helping us understand what data to collect and include in our models of *when* and *why* users want ad targeting explanations. Prior work in the latter category helped contextualize our findings to

diagnose why extant ad targeting explanations appear to be largely ineffective at addressing user concerns.

2.1. User Perceptions of OBA

Prior work suggests that user perceptions towards OBA are mixed and context-dependent. Ur et al. [3] found their participants had different attitudes under different browsing scenarios, e.g., they are more open to data collection for OBA when they are reading the news than when searching for medical treatments for a friend. Smit et al. [10] found users’ overall concerns about online privacy varied. Such differences were also found to be associated with their overall attitudes toward OBA. Advertisement content, repetition, or unexpected characterizations of the user were found to exacerbate users’ negative perceptions of targeted advertising [11], [12]. Still, across all demographics, the majority of users indicate they don’t want tailored advertisements [13]. To help explain why people feel negatively towards OBA, scholars have drawn on “Social Presence Theory” to argue that when the computer collects one’s data it has the same effect as when a person looks over one’s shoulder [14]. Over time, this sense of constantly having a “person over the shoulder” has made users more concerned about OBA practices and their privacy [15]. Some users even report that they have changed their online behavior because of this surreptitious data collection [16].

Prior work also suggests that a minority of users report positive perceptions of OBAs. Positive perceptions of OBAs arise from a contextual “privacy calculus” in which this minority of users view the benefits of OBA to outweigh its costs [17]. In many ways, these attitudes reflect the “free market” approach to privacy, wherein users make decisions of who they do business with to protect their privacy [18]. The resulting “privacy calculus” takes into account factors like the amount of trust users have in a provider, if the data exchange is for a service, and whether a third party is privy to the data collected [19], [20].

We built on this prior work by incorporating subjective perceptions as one of the underlying factors that we hypothesized might correlate with a user’s desire for an ad targeting explanation. In particular, we treat perceptual (subjective impressions of and reactions to specific ads), contextual (pertaining to the ad itself or the context in which it is seen), and user-level (general attitudes towards privacy) factors as inputs to build, to our knowledge, the first model to understand when and why users might want ad targeting explanations with comprehensive, in-situ field data.

2.2. Efforts to Improve OBA Transparency

Users generally desire greater transparency from OBAs and view OBAs more positively when made more transparent [13], [21]. Accordingly, much prior work has explored mechanisms to improve OBA transparency. OBA providers often provide privacy statements and informed consent requests because of government regulations. However, these statements and requests are often lengthy and overly complicated, so many users either do not comprehend them or

ignore them altogether [22]–[26]. Another approach is to use icons that identify targeted advertisements; prior work has shown that these OBA icons are rarely noticed and are unclear in purpose [3], [27], [28]. Stanaland found that “privacy trustmarks” — an icon pledging to maintain a set of standards — improve trust where OBA icons do not [29].

Ad targeting explanations are also widely used to help increase transparency in OBAs. In exploring how to better craft these explanations, Eslami et al. found that users prefer interpretable and non-creepy explanations with a clear connection to their identity. These explanations resulted in users finding targeting algorithms to be less capable than initially thought [7]. Similarly, Barbosa et al. found that providing insights about the targeting model as it evolves to help demystify the connection between behavior and advertisements. Additionally, study participants wanted a more collaborative relationship with the model when inaccuracies arose [30]. Lastly, Wei et al. investigated modes of ad explanation (e.g., Detailed Visual, Detailed Text, Creepy) that expand upon existing explanations provided by Twitter, finding those that had greater detail were preferred by participants [8].

To date, prior work on improving ad targeting explanations has focused on improving the general communication design of these explanations in situations where users are already curious. In contrast, our work is, to our knowledge, the first to systematically model *when* and *why* users want ad targeting explanations and the idiosyncratic variations across users thereof. This knowledge, in turn, is crucial to crafting more personalized and effective ad targeting explanations that users find effective and do not ignore.

3. Method

We conducted a mixed-methods study consisting of: (i) a four-week field study in which 60 participants installed a browser extension that tracked, with consent, the ads they encountered and asked them if they wanted ad targeting explanations for a random subset of these ads; and, (ii) exit interviews with 36 of these participants to gain more insight into why they did or did not want the explanations.

Specifically, to answer **RQ1** we created explanatory regression models correlating participants’ desires for ad explanations with three sets of factors that we hypothesized might correlate with a desire for ad targeting explanations: 1) **Contextual factors**: measures about a specific ad and its *in-situ* viewing context — e.g., the type of advertisement, a user’s browsing activity. 2) **Perceptual factors**: measures about users’ perceptions towards an ad and beliefs about how they were targeted — e.g., perceptions of creepiness and the covert surveilling of physical world conversations [3], [31], [32]. 3) **User factors**: measures about users — e.g., personality, attitudes towards security and privacy.

To answer **RQ2**, we conducted 36 semi-structured exit interviews with participants at the end of the data collection period to unpack the reasons for their (dis)interest in ad explanations.

3.1. Field Study

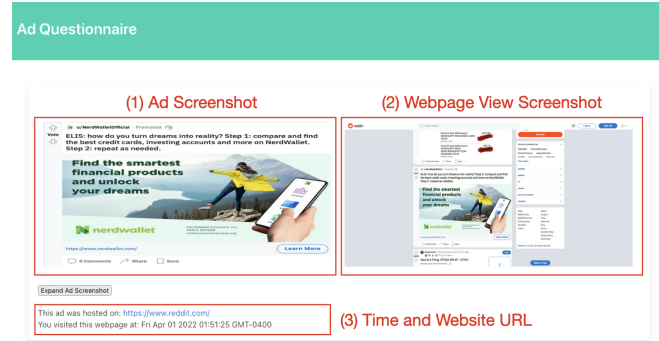


Figure 2. Screenshot of an ESM prompt for a sampled online ad and its associated webpage view.

3.1.1. Experience Sampling Method (ESM). We used the ESM to capture users’ perceptions towards specific ads they encountered in their day-to-day web browsing, as well as their desire for explanations thereof. We implemented a Firefox browser extension to record online ads that participants encountered during the four-week study period¹. The extension logs advertisement information (e.g., screenshot, website title, URL) to generate ESM prompts.

Up to six times per day, participants would be prompted to complete a questionnaire with at least an hour between prompts. Prompts were delivered by in-browser notifications or through the extension interface if the notification was missed. The ESM questionnaire asked for a description of the ad and for self-reported contextual and perceptual measures (e.g., browsing activity, creepiness). To refresh participants’ memory, the questionnaire presented the ad and the webpage they viewed (Figure 2). If participants could not recall the ad, they were asked to skip the questionnaire. To ensure that we only collected reactions to ads that were fresh on participants’ minds, the questionnaire could only be answered within two hours after the participant was first exposed to the ad. Each ESM questionnaire consisted of eleven multiple-choice items and three free-text response questions (Appendix B).

Browser extension: The browser extension: 1) detected ads, captured a screenshot of the viewing context, and documented associated meta-information of the hosting websites, and 2) collected participants’ self-report ESM responses. Ads were identified by referencing the EasyList² as in prior studies [6]. The extension captures a screenshot of each detected ad and participants’ current active webpage view via the `CanvasRenderingContext2D.drawWindow()`³ and `tabs.captureTab()`⁴ APIs, along with the current time, title and URL of the website. Due to privacy reasons, when participants submitted ESM prompts, the extension discarded the corresponding screenshot of the viewing con-

1. A Chrome extension was infeasible due to lack of support for the `CanvasRenderingContext2D` API and the “cross-domain policy,” prohibiting the logging of most ads [6].

2. <https://easylist.to/>

3. <https://developer.mozilla.org/en-US/docs/Web/API/CanvasRenderingContext2D/drawWindow>

4. <https://developer.mozilla.org/en-US/docs/Mozilla/Add-ons/WebExtensions/API/tabs/captureTab>

text of the ad. The research browser extension was tested in a one-week pilot study with eight participants to ensure the robustness of the system before we ran the actual study.

3.2. Data collected in the field study

To model the relationship between contextual factors, perceptual factors, and user factors as they relate to users' desires for ad targeting explanations, we collected the following data from each ESM questionnaire and from an entry questionnaire, which participants filled out at the outset of the study. As described in Section 2, we collected this data, in part, because prior work suggests that these factors correlate with users' general reactions to OBA.

3.2.1. Contextual Factors. We hypothesized that factors relating to the user's ad viewing context should correlate with a desire for ad targeting explanations — e.g., the content of an ad, the user's physical location, and the browsing activity in which they were engaged at the time of viewing the ad. In practice, many of these measures could be automatically inferred. However, to avoid intrusive data collection and to ensure high-fidelity data, we asked participants to self-report these measures in their ESM responses.

Ad topic: Prior research has noted a strong correlation between the topical content of an online ad and people's perceptions towards that ad [6], [32]. We asked participants to describe the topic of the ad through a free-text response. We later categorized these free-text responses into distinct categories adapted from a comprehensive ad topic list by Zeng et al. [32]. In total, we categorized the collected ads into 46 topics (e.g., Apparel, Insurance), including four ad topics that were not considered in the labels from Zeng et al. [ibid] (e.g., Place to Shop, Mixture; see Appendix Table 3). To ensure the reliability of the labeling process, two researchers coded 12% of the data independently, and they achieved an inter-coder agreement of 89%. The two discussed and resolved all disagreements between their labels until agreements were reached, and one researcher coded for the rest of the data.

Location and engaged browsing activity: Prior research has also noted that ad targeting that makes use of contextually identifying information — e.g., physical world location or browsing activity — can be perceived negatively [3], [30]. We asked participants to self-report their location and browsing activity with two multiple-choice questions. For location, we provided a set of pre-defined common thematic locations for personal computer usage: *Home*, *Work*, *Public*, and *Other* (free text). For browsing activity, we provided a set of pre-defined common web browsing activities following prior work by Sellen et al. [33]: *Fact finding and looking for specific information*, *Information gathering and researching some broader topic*, *Pure browsing for self's routine/habit/passing time/entertainment*, *Transactions*, *Communications*, *Maintaining and housekeeping the upkeep of web resources*, and *Other* (free text). We categorized free-text answers into the aforementioned categories when appropriate (e.g., label "Checking status of

package delivery" as *Fact finding and looking for specific information*). To do so, two researchers labeled the free-text responses individually, resolving disagreements as they arose through discussions.

3.2.2. Perceptual Factors. Prior studies have shown how one's attitudes towards and understanding of OBA are influenced by one's perceptions and subjective expectations [3], [5]. To that end, we hypothesized that subjective perceptions — e.g., users' beliefs about ads being creepy, or their beliefs about what data were used to target them — would correlate with the desire for ad explanations.

Unexpectedness, creepiness, and emotional response: For each targeted ad we asked about, we collected participants' perceptions of unexpectedness [31], [34], creepiness [3], [35], as well as their general emotional response to the ad [32], [36]. For unexpectedness and creepiness, participants responded to the statements "*I found this ad to be unexpected.*" and "*I found this ad to be creepy.*", respectively, on a five-point Likert scale ranging from "*strongly disagree*" (1) to "*strongly agree*" (5). For emotional response, we used the Self-Assessment-Manikin Scales (SAM) [37], which comprises the subscales of Valence, Arousal, and Dominance (Appendix B Q6). In presenting these subscales, we followed examples from prior work (e.g., [38]), and confirmed that participants understood how to select SAM figures that best reflected their feelings when they saw the ad through a pilot study.

Perceived information use: We asked participants to report on the type of information they felt was used to target them with a specific ad. We provided participants with a pick list of five options adapted from Leon et al. [39]: 1) browsing information (e.g., pages visited); 2) demographic information (e.g., age); 3) location information (e.g., ZIP code); 4) personal identification information (e.g., email address); and, 5) computer information (e.g., operating system). Participants were asked to select all information types that applied to the ad. We also added an "Other" item and allowed participants to provide free-text responses, and a "None" item to reflect cases when they felt that none of their personal information was used to serve the ad. For the purpose of analysis, we coded the Other responses into corresponded types based on their free-text responses.

Furthermore, we initially surveyed participants on their *willingness* to disclose information associated with the categories proposed by Leon et al. [ibid] for general OBA purposes. This measure was captured on a five-point Likert scale from "Strongly Disagree" (1) to "Strongly Agree" (5). We later binarized these measures to an *Unwillingness to Share* factor for each ad. In the set of categories of information felt to be used for ad targeting, if any categories were rated between "1" or "2," we labeled an ad as using information the participant was "Unwilling to Share" for OBA purposes.

3.2.3. User Factors. We hypothesized that participants' personality traits and their general attitudes toward security and privacy would affect their baseline desire to see ad

explanations. Accordingly, at the outset of the study, we had participants fill out an entry questionnaire to account for these effects. The questionnaire consisted of validated measurement instruments from prior work.

Personality traits: We measured participants' personality traits, which have been shown to correlate with information privacy concerns [40]–[42]. We measure participants' personality traits via a 10-item version of the Big Five Inventory [43], which are broken down into five dimensions: Extraversion, Agreeableness, Conscientiousness, Neuroticism, and Openness.

Security behavioral intention and privacy concern: We also measure participants' baseline attitudes and behaviors for privacy and security via the Internet Users' Information Privacy Concerns (UIIPC) [44] and Security Behavioral Intention Scale (SeBIS) [45]. The former reflects one's privacy attitudes and behaviors using ten items on a seven-point Likert scale; the latter reflects one's intention to adhere to expert-recommended security advice using 16-items on a five-point Likert scale.

3.2.4. Desire for Ad Targeting Explanations. Finally, the key dependent variable of interest in this study: *Desire for Ad Targeting Explanation* was collected as a binary yes/no value via the following question in each ESM questionnaire: “*Would you like to know more about why you were specifically targeted by this ad?*”

3.3. Exit Interview

To answer **RQ2** and understand why or why not participants were interested in receiving an explanation for a given targeted ad, we conducted semi-structured interviews after their participation in the four-week field study. Using participants' field study responses as a guide, the interviewer asked participants to explain their questionnaire responses (e.g., Why are you interested/ not interested in the targeting of this ad?). Each interview lasted 40-60 minutes, covering at least ten responses the participant provided during the study. Specifically, we selected some of their most recent ads to which they reported having stronger emotional reactions: e.g., ads that they felt were creepy. To refresh participants' memories, we showed participants both ad images and their ESM questionnaire responses associated with a specific line of questions as we interviewed them, and asked questions that helped interviewees contextualize and recall the ad (e.g., “What were you doing on the internet when you saw the ad?”). The interview protocol is in Appendix B.

3.4. Study Setup and Recruitment

Participants first completed the entry questionnaire to measure their personality traits, their general attitudes towards security and privacy, as well as their general willingness to disclose personal information to OBA. Then, they attended a remote system installation meeting with the research team, in which we introduced the study and installed the research browser extension on their browser.

The four-week ESM field study started immediately after the initial meeting. At the end of the field study, a subset of participants were invited to an exit interview with the research team. Participants qualified for the exit interview if they provided at least on average one ESM response per day during the field study. We invited qualified participants to the exit interview until we reached saturation — i.e., we stopped once we reached a point where new participants were not providing new insights. We interviewed a total of 36 participants, with other participants attending a shorter debriefing session to share their general study experience and to remove the research extension.

We recruited participants from Prolific⁵, an online crowdsourcing platform. We recruited 62 participants at first, but two of them (P6 and P51) terminated their participation right after they started. Of the remaining 60 participants, 58 participated throughout the four weeks, and two of them terminated their participation earlier. The 60 participants (30 men, 30 women) came from diverse age groups ($M=33.47;SD=9.16$) and backgrounds (Appendix Table 4). Due to limitations of the WebExtension API (Section 3.1.1) only Firefox users were recruited. Most participants reported not using an adblocker (53 out of 60), and all participants were informed to turn off their adblocker(s) during the study.

Participants were compensated with \$13 for the entry questionnaire, and received \$0.5 for each ESM prompt responded during the four-week ESM study, up to \$84, and an additional \$15 for attending the post-study interview or \$5 for the debriefing session.

To ensure participants' anonymity in the study, we did not collect any personally identifiable information from them, and we only communicated and arranged study logistics via the pseudo-anonymous messenger provided by Prolific — i.e., we could only see the unique participant IDs generated by Prolific, which we used to refer to specific participants throughout the study. We informed and received consent from participants to video and audio record the exit interview. The audio recordings were later transcribed for data analysis in a de-identified manner. Moreover, all data we collected from participants were encrypted in transit and stored in a secure server. The study protocol was also approved by our institution's IRB.

3.5. Analysis Procedure

We employed a mixed-methods approach to analyzing our data, guided by our research questions. We answered **RQ1** — when do people want ad targeting explanations? — through a quantitative analysis on the entry and ESM questionnaire data. We answered **RQ2** — why do people want ad targeting explanations? — using qualitative analysis on the exit interviews.

3.5.1. Quantitative Analysis. In total, our participants submitted 5,171 ESM responses, with a response rate of 70%. To reduce redundancy in the dataset, we removed 119

5. <https://prolific.co/>

responses for ads that were asked about more than once in the same hour. We also removed 56 responses on interface elements that our extension mistakenly identified as ads. Additionally, we removed the first three days of data for all participants to reduce novelty effects as recommended by prior work [46].

In all, we retained a total of 4,251 ESM responses for the quantitative analysis, with a per-participant average of 74 (SD=54). Overall, our participants wanted explanations for 1,281 (30%) of the ads that were brought to their attention during the study, with a per-participant average of 33% (SD=29). The top five ad topics represented in our dataset were Apparel (364, 9%), B2B products (362, 9%), Entertainment (343, 8%), Consumer Tech & Tech Services (309, 7%), and Food and Drink (222, 5%). The most frequently reported browsing activity when encountering these targeted ads was general web browsing (2,882, 68%). The most frequently reported location was at home (3,738, 88%).

To model the relationship between perceptual, contextual and user factors correlated with whether or not users wanted ad explanations, we fit four random-intercepts logistic regression models, one corresponding to each of the aforementioned factors and one combining all factors — the Context Model, the Perception Model, the User Model, and the Combined Model. To account for repeated measures, we included Participant ID as a random intercept term. For all categorical variables, we selected the most common factor level as the baseline reference. We pre-registered our random-intercepts logistic regression analysis using AsPredicted⁶ before collecting and analyzing the dataset.

3.5.2. Qualitative Analysis. Two researchers conducted the exit interviews. As mentioned, the interview sessions were audio recorded and transcribed for data analysis. All recordings were transcribed with Otter.ai⁷. One member of the research team validated and corrected the transcripts.

We conducted open coding (e.g., [47]) on participants' explanations of and elaborations on why they did or did not want ad targeting explanations for specific ads. One researcher performed the initial coding on 26 of the interview transcripts, and iteratively constructed a codebook in active discussion with two other researchers. The other two researchers coded the remaining ten transcripts with this codebook. Note that data would also be reviewed and re-coded each time the codebook was revised. The final codebook included reasons for participants *wanting* ad targeting explanations such as unexpectedness (e.g., the unexpected use of physical location for ad targeting), preconceptions (e.g., listening to physical-world conversations), and ad characteristics (e.g., fake ad or clickbait); and, reasons for participants *not* wanting ad targeting explanations such as confidence (e.g., confident in ad targeting demographic information), indifference (e.g., uninteresting ad contents), and helplessness (e.g., lack of accountability). We included the codebook in Appendix Table 2. We used the Dedoose⁸

6. https://aspredicted.org/blind.php?x=MTD_GH1

7. <https://otter.ai/>

8. <http://www.dedoose.com/>

TABLE 1. NON-STANDARDIZED COEFFICIENTS OF THE MIXED-EFFECTS LOGISTIC REGRESSIONS MODELING DESIRE FOR AD TARGETING EXPLANATIONS AGAINST CONTEXTUAL, PERCEPTUAL, USER-LEVEL, AND COMBINED FACTORS. SUBJECTIVE PERCEPTIONS WERE THE MOST PREDICTIVE OF PARTICIPANTS' DESIRE FOR EXPLANATIONS.

Model	M (SD) / Distribution	Context	User	Perception	Combined
(Pseudo) r square/ conditional r square		0.56	0.49	0.69	0.72
Intercept		-1.26	-1.13***	0.28	-0.02
Contextual Factors					
Ad Topic					
Apparel	8.57%	0 r			0 r
B2B Products	8.52%	0.71***			0.86**
Cell & Internet Service	2.33%	-0.83*			-0.41
Dating	0.19%	2.53**			1.54
Finance & Investment Pitch	4.33%	0.71**			0.54
Household Products	4.89%	0.28			0.59*
Humanitarian	1.11%	0.99*			0.90
Med Services & Rx	1.48%	1.27***			0.63
Miscellaneous	1.76%	1.06**			0.40
Public Relations	0.45%	1.36*			1.61*
Youtube Merch	2.35%	-1.10*			-0.31
Browsing Activity					
Pure Browsing	67.80%	0 r			0 r
Communications	5.76%	0.09			-0.06
Fact Finding	15.10%	-0.31*			-0.22
Information Gathering	8.12%	-0.48**			-0.25
Maintaining Systems	0.40%	0.57			0.58
Transactions	2.82%	-0.57			-0.10
Location					
Home	87.93%	0 r			0 r
Public	1.51%	0.33			-0.24
Work	10.56%	0.23			0.07
Time (hour)	13.20 (5.13)	-0.03			-0.01
Is Weekend	25.05%	0.004			0.07
User Factors					
Extraversion	5.85 (2.14)		0.07		0.01
Agreeableness	6.84 (1.66)		-0.22		-0.21
Conscientiousness	8.23 (1.75)		0.37		0.48
Neuroticism	5.88 (2.41)		0.38		0.25
Openness	7.33 (1.76)		0.33		0.05
IUIPC	16.56 (7.69)		-0.53*		-0.42
SeBIS	58.47 (7.79)		-0.16		0.06
Perceptual Factors					
Valence	3.05 (0.74)			0.38***	0.39***
Arousal	2.09 (1.06)			0.31***	0.33***
Dominance	2.79 (0.94)			0.09	0.08
Unexpectedness	2.15 (1.30)			1.20***	1.20***
Creepiness	1.62 (0.95)			0.43***	0.44***
Perceived Ad Targeting					
Browsing Info.	67.70%			0.05	0.11
Demographic Info.	37.80%			0.39**	0.43**
Location Info.	40.34%			0.12	0.17
Personal ID Info.	16.16%			0.06	0.06
Computer Info.	19.90%			-0.32	-0.27
Unwillingness to Share	40.32%			-0.20	-0.26

Significance: * p<.05; ** p<.01; *** p<.001; r: reference

software to code the transcripts.

4. RQ1: When do people want ad targeting explanations?

Overall, participants wanted ad targeting explanations for approximately 30% (1,281 out of 4,251) of the ads that were brought to their attention in our study. Through our quantitative modeling analysis, we found that perceptual factors were the most strongly correlated with participants' desire for ad targeting explanations, followed by contextual factors. User factors had little correlation.

Table 1 summarizes the four models we fit to the data. For numeric factors (i.e., Big 5, IUIPC, SeBIS), a positive coefficient implies that the log odds that a participant

wanted an explanation for a given ad is predicted to increase for every one standard deviation increase of that factor. A negative coefficient implies the opposite. For five-point Likert scale factors (i.e., perceptions of creepiness), a positive coefficient implies that the log odds that a participant wanted an explanation is predicted to increase for every one-point increase above the neutral score (3), and a negative coefficient implies the opposite. For categorical and binary factors (i.e., perceived ad targeting, ad topic, location, and browsing activity), the model coefficient presents the predicted difference in log odds that a participant wanted ad explanations for a given factor level relative to a baseline level ('false' for binary factors). Positive coefficients would imply an increased desire for ad explanations relative to the reference level and vice versa.

4.1. Model Comparisons

We first explored the relative model fit across all four models to understand how well perceptual, contextual, and user-level factors correlated with participants' desire for ad targeting explanations. Table 1 shows the pseudo- R^2 value for each model, with higher values indicating better model fits. Our results show that perceptual factors ($R^2=0.69$) outperform both contextual factors ($R^2=0.56$) and user factors ($R^2=0.49$). Furthermore, when all the three factors were combined in the Combined Model, the R^2 (0.72) increased only modestly compared to the Perception Model. It is noteworthy that all significant effects in the Perception Model remained in the Combined model. In short, our findings suggest that perceptual factors were the most strongly correlated with participants' desire for ad targeting explanations.

Nevertheless, it is worth exploring the significant effects in the User and Context models to better understand how user-level and contextual factors correlate with a broader desire for ad targeting explanations, even if perceptual factors dwarf these effects. In practice, after all, measuring perceptual factors is much harder than measuring contextual and user-level factors.

In the Context Model, we found unique main effects for ad topic and browsing activity. For ad topic, we selected the most common ad topic in the dataset — Apparel — as the baseline reference. Ad topics with positive main effects suggest that users' were more interested in explanations for ads about those topics than they were for ads about Apparel. From Table 1, we can see that most ad topics had statistically significant coefficients, suggesting that ad topic does appear to explain some of the variances in participants' desire for ad targeting explanations. Specifically, ads about Dating, Medical Services & Rx, and Public Relations had the highest coefficients, suggesting that participants were much more likely to want explanations for those ad topics than for Apparel. Conversely, ads about YouTube Merch and Cell & Internet Service had the most negative coefficients, suggesting that participants were much less likely to want explanations about those ads than Apparel.

We will discuss participant rationales in more detail in the next section that encompasses our qualitative analysis.

One possible explanation for why participants were more likely to want ad explanations for these topics, however, is that ads about Dating and Medical Services likely require sensitive information about participants' relationships and health, respectively.

For browsing activity, we selected the most commonly reported — pure browsing — as the reference level, and found negative main effects for the fact finding and information gathering activities, suggesting that participants were much less likely to want explanations for ads they encountered during these activities than they were for ads they encountered during general web browsing. One explanation, here, is that ads encountered during these activities were more topically relevant to the task at-hand.

In the User Model, we found a negative main effect for privacy concerns as measured by the UIIPC: i.e., our results show that those with greater privacy concern were less likely to want ad targeting explanations. One possible explanation is that users with higher UIIPC scores may have more knowledge about how OBA works, in general, resulting in less general interest in ad targeting explanations. This finding resonates with Boerman et al.'s [1] findings on the relationship between knowledge of marketing tactics and users' perceptions about the efficacy of OBA: "*The more people think they know about how OBA works (i.e., subjective persuasion knowledge), the more they tend to overestimate the effects of OBA on others and underestimate its effects on themselves.*"

While the Context and User models had significant effects when studied in isolation, many of these effects were no longer significant with the inclusion of perceptual factors. One reason for this is that the context and user-level factors are likely correlated with perceptual factors — e.g., privacy concern likely correlates with perceptions of creepiness and unexpectedness; ad topic likely correlates with emotional responses (medical ads vs. ads about household products). We will next take a closer look at the relative effects of the different factors in the Combined Model.

4.2. Factors that predict a desire for ad targeting explanations

Generally speaking, we found significant correlations mostly between perceptual and contextual factors (Table 1) when including all factors in the Combined Model. We discuss key findings in each type of factors, in turn.

Perceptual factors: Users' subjective perceptions of targeted ads were highly correlated with their desire for ad targeting explanations. In line with prior work on ad strategy and effectiveness with emotional response [36], we found that emotional valence and arousal correlate with participants' desire for ad targeting explanations (with coefficient $\beta_{valence}=0.4$; and $\beta_{arousal}=0.33$). We also found positive correlations for perceptions of creepiness and unexpectedness ($\beta_{creepiness}=0.44$; $\beta_{unexpectedness}=1.2$), both of which have been identified as common negative user reactions to OBA [3], and as factors that underlie people's perceptions of privacy violations [34], [48]. Figure 3 shows how as




Figure 3. Distribution of desire for targeting explanations (%) as they varied by perceived unexpectedness and creepiness. Higher perceived unexpectedness and creepiness correlated strongly with an increased desire for explanations. X-axis labels: SD=Strongly disagree; D=Disagree; N=Neither agree nor disagree; A=Agree; SA=Strongly agree

self-reported perceptions of unexpectedness and creepiness increased, so too did participants’ desire for explanations.

We did not find a main effect between users’ unwillingness to share information that they *believed* was used for ad targeting and their desire for ad targeting explanations. We did, however, observe a main effect correlating participants’ beliefs that demographic information was being used for targeting and their desire for explanations ($\beta_{demographic}=0.43$). This finding suggests that the perceived use of intrusive data, alone, does not necessarily correlate with users’ desire for ad targeting explanations. However, as we shall see in our qualitative analysis, when users had questions or preconceptions about *how* their personal data were collected and shared, they wanted ad targeting explanations. Relative to other forms of information, participants may have been less clear about how their demographic information was known to advertisers.

Contextual factors: When combined with perceptual and user factors, the only significant main effect among contextual factors was ad topic. Specifically, three ad topics, B2B Products, Public Relations and Household Products, had positive main effects ($\beta_{B2B}=0.86$; $\beta_{PR}=1.61$; $\beta_{household}=0.59$) — i.e., people were more likely to want explanations for these ad topics than the baseline topic of Apparel even after accounting for the effects of subjective perceptions. In short, participants are generally more likely to want explanations for these ad topics, all other things being equal.

User factors: We did not observe any main effect correlating user personality traits and attitudes towards security and privacy (i.e., Big 5, IUIPC, SeBIS) with a desire for ad targeting explanations. Apart from the fact that perceptual and contextual factors outperformed user factors, one possible explanation for this result is that these attributes may be partially correlated with subjective perceptions; accordingly, the variances that would have been explained by these general user-level factors may already be captured by the inclusion of perceptual and contextual factors in the Combined Model.

In sum, these findings help answer **RQ1** — “*When do people want ad targeting explanations?*” We found that users wanted explanations for about 30% of the ads that were brought to their attention (overall: 30%; per-participant

average: 33%), and that their subjective perceptions of an ad — whether an ad elicits a strong emotional response and whether they find it to be creepy or unexpected — are the strongest predictor of whether or not they want explanations. These perceptual factors far outweigh user-level and even contextual factors. Ad topic also appears to be correlated with desire for ad targeting explanations, as does the type of information that users perceive may be used for targeting a particular ad towards them.

5. RQ2: Why do users want or not want ad targeting explanations?

Our quantitative analysis helped us understand *when* participants wanted ad targeting explanations; we next qualitatively analyzed exit interview transcripts to understand *why*. We constructed — to our knowledge — the first systematic model to explain why users want targeting explanations for actual ads they encounter in their day-to-day web browsing, and uncovered six key reasons for why our participants wanted ad targeting explanations and four key reasons for why our participants did *not* want ad targeting explanations.

5.1. Why do users want ad targeting explanations?

Participants wanted ad targeting explanations because 1) they believed that the data used to target them pertained to their offline or online lives in ways they did not expect, 2) they wanted to confirm pre-conceived hypotheses of how they were targeted, 3) they felt that the targeting was grossly inaccurate or largely divergent from expectations, 4) they considered a specific ad abnormal or strange, 5) they noticed intriguing patterns across multiple targeted ads, and 6) they wanted agency over how their data was used for targeting.

5.1.1. Unexpected Use of Personal Data. Participants often wanted ad targeting explanations in order to understand how their personal data — commonly demographic and/or behavioral information — was collected and/or used for a specific advertisement.

Participants commonly expressed this reasoning when information about their physical world lives was exhibited in an ad, such as their **physical location**: “*I want to know how they know my location because I don’t remember that website ever asking me for permission.*” (P18) In a similar vein, P4 believed that her **demographic information** was used to target her with an ad for women’s shoes, and wanted to know how advertisers were able to collect that data: “*I want to know, if it was particularly targeted toward me [using] demographic information. How did they collect that information? How are they aware?*” Other participants wanted to know how advertisers inferred their **interests** such as shoe design preferences (P31). Participants also expressed concern that advertisers were able to access **sensitive, identifying information** such as medical history (P32).

Beyond information about their physical world lives, participants also wondered how advertisers came to know

of their online activities. We identified two major causes for such uncertainty: data provenance and data flow. With respect to **data provenance**, participants wanted to know how their online activities were collected for ad targeting. For example, P23 was curious about how YouTube was able to serve an ad about national grants and stimulus, which was highly related to his recent web searches. Similarly, P28 was surprised an ad about razors was so tightly linked to his recent browsing activities: *“how’d they know so much about me? I was looking up like razors and then I saw this”*. With respect to **data flow**, participants wanted to know more about how their data was shared between different web services and advertisers in the backend: *“[...] based on the fact that these are both Google. So, it was from my Gmail account to YouTube. So I’m sort of curious as to how much Google is compiling about me and how they’re passing that along to their partners”* (P46).

5.1.2. Testing Preconceptions. Another reason participants wanted ad targeting explanations was to *test preconceptions* about how ad targeting works and why they were targeted by a specific ad. Participants implicitly or explicitly constructed three types of preconceptions about: how personal data was used for targeting, the motivation and strategy of advertisers, and how personal data was tracked.

Preconceptions about how personal data was used for targeting are guesses about how one’s personal data was used when being targeted by a specific ad. Some participants harbored notions of how their **browsing history and activity** might have impacted the ads they saw: *“I’m wondering if they’re... tracking my purchasing activity, buying vape products and going to vaping related sites”* (P32). Similarly, when seeing an ad from Reddit about hair loss treatment, P27 explained that Reddit inferred he was bald *“because they have access to my subscriptions on Reddit”*.

Participants also held preconceptions about how advertisers utilized information about their physical-world lives and experiences — such as their age and location. For example, P26 wondered: *“are they targeting me because I’m young?”*. P18 wondered if he was targeted with an ad about movers because he was living in a college town: *“[I’m] maybe wondering if it’s using some kind of location information, because I live in a college town. So there’s always people kind of moving in and outside and maybe that was why it was targeting.”*

Preconceptions about the motives and strategies of advertisers were also common reasons for wanting ad explanations. For example, P31 wanted to assess his hypothesis that game manufacturers advertise broadly close to new game launch dates: *“A lot of times, they’re products that I’ve never seen before. So I’m just wondering if they throw these out when they launch new products?”* Likewise, P23 wanted to know more when he was targeted by non-“mainstream” brands, believing that he must have been very specifically targeted: *“a key factor is something that doesn’t appear like mainstream. Like if it was something like Google on here, or some big company, I wouldn’t think anything strange*

about it... doesn’t seem like a lot of people would get [this brand]... makes me wonder why I’m getting this.”

The two preconceptions mentioned above are, to our knowledge, novel in the literature. We also identified **Preconceptions about how personal data was collected and shared** that extend and provide supporting evidence for the concepts relating to folk models of OBA identified by Yao et al. [5]. Specifically, we found that when our participants harbored one of the following folk mental models about how their data was collected and shared, they wanted explanations to confirm or deny their understanding:

Data breaches of, and sharing by, first-party websites is similar to Yao et al.’s [ibid] connected first-party folk model, in which participants believed that websites and applications that they had directly interacted with shared their data with others for ad targeting. For example, P13 stated *“I have looked at different survey sites. So maybe they all like sell information or sell the cookies, or whatever they do. And it’s all just a big circular round and round selling information fast”*. Similarly, P10 wanted an explanation because he was concerned about what, exactly, was shared among these first-party websites: *“if I watched that another live show on Netflix on my laptop, are they pulling that data? Or did Netflix sell that data to YouTube? And, you know, YouTube buddies? Like they’re all interconnected these days?”*

Cross-device tracking could be seen as an extension of Yao et al.’s [ibid] third-party folk model, in which participants explicitly expressed their preconception of how their online activities on different devices were tracked and used by ad targeting across devices. For example, P24 stated: *“are they tracking me with my phone to serve ads on my computer at home. And what is tracking me on my phone? What is tracking me on my phone that’s telling them all behind the scenes to serve me ads at home on the MacBook”*. P18 wanted an ad targeting explanation to test his hypothesis that because he had used Facebook on multiple devices, it must be the culprit in sharing information about him to advertisers: *“Because I shopped on this website and I use multiple devices for Facebook, I don’t just use the same computer, I use my phone, I have an iPad, and I’ve seen this ad on all of my different devices. So I think that would be why I’d want to know, like how is it following me on every device I have.”*

Tracking via spyware and malware refers to preconceptions that personal information used for ad targeting is collected and tracked via malicious software running on participants’ computers. This preconception builds on misconceptions and speculations about OBA trackers identified by Yao et al. [ibid]. For example, P21 expressed concern that websites may install spyware on his computers to facilitate the tracking and sell his personal data to institutional actors like AT&T: *“So it might be that there’s like spyware on my computer. And I’m targeted because somebody is seeing exactly what I’m doing.”* P1, after seeing an ad that she felt was different from other mainstream ads, was suspicious of malware. This suspicion led her to wanting an explanation: *“Is this really targeting me to click on the ad and maybe give me a virus or something? Or is it gonna gather even*

more confidential information about me if I click that?”

We also found a folk model not previously discussed by Yao et al. [ibid], *listening to physical-world conversations*, or preconceptions that participants’ had about their conversations in the physical world being collected and used for ad targeting. “I would like to know, just because I hadn’t used anything on my computer or my browser that relates to these tacos. But I did like, by word of mouth and through my phone... So I was curious about why I was being shown it on my computer, and what information was, like crossed over for me to see it” (P17). Participants appeared to develop this preconception after they ruled out other alternatives: “there were definitely a couple of them where I wondered because I couldn’t recall looking for something... And something about it came up... I don’t know if it can hear me or if it was that was just in my head” (P31).

5.1.3. Unexpected Targeting. Participants also wanted explanations when they encountered *ad targeting that diverged from their expectations*: for example, when they encountered ads that were misaligned with their perception of their own digital footprints, their actual interests, or their actual needs. Indeed, prior work suggests that many users believe ad targeting algorithms to be smart, useful, and accurate [3], [7]; our findings suggest that when users encounter evidence that contradicts these beliefs, they want to know what led the algorithm astray.

Mismatches with self-aware digital footprints were identified when participants encountered ads that diverged from or even contradicted their recent online activities. For example, when seeing an ad for a health supplement, P21 stated: “Well, I was not searching for anything health-related and recently, nor is I watching any health-based videos, which is why I was wondering why I was sent this.” P4, in contrast, was impressed by when an ad she encountered matched her interests because she could not recall searching for anything pertinent: “Yeah, I was interested in this [ad]...but I’m not sure why I saw this ad in particular, I’d never searched for this.”

Mismatch with interests and conceptions of self were identified when participants were served with ads that did not fit their actual interests or profiles, and saw such targeting as “wrong” (P12) or “a glitch” (P34). For example, P23 and P1 expressed interests in knowing why they were seeing ads promoting items that they “didn’t even want” and were “not interested in”. P10 also expressed a desire for ad explanations when he saw an ad about women’s underwear: “that’s shitty advertising there.” Other participants viewed the mismatch as offensive, driving them to want more information for why the targeting occurred. For example, P35 felt insulted when they were targeted with an ad for free groceries because of the implication that they would need the service: “...this being targeted at me, is just to me insulting, and I would be interested to know what is drawing this ad to me.”

5.1.4. Abnormal Ad Characteristics. Another reason participants wanted ad targeting explanations related to *abnor-*

mal ad characteristics — i.e., when participants believed the content or design of an ad to be unusual compared to their other ads. Prior work has studied what makes ads bad or problematic [6], [32]. Our results extend those findings and provide insights on how negative perceptions of ads factor into users’ desire for explanations. We also find that even ads that are not perceived as “bad” might trigger participants’ desire for explanations.

Participants brought up ads being unclear or poorly designed, or otherwise bad as a factor for wanting ad targeting explanations. For example, when discussing why she wanted an explanation for an unclear ad, P1 explained: “... this one is just general... It’s very vague... In the sense of like, a more intrigue about it, ... it interested me”. P25, on the other hand, mentioned an ad was not only unclear, but also looked **fake** and like **clickbait**: “[the ad] doesn’t even look like a real thing... it looks like one of those fake ads, ... I don’t know why I would be getting this random ad”.

5.1.5. Intriguing Ad Patterns. People also wanted ad targeting explanations when they uncovered *intriguing or troubling patterns* about the context in which an ad was delivered or across multiple ads that they encountered. In other words, their desire for an explanation had less to do with the ad itself and more to do with how they situated the ad in their broader experience with ads. To date, prior work has focused on user reactions to specific targeted ads or to targeted ads in general, but little attention has been paid to how users perceive a specific ad as interplaying with a broader collection of targeted ads, nor the resulting user-perceived implications for privacy.

One such pattern was **repetition**: participants wanted to know why they would keep seeing the same ad: “just not knowing why I keep getting this ad. And I kept getting it so much... And I’m like, Why? Why so much?” (P25). P26 further expressed how such repetitiveness negatively affected his perception of the ad: “at first I [would] maybe... be more curious, but... after seeing [the ad] 10 times, or whatever, it’s like, you just kind of become mad. So I would say I’m curious.”

Other than patterns of repetition, some participants expressed wanting ad targeting explanations owing to **meta patterns** about targeted ads that they found unsettling. P5, for example, elaborated on how she wanted an explanation for an ad she saw on weight loss owing to recent revelations from the Facebook whistleblower Frances Haugen “[Haugen] talked about how if someone searches healthy recipe... the next thing in their algorithm would be for like weight loss. And then it keeps on getting more and more extreme to where it goes into like pro eating disorder sites. So it seems like it’s kind of going in that direction, which I don’t like.”

Ads that were unique relative to other ads or appeared in unique contexts also piqued participants’ interest. For example, the first time P1 saw an ad that she found interesting, she stated: “I wanted to know more about it, because I had never seen the ad, not only through the four weeks, but not even before”. P46 stated that she wanted an ad explanation because she encountered the ad on an unfamiliar website:

“I think it was because this was a new site. And I didn’t know what the privacy sort of protocols were. And so to see advertising there... made me wonder how much they’re targeting. And how much privacy we’re giving up by moving to this format” (P46).

5.1.6. Agency & Autonomy. Finally, some participants also expressed a desire to gain more agency and autonomy (i.e., control over their own information used by ad targeting) through ad targeting explanations. For example, when asked about their main reason for wanting ad explanations, P10 responded: *“I’m just curious on how that drew from all my metadata. How I got this ad, so that I can delete all that stuff. So I never see that.”* In a similar vein, P34 answered: *“So I can do something to prevent this ad from being shown up there again to me... Like I don’t want to actively try to hunt to turn this ad off. I just want it to be known like hey I’m not interested... stop.”* This finding helps bridge prior literature in OBA and human-centered explainable AI (HCXAI): indeed, prior work on HCXAI has found that explanations, alone, are unsatisfactory for users [49], [50] — users want agency over the model or the data therein when explanations reveal faulty reasoning or unsettling data collection practices.

5.2. Q2: Why do users not want ad targeting explanations?

For 70% of the ads they encountered, our participants reported that they did not want targeting explanations. Through our exit interviews, we uncovered four reasons for why — participants did not want ad explanations because 1) they felt confident in their understanding of why they were targeted by a specific ad, 2) they situated an ad within a trend that they recognized, 3) they considered themselves powerless to effect change, and 4) they were indifferent to the ad.

5.2.1. Confident Understanding of How They Were Targeted. The first reason our participants did not want ad targeting explanations was because they reported having a *confident understanding* of how they were targeted with a specific ad. When an ad closely matched participants’ preconceptions of what data advertisers could collect about them and what inferences they could make from that data, they did not want an explanation. For example, P26 stated: *“I guess I kind of already have a good guess as to why I was being targeted by this time...I guess...that data is just location and demographic data. So I wasn’t really curious for that reason.”*

Participants also did not want ad targeting explanations when they perceived a clear link between a targeted ad and their search history, browsing history, online shopping, and other online activities. P4 stated: *“I think they just see that I’m constantly like looking at books on these or Amazon or the library webpage or whatever book site there is. So it seems self-explanatory like this guy reads. So let’s send them, you know, the new New York Times bestseller”.*

Some participants also pieced together how their data might be correlated to people they lived with, such as their spouses. For example, when P29 saw a job advertisement, she associated the ad targeting with her husband’s browsing activities: *“I feel like some of this browsing history seems to be connected to like my husband’s searches. So I feel like it might be more that influenced as opposed to mine.”*

Generally, our participants’ confidence was rooted in their belief that ad targeting is accurate and reasonable. Still, participants believed some ads were not specifically targeted at them: *“it was just targeted to everyone” (P26).* Others attributed an ad to a broader context — P23 believed that he saw a video game ad because Halloween was coming: *“I figured out why I got it just from the time of the year.”* Likewise, P16 believed he saw an ad that accorded with his political ideology *“because I’m a Republican, so I just kind of assumed that’s probably one of the top things they know. And the website once again [a Republican news website]... this led me to make sense I’m target audience.”*

5.2.2. Repetition and Familiarity. The second reason why participants did not want ad targeting explanations was because of *repetition and familiarity*. While some participants stated that repetitiveness was a reason for why they *wanted* explanations, others gradually felt more at ease with particular ads through increased exposure. We note, however, that repetitiveness may not always itself be the underlying reason why users did not want targeting explanations — rather, repetitiveness might correlate with other reasons for not wanting targeting explanations (e.g., users developing folk models for why they might be targeted by a given ad, as we describe in Section 5.2.1).

For example, when explaining her lack of interest in receiving a targeting explanation for an ad she saw frequently, P1 described: *“it’s mainly not only because I see it a lot, so I understand the targeting, but also because it’s a very common all the time type of ad that you see in the websites that I frequently go to.”* Others recognized repetitive exposures as a known marketing strategy: *“the more impressions, the more likelihood, likely they are to get a sale” (P27).* Still others experienced habituation effects, and were no longer bothered by the ad: *“After seeing so many times, you kind of already thought about like, why that ad targeted for this? ... there’s not to say after seeing it so many times, that is really not many additional thoughts that I could get” (P26).*

5.2.3. Helplessness and Resignation. The third reason participants did not want ad explanations was because of a sense of *helplessness and resignation*; they felt that they were powerless to change anything about ad targeting. Participants said, for example, that *“you just get numb to it” (P24)* and *“I have to accept these types of things” (P10).* We also uncovered some underlying reasons that led to these feelings of helplessness and resignation. For example, P10 specifically pointed out the **lack of accountability** for ad targeting, and how ad explanations would still provide him with no agency to make changes: *“not really interested...”*

in how they collected my data, because... they're gonna do it anyways... the only purpose it [an explanation] serves more is just [to] make sure you're understanding where your browsing history is going." P1, on the other hand, simply accepted that her data was already collected and used for ad targeting: *"because I guess that my browsing information is already breached for them. So I'm fine... by being targeted. I don't I kind of don't need more information about it."*

5.2.4. Indifference. The fourth and final reason why our participants did not want ad targeting explanations was *indifference*. Different from *helplessness and resignation*, where participants might have preferred a different state of the world but felt that change was impossible, participants who expressed indifference were not at all concerned about the ad or how it was targeted. P4, for example, stated: *"I wasn't interested in the ad. And I didn't particularly feel interested in knowing why [I] was targeted for this ad."* P28 reported indifference because he was already a customer of a similar product and not considering a change: *"I'm already with RBC...for like eight years. I'm not really gonna switch to a new bank, unless I have a reason to... So I don't really... want to know why I was being targeted."* Other participants were indifferent to explanations because the ad, itself, seemed low-effort or uninteresting. P34 stated, for example, *"it was so basic and just now that I was like, if you're not going to put energy in it, why bother?"*

To summarize, our qualitative analysis helps answer **RQ2** — *"why do people want ad targeting explanations?"* We presented six key reasons why users wanted ad targeting explanations: Unexpected Use of Personal Data, Testing Pre-conception, Unexpected Targeting, Abnormal Ad Characteristics, Intriguing Ad Patterns, and Agency & Autonomy. We also identified four key reasons why participants did *not* want ad targeting explanations: A Confident Understanding of Specific Ad Targeting, Repetition and Familiarity, Helplessness and Resignation, and Indifference.

6. Discussion

6.1. Subjective perceptions determine when and why people want ad targeting explanations

Participants wanted ad targeting explanations for about 30% of the ads that were brought to their attention. Our results suggest that subjective perceptions were the key determinant for *when* and *why* participants wanted these explanations. Indeed, in our regression model correlating the effects of contextual, perceptual, and user-level factors on participants' desire for ad targeting explanations, we found that perceptual factors far outweighed other factors. Participants wanted explanations when, for example, ads elicited strong emotional responses, when they perceived an ad to be unexpected or creepy, when they wanted to confirm preconceptions of how their data was collected and shared, or when they wanted to situate an advertisement within a noticed trend. In other words, participants wanted

explanations when they had idiosyncratic concerns about specific advertisements in order to better understand or confirm their suspicions about how their data was harvested or in order to better understand the motives of advertisers and ad brokers.

6.2. Designing more human-centered ad targeting explanations

Existing approaches to crafting ad targeting explanations are inaccurate, incomplete, and insufficient for end-users to construct an accurate and realistic understanding of how OBA works [5], [8], [51]. This knowledge gap could also prevent users from making effective decisions to protect their own privacy [16]. Our results suggest that these shortcomings may be due, in part, to the fact that existing approaches to crafting explanations are one-size-fits-all, and focus on surface-level properties of the ad and the data that may have been used for targeting [7]: for example, that an ad was meant to be seen by users of a particular age group or who have a particular interest. There is little information in these explanations about the provenance and flow of data, the motives of the advertisers, or how an ad might relate to other ads the user has seen. In short, there is a disconnect between the information provided in the explanations and the information users actually want to see and act on. Based on these findings and building on the vision for human-centered explainable AI put forth by prior work [50], [52], we provide suggestions below on how to craft ad explanations that better cater to what people desire to know.

6.2.1. Personalized and interactive ad explanations. Our results suggest that ad targeting explanations should be personalized and account for the idiosyncratic concerns about ad targeting that might lead users to want those explanations.

Doing so may be beneficial to all involved parties — users, ad brokers, and advertisers. Users benefit from increased transparency. Ad brokers and advertisers benefit from curbing incorrect beliefs about how they collect, use, and process user data, which affects user trust. Indeed, users often harbor inaccurate beliefs about how their data was collected; some of these beliefs may be more nefarious than reality, such as the commonly held belief that the Facebook smartphone app eavesdrops on people's conversations in the physical world [53]. Improving the effectiveness of ad targeting explanations can help ad brokers and advertisers comply with regulatory mandates — such as the "right to explanation" for solely-automated decision making under GDPR. Finally, creating effective explanations can also be seen as a form of competitive differentiation in that users may be more comfortable with advertisers and ad brokers that are effectively transparent over their data collection and processing practices. Indeed, prior work suggests that users often find OBA smart and useful, but dislike covert data collection and opaque targeting algorithms (e.g., [3], [5]).

Proactively accounting for individual preferences for ad explanations based on subjective perceptions can be

difficult. One possibility for addressing this challenge is with a domain-specific question-and-answer chatbot and/or personalized privacy assistant (e.g., [54], [55]). We envision, for example, users being able to ask questions about things they find unsettling or intriguing about a specific ad such as “Which websites or applications shared my location information with the advertiser?” This chatbot could use a combination of the user’s own browsing history, public content (e.g., terms of service), and previous answers to similar questions as source material. It could also opportunistically suggest the use of privacy controls to enhance user agency over their personal data. While ad brokers may not want to make all relevant data available, exploring the design space of question-and-answer-based ad explanations that take into account the perspectives of users, advertisers, and ad brokers presents a ripe opportunity for future work.

6.2.2. Proactive ad explanations. Ad transparency mechanisms are notoriously obtuse. OBA disclosure icons, for example, are commonly misunderstood and ignored by users [3]. Our interview results suggest ad targeting explanations could be designed to be more proactive in scenarios that may elicit heightened end-user privacy concerns: e.g., ads that are targeted based on demographic information such as physical world location. While advertisers and ad brokers may not want to draw attention to their data collection and processing practices, ad explanations may help curb misconceptions that are potentially more unsettling than reality: e.g., such as advertisers eavesdropping on their conversations, or ad brokers sharing personal data with third parties.

Future work might explore identifying “teachable moments” where targeting explanations are actively pushed to end-users. A naive rule-based algorithm could be a good enough baseline approach — e.g., show explanations for ads that use fine-grained location. Future work could also explore allowing users to create their self-defined rules for when they might want ad targeting explanations to mitigate habituation effects. There is also a need to explore more effective explanation presentations, e.g., just-in-time embedded explanations alongside ads. More generally, there is a ripe opportunity for exploring the design space and the effectiveness of proactive ad explanations.

6.2.3. Increasing agency and autonomy. Echoing recommendations from prior work in human-centered explainable AI [49], [50], explanations alone are not enough — participants expressed the desire for agency. Explanations are simply a means to an end. Ultimately, participants wanted avenues to do something about ad targeting and data harvesting practices that they found unsettling, creepy, or even inaccurate. Future work might consider pairing targeting explanations with controls to help users change unsettling data collection or inaccurate inferences. We hypothesize that pairing explanations with privacy controls may result in greater use of OBA controls and decreased user concern.

6.3. Limitations

Our work has several limitations. Our dataset only comprises a subset of target ads that participants encountered on their primary personal computers. There is room, thus, for future work to expand on our investigation by exploring a broader milieu of contexts in which targeted ads may be encountered (e.g., on smartphones). Second, owing to technical limitations, we could only deploy our research extension on the Firefox browser, excluding a large number of potential participants. Nevertheless, we recruited a broad array of participants who varied in demographic backgrounds, privacy concerns, security attitudes, and technical literacy. Third, in order not to overwhelm our participants, we randomly sampled at most six of the ads they encountered on a given day; thus, we did not ask participants about many of the targeted ads they might have encountered during the study period. Fourth, our extension did not distinguish between targeted and untargeted advertisements. While there’s no clear measurement of the percentage of online ads that are targeted, prior studies have suggested that targeted ads generate more revenue than untargeted ads, and that a majority of ads served in the digital advertisement market are targeted [56], [57]. Thus, we assume that most, if not all, of the ads our participants encountered were targeted. Fifth, the contextual factors we collected for our quantitative analysis were self-reported. In theory, many of these factors could be automatically inferred; however, since our goal was to create explanatory models, we opted for self-report to ensure higher fidelity in our data. Finally, our participants were based in either Canada or the United States of America, limiting the cultural scope of our findings and leaving open the potential for cross-cultural work in the future.

7. Conclusion

We conducted a four-week, mixed-methods field study with 60 participants to model *when* and *why* users want ad targeting explanations. We found that users wanted explanations for about 30% of the 4,251 ads we brought to their attention. Subjective perceptions about how their data was collected and shared were highly correlated with when users wanted explanations — more than general attitudes, browsing context, or the nature of the ad itself. But the factors that led to such subjective perceptions varied considerably across users and contexts, making it hard to predict when users will want ad targeting explanations based strictly on general knowledge about the users and the ads. We also found that users typically wanted ad targeting explanations in order to: (i) test hypotheses, (ii) investigate advertiser motivations, or (iii) situate an ad within a noticed trend. These insights suggest that one-size-fits-all approaches to crafting ad targeting explanations are unlikely to be satisfactory to users when they do indeed want targeting explanations. Instead, we envision a future where ad targeting explanations afford users make clearer the provenance and flow of the data used for targeting, the motivations of the advertiser in targeting the user, and controls to help mitigate unwanted collection of personal data.

Acknowledgment

This work was generously supported by the ‘Improve Ad Experiences’ Grant program by Facebook (now Meta). We thank our reviewers for being generous with their time and feedback to help us improve the paper. We also thank Chih-Heng Lin for assisting with system implementation, Avinandan Basu, Upol Ehsan, and Devansh Jatin Ponda for assisting with discussion and ideation early in the project, and Youngwook Do for supporting data analysis.

References


- [1] S. C. Boerman, S. Kruijkeimer, and F. J. Zuiderveen Borgesius, “Online Behavioral Advertising: A Literature Review and Research Agenda,” *Journal of Advertising*, vol. 46, no. 3, pp. 363–376, Jul. 2017.
- [2] P. Verna, “Insider intelligence’s digital advertising trends to watch in 2022,” 2021.
- [3] B. Ur, P. G. Leon, L. F. Cranor, R. Shay, and Y. Wang, “Smart, useful, scary, creepy: perceptions of online behavioral advertising,” in *proceedings of the eighth symposium on usable privacy and security*, 2012, pp. 1–15.
- [4] B. Auxier, L. Rainie, M. Anderson, A. Perrin, M. Kumar, and E. Turner, “Americans and privacy: Concerned, confused and feeling lack of control over their personal information,” 2019.
- [5] Y. Yao, D. Lo Re, and Y. Wang, “Folk models of online behavioral advertising,” in *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing*, 2017, pp. 1957–1969.
- [6] T. K. Eric Zeng and F. Roesner, “Bad news: Clickbait and deceptive ads on news and misinformation websites,” in *Workshop on Technology and Consumer Protection*, 2020.
- [7] M. Eslami, S. R. Krishna Kumaran, C. Sandvig, and K. Karahalios, “Communicating algorithmic process in online behavioral advertising,” in *Proceedings of the 2018 CHI conference on human factors in computing systems*, 2018, pp. 1–13.
- [8] M. Wei, M. Stamos, S. Veys, N. Reitingger, J. Goodman, M. Herman, D. Filipczuk, B. Weinschel, M. L. Mazurek, and B. Ur, “What twitter knows: Characterizing ad targeting practices, user perceptions, and ad explanations through users’ own twitter data,” in *29th {USENIX} Security Symposium ({USENIX} Security 20)*, 2020, pp. 145–162.
- [9] A. Vance, D. Eargle, J. L. Jenkins, C. B. Kirwan, and B. B. Anderson, “The fog of warnings: how non-essential notifications blur with security warnings,” in *Fifteenth Symposium on Usable Privacy and Security (SOUPS 2019)*, 2019, pp. 407–420.
- [10] E. G. Smit, G. Van Noort, and H. A. M. Voorveld, “Understanding online behavioural advertising: User knowledge, privacy concerns and online coping behaviour in Europe,” *Computers in Human Behavior*, vol. 32, pp. 15–22, Mar. 2014.
- [11] T. H. Baek and M. Morimoto, “Stay Away From Me,” *Journal of Advertising*, vol. 41, no. 1, pp. 59–76, Apr. 2012, publisher: Routledge _eprint: <https://doi.org/10.2753/JOA0091-3367410105>.
- [12] C. A. Summers, R. W. Smith, and R. W. Reczek, “An Audience of One: Behaviorally Targeted Ads as Implied Social Labels,” *Journal of Consumer Research*, vol. 43, no. 1, pp. 156–178, Jun. 2016.
- [13] J. Turow, M. X. D. Carpini, N. A. Draper, and R. Howard-Williams, “Americans Roundly Reject Tailored Political Advertising,” p. 30, 2009.
- [14] C. Phelan, C. Lampe, and P. Resnick, “It’s Creepy, But it Doesn’t Bother Me,” in *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, ser. CHI ’16. New York, NY, USA: Association for Computing Machinery, May 2016, pp. 5240–5251.
- [15] A. I. Anton, J. B. Earp, and J. D. Young, “How internet users’ privacy concerns have evolved since 2002,” *IEEE Security Privacy*, vol. 8, no. 1, pp. 21–27, Jan. 2010, conference Name: IEEE Security Privacy.
- [16] A. McDonald and L. F. Cranor, “Beliefs and Behaviors: Internet Users’ Understanding of Behavioral Advertising,” Social Science Research Network, Rochester, NY, SSRN Scholarly Paper ID 1989092, Aug. 2010.
- [17] T.-M. C. Jai, L. D. Burns, and N. J. King, “The effect of behavioral tracking practices on consumers’ shopping evaluations and repurchase intention toward trusted online retailers,” *Computers in Human Behavior*, vol. 29, no. 3, pp. 901–909, May 2013.
- [18] M. K. Ohlhausen and A. P. Okuliar, “Competition, Consumer Protection, and the Right [Approach] to Privacy,” *Antitrust Law Journal*, vol. 80, no. 1, pp. 121–156, 2015.
- [19] J. Sutanto, E. Palme, C.-H. Tan, and C. W. Phang, “Addressing the Personalization-Privacy Paradox: An Empirical Assessment from a Field Experiment on Smartphone Users,” *MIS Quarterly*, vol. 37, no. 4, pp. 1141–1164, 2013, publisher: Management Information Systems Research Center, University of Minnesota.
- [20] A. Bleier and M. Eisenbeiss, “The Importance of Trust for Personalized Online Advertising,” *Journal of Retailing*, vol. 91, no. 3, pp. 390–409, Sep. 2015.
- [21] J. Gomez, T. Pinnick, and A. Soltani, “KnowPrivacy,” Oct. 2009.
- [22] A. M. McDonald and L. F. Cranor, “The Cost of Reading Privacy Policies 2008 Privacy Year in Review,” *I/S: A Journal of Law and Policy for the Information Society*, vol. 4, no. 3, pp. 543–568, 2008.
- [23] C. Jensen and C. Potts, “Privacy policies as decision-making tools: an evaluation of online privacy notices,” in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ser. CHI ’04. New York, NY, USA: Association for Computing Machinery, Apr. 2004, pp. 471–478.
- [24] G. R. Milne, M. J. Culnan, and H. Greene, “A Longitudinal Assessment of Online Privacy Notice Readability,” *Journal of Public Policy & Marketing*, vol. 25, no. 2, pp. 238–249, Sep. 2006, publisher: SAGE Publications Inc.
- [25] H. Marreiros, R. Gomer, M. Vlassopoulos, M. Tonin, and m. c. schraefel, “Exploring user perceptions of online privacy disclosures,” Oct. 2015, num Pages: 12.
- [26] F. Z. Borgesius, “Informed Consent: We Can Do Better to Defend Privacy,” *IEEE Security Privacy*, vol. 13, no. 2, pp. 103–107, Mar. 2015, conference Name: IEEE Security Privacy.
- [27] P. G. Leon, J. Cranshaw, L. F. Cranor, J. Graves, M. Hastak, B. Ur, and G. Xu, “What do online behavioral advertising privacy disclosures communicate to users?” in *Proceedings of the 2012 ACM workshop on Privacy in the electronic society*, 2012, pp. 19–30.
- [28] G. van Noort, E. G. Smit, and H. A. M. Voorveld, “The Online Behavioural Advertising Icon: Two User Studies,” in *Advances in Advertising Research (Vol. IV): The Changing Roles of Advertising*, S. Rosengren, M. Dahlén, and S. Okazaki, Eds. Wiesbaden: Springer Fachmedien, 2013, pp. 365–378.
- [29] A. J. Stanaland, M. O. Lwin, and A. D. Miyazaki, “Online privacy trustmarks,” *Journal of Advertising Research*, vol. 51, no. 3, pp. 511–23, 2011.
- [30] N. a. M. Barbosa, G. Wang, B. Ur, and Y. Wang, “Who am i? a design probe exploring real-time transparency about online and offline user profiling underlying targeted ads,” *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.*, vol. 5, no. 3, sep 2021.
- [31] E. Rader, S. Hautea, and A. Munasinghe, “‘i have a narrow thought process’: Constraints on explanations connecting inferences and self-perceptions,” in *Sixteenth Symposium on Usable Privacy and Security ({SOUPS} 2020)*, 2020, pp. 457–488.
- [32] E. Zeng, T. Kohno, and F. Roesner, “What makes a “bad” ad? user perceptions of problematic online advertising,” in *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, ser. CHI ’21. New York, NY, USA: Association for Computing Machinery, 2021.

- [33] A. J. Sellen, R. Murphy, and K. L. Shaw, "How knowledge workers use the web," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ser. CHI '02. New York, NY, USA: Association for Computing Machinery, 2002, p. 227–234.
- [34] J. Lin, S. Amini, J. I. Hong, N. Sadeh, J. Lindqvist, and J. Zhang, "Expectation and purpose: understanding users' mental models of mobile app privacy through crowdsourcing," in *Proceedings of the 2012 ACM conference on ubiquitous computing*, 2012, pp. 501–510.
- [35] I. Shklovski, S. D. Mainwaring, H. H. Skúladóttir, and H. Borgthorsson, "Leakiness and creepiness in app space: Perceptions of privacy and mobile app use," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 2014, pp. 2347–2356.
- [36] J. D. Morris, "Observations: Sam: the self-assessment manikin; an efficient cross-cultural measurement of emotional response," *Journal of advertising research*, vol. 35, no. 6, pp. 63–68, 1995.
- [37] P. Lang, "Behavioral treatment and bio-behavioral assessment: Computer applications," *Technology in mental health care delivery systems*, pp. 119–137, 1980.
- [38] H.-J. Suk, "Color and emotion- a study on the affective judgment across media and in relation to visual stimuli," Ph.D. dissertation, November 2006. [Online]. Available: <https://madoc.bib.uni-mannheim.de/1336/>
- [39] P. G. Leon, B. Ur, Y. Wang, M. Sleeper, R. Balebako, R. Shay, L. Bauer, M. Christodorescu, and L. F. Cranor, "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*, 2013, pp. 1–12.
- [40] I. Junglas and C. Spitzmuller, "Personality traits and privacy perceptions: An empirical study in the context of location-based services," in *2006 International Conference on Mobile Business*, 2006, pp. 36–36.
- [41] M. L. Korzaan and K. T. Boswell, "The influence of personality traits and information privacy concerns on behavioral intentions," *Journal of Computer Information Systems*, vol. 48, no. 4, pp. 15–24, 2008.
- [42] B. Osatuyi, "Personality traits and information privacy concern on social media platforms," *Journal of Computer Information Systems*, vol. 55, no. 4, pp. 11–19, 2015.
- [43] B. Rammstedt and O. P. John, "Measuring personality in one minute or less: A 10-item short version of the big five inventory in english and german," *Journal of research in Personality*, vol. 41, no. 1, pp. 203–212, 2007.
- [44] N. K. Malhotra, S. S. Kim, and J. Agarwal, "Internet users' information privacy concerns (iupc): The construct, the scale, and a causal model," *Information systems research*, vol. 15, no. 4, pp. 336–355, 2004.
- [45] S. Egelman and E. Peer, "Scaling the security wall: Developing a security behavior intentions scale (sebis)," in *Proceedings of the 33rd annual ACM conference on human factors in computing systems*, 2015, pp. 2873–2882.
- [46] E. M. Redmiles, Y. Acar, S. Fahl, and M. L. Mazurek, "A summary of survey methodology best practices for security and privacy researchers," Tech. Rep., 2017.
- [47] J. Corbin and A. Strauss, "Basics of qualitative research. 3rd edn thousand oaks," 2008.
- [48] D. J. Solove, "A taxonomy of privacy," *U. Pa. l. Rev.*, vol. 154, p. 477, 2005.
- [49] A. Smith-Renner, R. Fan, M. Birchfield, T. Wu, J. Boyd-Graber, D. S. Weld, and L. Findlater, "No explainability without accountability: An empirical study of explanations and feedback in interactive ml," in *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, 2020, pp. 1–13.
- [50] U. Ehsan and M. O. Riedl, "Human-Centered Explainable AI: Towards a Reflective Sociotechnical Approach," in *HCI International 2020 - Late Breaking Papers: Multimodality and Intelligence*, C. Stephanidis, M. Kurosu, H. Degen, and L. Reinerman-Jones, Eds. Cham: Springer International Publishing, 2020, pp. 449–466.
- [51] A. Andreou, G. Venkatadri, O. Goga, K. Gummadi, P. Loiseau, and A. Mislove, "Investigating ad transparency mechanisms in social media: A case study of facebook's explanations," in *NDSS 2018-Network and Distributed System Security Symposium*, 2018, pp. 1–15.
- [52] Q. V. Liao and K. R. Varshney, "Human-Centered Explainable AI (XAI): From Algorithms to User Experiences," *arXiv:2110.10790 [cs]*, Jan. 2022, arXiv: 2110.10790.
- [53] A. G. Martinez, "Facebook's not listening through your phone. it doesn't have to," *Wired*. November, vol. 10, 2017.
- [54] E. Fast, B. Chen, J. Mendelsohn, J. Bassen, and M. S. Bernstein, "Iris: A conversational agent for complex tasks," in *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, ser. CHI '18. New York, NY, USA: Association for Computing Machinery, 2018, p. 1–12. [Online]. Available: <https://doi.org/10.1145/3173574.3174047>
- [55] B. Liu, M. S. Andersen, F. Schaub, H. Almuhiemedi, S. A. Zhang, N. Sadeh, Y. Agarwal, and A. Acquisti, "Follow my recommendations: A personalized privacy assistant for mobile app permissions," in *Twelfth Symposium on Usable Privacy and Security (SOUPS 2016)*. Denver, CO: USENIX Association, Jun. 2016, pp. 27–41. [Online]. Available: <https://www.usenix.org/conference/soups2016/technical-sessions/presentation/liu>
- [56] I. Markit, "The economic value of behavioural targeting in digital advertising," research report (accessed January 21, 2019), <https://datadrivenadvertising...>, Tech. Rep., 2017.
- [57] V. Marotta, V. Abhishek, and A. Acquisti, "Online tracking and publishers' revenues: An empirical analysis," in *Workshop on the Economics of Information Security*, 2019.

Appendix

1. ESM Questionnaire

- 1) Please use a few words to describe what the ad is about. (e.g., shopping related to clothes; car insurance; promotion of a new smartphone) (Free response)
- 2) Did you click on this ad after seeing it? (Yes/No/Not sure)
- 3) Were you interested in this ad? (Yes/No)
- 4) Would you like to know more about why you were specifically targeted by this ad? (Yes/No)
- 5) Please explain why you would or would not like to know more about the targeting of this ad towards you: (Free response)
- 6) Please select an image from each of the three groups that best reflects your feelings when you saw the ad



- 7) How did you feel about the ad? Please indicate the degree to which you agree on the statements: (1 = strongly disagree; 5= strongly agree)
- a) I found this ad to be unexpected.
 - b) I found this ad to be creepy.
- 8) Please provide a brief notation about your feeling when you saw this ad: (e.g., The ad felt creepy since the ad was exactly the same item I browsed on Amazon earlier today; I did not expect to see the ad at all.) (Free response)
- 9) To present this ad, what kind of information do you think were collected about you: (choose all that apply)
- a) Browsing Information (e.g. pages visited, search terms entered)
 - b) Demographic Information (e.g. age, sexual orientation, income bracket)
 - c) Location Information (e.g. state, ZIP code)
 - d) Personal Identification Information (e.g. email address, name)
 - e) Computer Information (e.g. computer OS, Web browser version)
 - f) Others (Free response)
 - g) None (I thought none of my personal information was collected to present this ad)
- 10) Where were you when you saw the ad? (Home/Work/Public/Other (Free response))
- 11) Select the option that most closely reflects what you were doing online when you saw this ad:
- a) Fact finding and looking for specific information (e.g. weather, location)
 - b) Information gathering and researching some broader topic (e.g. job hunting)
 - c) Pure browsing for self's routine/habit/passing time/entertainment
 - d) Transactions (e.g. online banking, filling a survey or application)
 - e) Communications (e.g. email, blog and post updates, messaging)
 - f) Maintaining and housekeeping the upkeep of web resources (e.g. maintaining web pages)
 - g) Others (Free response)

2. Exit Interview Protocol For Each Ad Asked

- 1) What were you doing on the internet when you saw the ad?
- 2) On which website did you see the ad?
- 3) When did you see this ad?
- 4) Where were you physically?
- 5) How did you feel about the ad when you saw it?
 - a) What made you have that feeling?
 - b) *(If the feeling is negative)*
 - i) Did such a feeling affect how you saw ads right after that ad?
 - ii) Did you or had you thought about doing something about it in response to the negative feeling of the ad?
- c) Is this the type of ad that you would typically see?
 - i) *(If yes)* What makes you feel that the ad is typical?

TABLE 2. CODEBOOK FOR THE QUALITATIVE ANALYSES.

Why do users want ad targeting explanations?		
Unexpected Use of Personal Data Physical location Demographic information Personal profiles and interests Sensitive identifying information Data provenance Data flow	Testing Preconceptions Preconceptions about how personal data is used for ad targeting – <i>browser history and activity</i> – <i>browser information</i> – <i>demographic</i> – <i>physical location</i> – <i>personal profiles and interests</i> Preconceptions about how personal data was collected and shared – <i>data breaches of, and sharing by, first-party websites</i> – <i>cross-device tracking</i> – <i>listening to physical-world conversations</i> – <i>tracking via spyware and malware</i>	Unexpected Targeting Mismatch with self-aware digital footprint Mismatch with interests and conceptions of self – <i>mismatched profiles</i> – <i>mismatched interests</i>
Abnormal Ad Characteristics Unclear ad Fake ad or clickbait Ad content that looks unusual Advertiser characteristics		Intriguing Ad Patterns Repetition Meta patterns Agency & Autonomy Desire to gain agency and autonomy
Why do users not want ad targeting explanations?		
Confident Understanding Targeting demographic Targeting physical location Targeting browsing activities Reasonable interest targeting Not targeting / targeting broadly Targeting temporal factors Hosting website characteristics Advertiser characteristics	Repetition and Familiarity Feel less invasive overtime Marketing strategy No additional thoughts on the ad Indifference Unclear ad Uninteresting ad content	Helplessness and Resignation Lack of accountability Personal data was already breached and collected

- ii) (If no) What makes you feel that the ad is different?
- a) Why are you not interested in the ad targeting of this ad?
- 6) When seeing the ad, did you have any idea why you saw this ad?
- 7) (If participants answered “Yes” to the question “Would you like to know more about why you were specifically targeted by this ad?”)
- a) Why are you interested in the ad targeting of this ad?
- b) What specific ad targeting information would you want to know about?
- 8) (If participants answered “No” to the question “Would you like to know more about why you were specifically targeted by this ad?”)

TABLE 3. LABELS FOR AD TOPICS (*) DENOTES AD TOPICS NOVEL FROM ZENG ET AL. [32]

Label	Definition
Apparel	Ads for clothes, shoes, and accessories
B2B Products	Ads for any product intended to be sold to businesses
Banking and Debt/Credit Card	Ads for financial services that banks or debt/credit card companies provide to consumers or financial services directly related to banks or debt/credit card, exclude loans or mortgages
Beauty Products	Ads for cosmetics, skincare, haircare and styling products
Cars	Ads for automobiles, vehicle, motorcycles, gas, and related services (e.g. car repair)
Cell and Internet Service	Ads for mobile phone and internet plans for consumers
Celebrity News	Ads for articles about celebrities; gossip
Consumer Tech & Tech Services	Ads promoting smartphones, laptops, smart devices; accessories for consumer electronics or services for these electronics
Contest	Ads for giveaways, lotteries, gambling etc.
COVID	Ads for masks, hand sanitizer, health measures, and vaccination for COVID
Dating	Ads for dating apps and services
Education	Ads for colleges, degree programs, training, etc.
Employment	Ads about job listings, services related to job searching, and micro-labor tasks
Entertainment	Ads for entertainment content, e.g., TV, books, movies, etc.
Finance and Investment Pitch	Ads promoting a specific investment product, opportunity, newsletter, and services related to finance advice or investing
Food and Drink	Ads regarding anything food related, e.g., recipes and restaurants
Games and Toys	Ads for video games, board games, mobile games, toys
Genealogy	Ads for genealogy services/social networks
Gifts	Ads for gifts, gift cards
Health and Supplements	Ads for supplements, wellness advice, fitness, and personal cares, excludes medical services
Household Products	Ads for furniture, home remodeling, any other home products
Humanitarian	Ads for charities and humanitarian efforts, public service announcements
Human Interest	Ads for articles that are generic, evergreen, baseline appealing to anyone
Insurance	Ads for any kind of insurance product — home, car, life, health, etc.
Journalism	Ads from journalistic organizations — programs, newsletters, etc.
Legal Services	Ads for law firms, lawyers, or lawyers seeking people in specific legal situations
Medical Services and Prescriptions	Ads for prescription drugs, medication, doctors and specific medical services
Mixture*	Ads with a collection of promoted content and links
Mortgages and Loans	Ads for mortgages, mortgage refinancing, reverse mortgages, or loans
Pets	Ads for pet products
Place to Shop*	Ads from a shopping platform or store that sell items of various category, and the ad does not promote a specific item or a specific type of items
Political Campaign	Ads from an official political campaign
Political Memorabilia	Ads for political souvenirs/memorabilia, like coins
Public Relations	Ads intended to provide information about a company to improve public perceptions
Real Estate	Ads for property rentals/sales
Recreational Drugs	Ads for alcohol, tobacco, marijuana, or other drugs
Religious	Ads for religious news, articles, or books
Social Media	Ads for social media services
Software and Application	Ads promoting consumer-facing online/offline software and application
Sports	Ads with anything sports-related - sports leagues, sports equipment, etc.
Senior Caring and Living	Ads for senior living and caring services
Travel	Ads for anything travel related - destinations, lodging, vehicle rentals, flights
Weapons	Ads for firearms or accessories like body armor
Wedding Services	Ads for any services or products specifically for weddings, like photographers
YouTube Merchandise* Catalog	Ads for a collection of merchandise provided by a YouTube Channel
Miscellaneous*	Specific ads that could not be categorized by the aforementioned topics (e.g., tattoos, boat motors, psychic tarot)

TABLE 4. GENERAL PARTICIPANT DEMOGRAPHICS

#	Age	Gender Identity	Education	Occupation	Interview
P1	33	Woman	Master's degree	Art, writing, or journalism (e.g., author, reporter)	V
P2	29	Man	Master's degree	Engineer in other fields (e.g., civil engineer, bio-engineer)	
P3	32	Woman	Master's degree	Administrative support (e.g., secretary, assistant)	
P4	29	Woman	Bachelor's degree	Medical (e.g., doctor, nurse, dentist)	V
P5	33	Woman	Professional degree	Medical (e.g., doctor, nurse, dentist)	V
P7	48	Woman	Bachelor's degree	Business, management, or financial (e.g., manager, accountant, banker)	
P8	34	Woman	Some college	Administrative support (e.g., secretary, assistant)	
P9	37	Woman	Bachelor's degree	Business, management, or financial (e.g., manager, accountant, banker)	V
P10	40	Man	Trade, technical, or vocational training	Healthcare/Holistic Coach	V
P11	24	Man	Bachelor's degree	Computer engineer or IT professional (e.g., systems administrator, programmer, IT consultant)	V
P12	42	Man	Master's degree	Transportation Planner	V
P13	57	Woman	Bachelor's degree	Education (e.g., teacher)	V
P14	19	Woman	High school	Student	
P15	28	Woman	High school	Legal (e.g., lawyer, law clerk)	V
P16	23	Man	Some college	Unemployed	V
P17	22	Woman	Some college	Student	V
P18	34	Woman	Master's degree	Education (e.g., teacher)	V
P19	31	Man	Bachelor's degree	Computer engineer or IT professional (e.g., systems administrator, programmer, IT consultant)	V
P20	20	Man	High school	Student	
P21	49	Man	Master's degree	Medical (e.g., doctor, nurse, dentist)	V
P22	24	Woman	Some college	Art, writing, or journalism (e.g., author, reporter)	
P23	19	Man	Some college	Student	V
P24	52	Woman	Bachelor's degree	Business, management, or financial (e.g., manager, accountant, banker)	V
P25	31	Woman	Bachelor's degree	Art, writing, or journalism (e.g., author, reporter)	V
P26	21	Man	Bachelor's degree	Student	V
P27	34	Man	Master's degree	Business, management, or financial (e.g., manager, accountant, banker)	V
P28	36	Man	Master's degree	Scientist (e.g., researcher, professor)	V
P29	31	Woman	Bachelor's degree	Medical (e.g., doctor, nurse, dentist)	V
P30	21	Woman	High school	Scientist (e.g., researcher, professor)	V
P31	28	Woman	Bachelor's degree	Homemaker	V
P32	37	Man	Associate's degree	Unemployed	V
P33	29	Man	Master's degree	Business, management, or financial (e.g., manager, accountant, banker)	
P34	29	Man	Bachelor's degree	Technician	V
P35	49	Man	Master's degree	Education (e.g., teacher)	V
P36	34	Man	Master's degree	Business, management, or financial (e.g., manager, accountant, banker)	V
P37	25	Woman	Master's degree	Education (e.g., teacher)	V
P38	44	Man	Master's degree	Education (e.g., teacher)	V
P39	40	Man	Bachelor's degree	Pastor	V
P40	40	Woman	Bachelor's degree	Education (e.g., teacher)	V
P41	31	Man	Some college	Computer engineer or IT professional (e.g., systems administrator, programmer, IT consultant)	V
P42	47	Woman	Some college	Service (e.g., retail clerks, server)	
P43	36	Woman	Bachelor's degree	Digital marketing consultant	
P44	33	Man	Some college	Apartment Management	V
P45	26	Woman	Master's degree	Administrative support (e.g., secretary, assistant)	V
P46	50	Woman	Master's degree	Business, management, or financial (e.g., manager, accountant, banker)	V
P47	28	Man	Some college	Service (e.g., retail clerks, server)	
P48	24	Man	Bachelor's degree	Engineer in other fields (e.g., civil engineer, bio-engineer)	
P49	49	Man	Some high school	Homemaker	
P50	33	Woman	Master's degree	Art, writing, or journalism (e.g., author, reporter)	
P52	33	Man	Some college	Administrative support (e.g., secretary, assistant)	
P53	33	Man	Associate's degree	Skilled labor (e.g., electrician, plumber, carpenter)	
P54	22	Man	Some college	Student full time and Delivery Service driver	V
P55	33	Man	Master's degree	Scientist (e.g., researcher, professor)	
P56	29	Man	Professional degree	Legal (e.g., lawyer, law clerk)	
P57	39	Woman	Bachelor's degree	Administrative support (e.g., secretary, assistant)	
P58	45	Man	Master's degree	Business, management, or financial (e.g., manager, accountant, banker)	
P59	41	Woman	Bachelor's degree	Education (e.g., teacher)	
P60	24	Woman	Some college	Art, writing, or journalism (e.g., author, reporter)	
P61	39	Woman	Master's degree	Computer engineer or IT professional (e.g., systems administrator, programmer, IT consultant)	
P62	25	Woman	Bachelor's degree	Engineer in other fields (e.g., civil engineer, bio-engineer)	