Unpacking Perceptions of Data-Driven Inferences Underlying Online Targeting and Personalization

Claire Dolin*, **Ben Weinshel***, **Shawn Shan**,
Chang Min Hahn, **Euirim Choi**, **Michelle L. Mazurek**[†], **Blase Ur**
University of Chicago and University of Maryland (†)
{cdolin, weinshel, shansixioing, changhahn, euirim, blase}@uchicago.edu, mmazurek@cs.umd.edu (†)

ABSTRACT
Much of what a user sees browsing the internet, from ads to search results, is targeted or personalized by algorithms that have made inferences about that user. Prior work has documented that users find such targeting simultaneously useful and creepy. We begin unpacking these conflicted feelings through two online studies. In the first study, 306 participants saw one of ten explanations for why they received an ad, reflecting prevalent methods of targeting based on demographics, interests, and other factors. The type of interest-based targeting described in the explanation affected participants' comfort with the targeting and perceptions of its usefulness. We conducted a follow-up study in which 237 participants saw ten interests companies might infer. Both the sensitivity of the interest category and participants' actual interest in that topic significantly impacted their attitudes toward inferencing. Our results inform the design of transparency tools.

ACM Classification Keywords
H.5.2 User Interfaces: Evaluation/methodology

Author Keywords
Transparency; inferencing; tracking; OBA; targeting; privacy

INTRODUCTION
As users browse the internet, their online activity is tracked by the website they are visiting, as well as by third-party advertising and analytics companies using techniques ranging from HTTP cookies to browser fingerprinting [28]. Companies use these logs of browsing behavior to infer that user's interests, preferences, and demographics [30]. Both first-party and third-party companies tailor a user's internet experience in part based on these inferences, impacting the search results [47], ads [30, 42], and social feeds [14, 37] a user sees.

While users may be aware in general terms that their web experience is personalized, they often do not understand the mechanics of this personalization [30, 40]. Users find targeting useful because it highlights relevant information, yet simultaneously find the underlying data collection scary and creepy [46]. Many users incorrectly assume that targeting relies on straightforward inferences or information explicitly provided to companies [47]. As a result, users are surprised when inferences buried in large amounts of data reveal otherwise hidden information, as in the case in which the retailer Target knew a teen was pregnant before her parents did [20].

Users' discomfort with personalization and its associated tracking stems from multiple factors. Users consider targeting based on demographics to be discriminatory [36, 42]. Automatically generated interest categories on Facebook [3] and Google [21] that encapsulated racism and anti-semitism recently highlighted this issue. The accuracy of inferences also influences perceptions [7]. Nonetheless, the morass of users' conflicted feelings about online targeting and personalization remains complicated by users' poor understanding of the types of possible inferences [47] and how they are made [40, 46].

In this paper, we begin unpacking users' mixed feelings about online targeting and personalization. To do so, we conducted complementary between-subjects experiments on Mechanical Turk. These studies investigate how the method by which ads are targeted, the precise inferences made about a user, and the accuracy of those inferences impact user comfort with online targeting and personalization, as well as how they influence users' perceptions of the fairness and utility of these practices.

In our first study, we focused on methods of targeting ads. We showed 306 participants an ad for a product they might be interested in, along with one of ten explanations for why they had seen that ad. Participants then answered a series of questions about their reactions to the targeting method. Many participants' reactions varied based on the targeting method, confirming that the method of ad targeting does matter to consumers. In particular, ad targeting was perceived as more useful when the ad shown was directly related to the interest used for targeting, compared to when the ad was unrelated to the interest used for targeting, yet still relevant to the participant's shopping preferences. Participants also perceived interest-based targeting as more fair when it was based on the aggregate interests of all visitors to a website, rather than the interests of an individual visitor. Furthermore, participants were more comfortable with ads being targeted based on the aggregate interests of all visitors to a website, rather than an individual user's interests.

* Co-authors Dolin and Weinshel contributed equally to this work.

CHI 2018, April 21–26, 2018, Montréal, QC, Canada.
ACM ISBN 978-1-4503-5620-6/18/04.
http://dx.doi.org/10.1145/3173574.3174067

We then conducted a complementary study to understand how the accuracy and sensitivity of the specific interest on which targeting was based impacted perceptions. Each participant saw ten topics selected from among the actual topics used in Google's AdWords product [17]. We then asked participants to evaluate their comfort with companies making such an inference about them, as well as using that inference to personalize their online experience. To better understand the sensitivity of the inference, we also queried participants' comfort with other people (e.g., a significant other, a work supervisor) learning such an inference had been made about them.

We found large differences in the perceived sensitivities of inference categories. Predictably, participants found inferences about topics like sex, dating, divorce, alcohol, and religion to be sensitive. Whereas many advertising companies currently distinguish only between sensitive and non-sensitive categories (as in Google's AdWords [18]), we found that topics were not merely sensitive or insensitive. Instead, we observed many intermediate gradations of sensitivity. Further, the accuracy of an inference had a strong correlation with participants' comfort, regardless of sensitivity.

Our results confirm that the details of inferencing and targeting matter. How an inference is made, whether it is accurate, precisely what interest is inferred, and who is likely to find out about it all play a role in users' resulting comfort (or lack thereof) with online targeting and personalization. We discuss how these results can help inform the next generation of transparency tools, laws, and regulations.

RELATED WORK

We discuss the technical mechanisms of OBA, transparency tools, and user perceptions of algorithmic personalization.

Technical Mechanisms of Online Tracking

It is well-known that advertisers can set HTTP cookies with a unique identifier to correlate browsing activity with that unique identifier [23]. In recent years, companies have more widely adopted stateful and stateless ("fingerprinting") technologies [11]. In a 2010 sample of nearly 500,000 browsers, 83.6% could be uniquely fingerprinted [30]. Furthermore, researchers found that 78% of sites contained trackers that attempted to transfer unsafe, personal data [50].

Roesner et al. estimated that several trackers can each capture more than 20% of a user's browsing behavior [38], while Englehardt et al. reported that some companies can reconstruct 62-73% of a user's browsing history [12]. FPDetective [1], XRay [24], UAframwork [44], MyAdChoices [35], and Sunlight [25] respond to this tracking by using correlations and cookie matching to measure what companies gather.

User-Facing Transparency Tools

Some tools try to provide transparency about online tracking. These tools include browser extensions (e.g., Ghostery [6], Privacy Badger [10], and Lightbeam [9, 33]) that detect, display, and block third-party trackers. These tools increase awareness of tracking, but also introduce confusion [40].

Other options to avoid third-party tracking include opt-out cookies, Do Not Track headers, and third-party cookie blocking. With the exception of DNT headers, these methods reduce the amount of behavioral targeting [4, 22]. However, Leon et al. found these tools have significant usability flaws, including minimally protective defaults and confusing interfaces [26]. Melicher et al. found users are concerned about tracking in the abstract, yet current tools do not mitigate their concerns [32].

Some companies provide "privacy dashboards," where users can see some of the information the company has inferred about them and customize their interest profiles [15, 16, 34]. Unfortunately, automated experiments have found that the information on the dashboards is both inaccurate and opaque, and that online targeting currently discriminates on the basis of gender [8]. Similarly, Lecuyer et al. found that Google's practices contradict its own statements by targeting based on sensitive and prohibited topics [25]. Wills and Tatar also found that Google did not disclose all interests in the Ad Settings page and targeted based on sensitive topics [48].

User Perceptions of Algorithmic Personalization

Prior work has addressed users' attitudes and understanding of online targeting and personalization. Researchers have found that users do not want targeted advertising when made aware of how advertisers collected the underlying data [45]. Many users find targeting and personalization creepy or invasive [29, 31, 46]. Targeted ads can act as social labels that cause consumers to adjust their own self-perceptions to match the implied labels, while awareness of targeted advertising can worsen attitudes toward the product advertised [39, 41].

Leon et al. found that users were comfortable sharing certain classes of information with advertisers, but uncomfortable sharing other classes. These perceptions varied based on the advertiser's privacy policies [27]. Yao et al. showed that people are more concerned with the types of personal information collected than who was collecting it [49], while Agarwal et al. found that users were more concerned with the sensitivity of the ads being shown than the associated tracking [2]. Plane et al. found users were more concerned if an ad was targeted based on demographic information, such as age, gender, or race, than based on interests [36]. Attitudes also depend on the accuracy of inferences and the demographic attributes used [7], as well as on feelings of control over data collection [5].

The process of data-driven inferencing is widely misunderstood. Warshaw et al. found that users could be divided into two clusters: those who believed inferences were based on stereotypes related to directly provided demographic information, and those who believed that inferences were made with behavioral data, albeit via straightforward and intuitive logic [47]. Algorithmic personalization in social feeds is similarly misunderstood. In one study, 62.5% of participants were unaware Facebook personalizes its News Feed; learning about personalization led to greater feelings of control [14]. Users have developed "folk theories" of algorithmic personalization, modifying their behaviors in an attempt to manipulate the algorithms to their benefit [13, 37]. Taken together, this prior work informed the factors we considered in our experiments: both interest-based and demographic-based personalization.

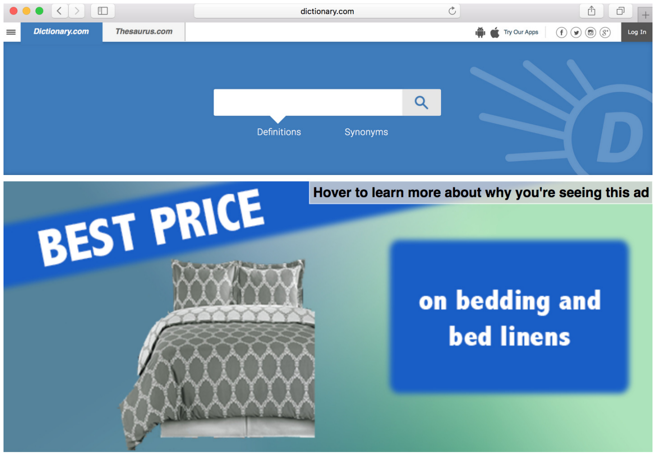

Figure 1: Sample ad (Study 1). Hovering covered the ad with an explanation of why that ad was chosen for that user.

Useful	I would find it useful to have ads targeted to me for this reason, as opposed to any other reasons.
Informative	This notification gives me enough information to understand why an advertiser would show me this ad.
Like to Know	I would like to know whenever an ad is targeted to me for this reason.
Comfortable	Overall, how comfortable or uncomfortable are you with companies advertising to you for this reason?
Fair to Target	I think it is fair for a company to target ads for this reason.
Fair to Collect	I think it is fair for a company to collect the information necessary to target ads for this reason.
Annoyed	I would be annoyed by this type of ad targeting.

Table 1: **Reaction Statements** for Study 1. Participants responded using seven-point Likert scales.

STUDY 1: IMPACT OF TARGETING MECHANISM

The goal of our first study was to learn how the method of targeting an advertisement (targeting based on demographics, inferences, or other factors) impacts privacy attitudes. We showed participants an example ad and an explanation, varying across conditions, for why they received that specific ad.

Methodology

For both Study 1 and Study 2, we recruited participants on Amazon's Mechanical Turk. Both studies used a between-subjects design, and both studies were approved by our institution's IRB. We include the full survey instruments in our online supplementary materials. The remainder of this section describes the procedure specific to Study 1, for which we had 306 participants. The survey took approximately 15 minutes, and we compensated participants $2.50.

Because the accuracy of an inference can affect the subject's feelings about the fairness of the inference [7], we controlled for accuracy by dynamically generating the advertisement and explanatory notification based on a participant's actual demographics and interests. To learn this information about each participant, we began Study 1 with (optional) questions about demographics and interests. Were a participant not to provide an age or gender, the survey would randomly select demographics and display "You chose not to provide your age or gender. For the purposes of the following questions, you may pretend that you are..." That said, all of our participants provided their age and gender. We wanted the example ads and products to be neutral and not overly associated with any one demographic. As such, we selected twelve possible shopping categories from Amazon of products costing approximately $50. Participants chose multiple categories they were interested in, and the example ad used a product from one of these selected categories. Figure 1 shows an example ad from the category "bedding and bed linens." Participants also selected topics that they were likely to read about online from a possible eighteen choices, taken from Google's list of topics used for personalized ads [19]. These topics served as an unrelated interest in some of the conditions, as described below. We did not show participants possible interest categories closely related to the chosen shopping category.

We used a between-subjects design to compare reactions to different kinds of ad explanations. All explanations begin with "You are seeing this ad because...," followed by one of ten possible explanations (Figure 2). Two of these possibilities are control conditions, stating either that "the advertiser decided to purchase an ad on this site" (the *Control* condition) or "the advertiser's computer algorithms have determined this ad would be effective" (the *Algorithm* condition). The eight other conditions, termed *inference explanations*, have a full-factorial design along two dimensions we varied (Figure 2). The inference explanations state that "the advertiser has inferred" something leading to the ad. These conditions test several variables: the target of the inference (the user specifically, or visitors to this website), the mechanism of the inference (age/gender demographics or browsing interest), and whether the browsing interest information is intuitively related or not.

Using these ten conditions, we investigated five research questions about how attitudes varied with the method of targeting:

- **RQ1**: Does it matter whether targeting is determined by generic "computer algorithms" or is specific to the participant? To answer this question, we compared Algorithm vs. You-Demographic and Algorithm vs. You-Interest.
- **RQ2**: Does it matter whether targeting is based on all visitors to a site, rather than specific to one user? We compared Visitors-Demographic vs. You-Demographic and Visitors-Interest vs. You-Interest.
- **RQ3**: Does it matter whether targeting is based on a correct or incorrect demographic inference? We compared Visitors-Demographic vs. Visitors-WrongDemographic and You-Demographic vs. You-WrongDemographic.
- **RQ4**: Does it matter whether targeting is based on an inference about an interest directly related to the product shown or an interest with no apparent relation? We compared Visitors-Interest vs. Visitors-UnrelatedInterest and You-Interest vs. You-UnrelatedInterest.
- **RQ5**: Does it matter whether or not targeting is performed by an algorithm? We compared Control vs. Algorithm.

Participants then responded to reaction statements (Table 1) about their comfort level with, and perceptions of, the targeting method. Responses were on seven-point agreement Likert scales, except for the Comfortable statement, which used a

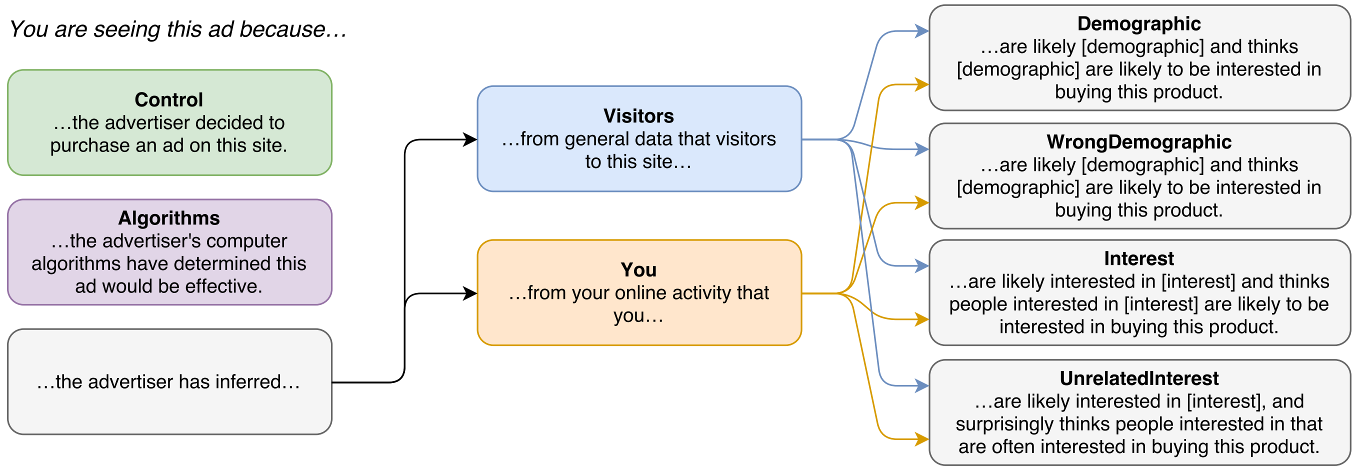

Figure 2: Study 1 conditions. The notifications shown were *Control*, *Algorithms*, and the eight combinations possible from selecting one setting from the second column and one setting from the third column (e.g., *You-Interest*).

seven-point comfort Likert scale. We designed the reaction statements to capture the same sentiments as the scale validated by Samat et al. [39]. After each question, participants explained their answer. Finally, we asked general questions about participants' knowledge about targeted advertising.

Analysis Methods and Metrics

All statistical tests use $\alpha = .05$. For each of the seven reaction statements (Table 1), we created a proportional-odds logistic regression with the reaction statement (ordinal) as the dependent variable. The independent variables were the explanation the participant saw (nominal), as well the participant's age category (ordinal), gender (nominal), technical expertise (nominal), and race (nominal, binned as "white" or "non-white"). For each of these seven models, we first performed an omnibus test comparing the model with and without the independent variable specifying the explanation the participant saw. To investigate the five specific research questions described above, we ran planned contrasts (shown in the second column of Table 2) using the glht function in the multcomp R package. If a pairwise contrast is not noted as significant in the results section, it was not found to be statistically significant. We ran contrasts only where the omnibus test was significant. Because these contrasts are not orthogonal, we used Holm correction to control the family-wise error rate.

Two researchers independently coded free-text responses using a joint codebook. Their average agreement across the seven free-text questions (Cohen's κ) was 0.697. The question with the lowest agreement (understand inferences) had $\kappa = 0.667$.

Results

We had 306 participants in Study 1. Reflective of the Mechanical Turk population, our sample was younger and more highly educated than the general population; 60% of participants were between 18 and 34 years old. Among participants, 42% identified as female, 58% identified as male, and 1% identified as non-binary. Because the impact and effects of demographic-based inferencing varies across groups [36, 42], we also asked about participants' race. In total, 80% of participants identified as white, 9% identified as black or African-American, and 7% identified as Asian. Across groups, 10% of participants identified as Hispanic or Latinx. Overall, 18% of participants held a degree or job in computer science or a related field. We used an attention check question "I have used a web browser to access the internet" in the middle of the survey, excluding responses that did not agree.

The particular explanation for why the ad was shown impacted participants' comfort and perceptions of the fairness of the targeting. It also impacted their perceptions of the usefulness of receiving ads targeted in that way and of the informativeness of the notice. More formally, for four of the seven reaction statements (see Table 1), participants' responses varied significantly across conditions. Participants' responses to these four reaction questions are summarized in Figure 3, while the results of the planned comparisons are shown in Table 1. Appendix C of our online supplementary materials presents the full regression tables.

Participants found notifications invoking the user's own activity more informative than those that explained targeting as a decision by an "algorithm" (RQ1). They felt that interest-based targeting of all site visitors in aggregate was more fair than interest-based targeting of specific users, and they were also more comfortable with it (RQ2). We did not observe significant differences regarding the accuracy of demographic inferences (RQ3). Unsurprisingly, participants considered targeting based on related interests more useful, and the explanations more informative, than targeting when the interests had no apparent relation to the product shown (RQ4). Supporting RQ5, participants considered ads they were told were targeted by algorithms more useful, but less fair, than ads they were told were selected simply based on an advertiser's purchase.

General Perceptions of Targeting

To contextualize participants' answers, we begin by reporting their general perceptions of online targeting. Most participants felt they understood the mechanisms of OBA, including 83% who at least somewhat agreed that they understood "how ad-

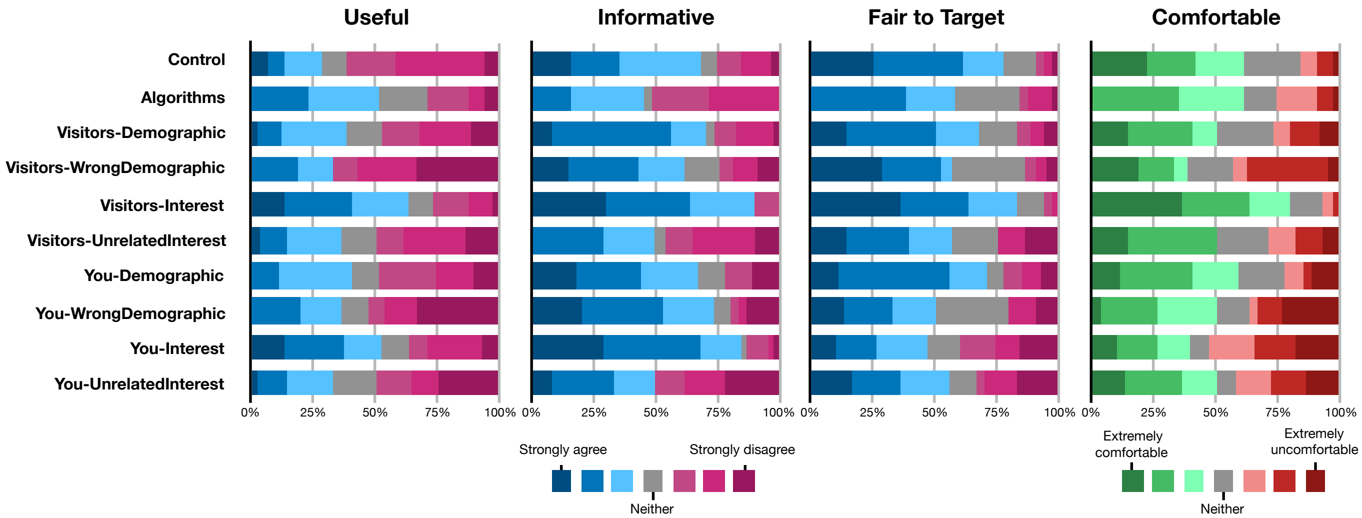

Figure 3: Agreement with Study 1 reaction statements that differed significantly by condition.

RQ	Comparison	More Useful	More Informative	More Fair to Target	More Comfortable
RQ1	Algorithm vs. You-Interest	–	You-Interest ($<$ .001)	–	–
	Algorithm vs. You-Demographic	–	–	–	–
RQ2	Visitors-Interest vs. You-Interest	–	–	Visitors-Interest (.003)	Visitors-Interest ($<$ .001)
	Visitors-Demographic vs. You-Demographic	–	–	–	–
RQ3	You-Demographic vs. You-WrongDemographic	–	–	–	–
	Visitors-Demographic vs. Visitors-WrongDemographic	–	–	–	–
RQ4	You-Interest vs. You-UnrelatedInterest	You-Interest (.045)	You-Interest ($<$ .001)	–	–
	Visitors-Interest vs. Visitors-UnrelatedInterest	Visitors-Interest (.037)	Visitors-Interest (.002)	–	–
RQ5	Control vs. Algorithm	Algorithm (.049)	–	Control (.022)	–

Table 2: Results of planned contrasts investigating the five research questions in Study 1.

vertising companies target advertisements" and 81% who at least somewhat agreed they understood "how advertisers make inferences." These questions were shown near the end of the survey, so it is possible that these responses were primed.

In total, 35% of participants mentioned either their browsing history or their most recent search query as a source of inferences. 20% mentioned activity tracked on specific websites (e.g., P-199: "they use the sites that I have been to previously to form a profile."). 19% believed that inferences were based on demographic information (e.g., age, gender, race, location). Asked how advertisers choose ads, 49% of participants reported that they were selected based on browsing history and recent search queries. 22% mentioned cookies. Demographics were also mentioned by 22% of participants. Of the 11% of participants who said they did not understand how advertisers choose ads, many found the process creepy. For example, P-80 said "it feels like voyeuristic spying."

Participants expressed frustration with the lack of transparency in advertising: 64% disagreed that "advertising companies are transparent about how they target advertisements to me." Of those 198 participants, 34 said that the advertising companies were intentionally hiding information or misleading users.

Useful

In this section and the three that follow, we detail the four reaction statements for which responses differed significantly

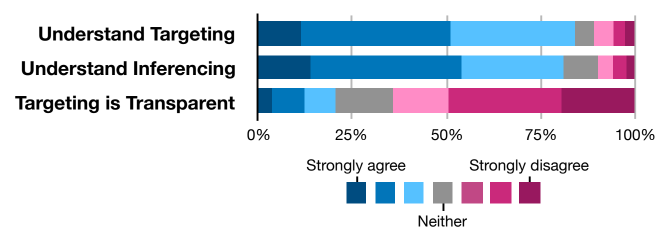

Figure 4: Participants' general perceptions of targeting.

by condition. First, participants' perception of the usefulness of targeting using a specific method (the Useful reaction statement), varied significantly by condition ($p < .001$). This echoes the tension between the utility of targeting and associated privacy risks identified in prior work [46].

In particular, we found that targeting based on directly related interests was perceived to be more useful than targeting based on unrelated interests (RQ4). Participants considered ads targeted based on Visitors-Interest to be more useful than those targeted based on Visitors-UnrelatedInterest ($p = 0.037$), with 63% agreeing that this kind of targeting is useful when the inference was directly related to the ad, as opposed to 36% when the inference was unrelated. Likewise, 52% of participants in the You-Interest condition found such targeting useful, while only 34% in You-UnrelatedInterest did ($p = 0.045$). Many free responses in the UnrelatedInterest conditions questioned the seemingly useless connection between the interest and the product shown. For example, P-253 wrote, "I don't see how it makes any sense. I don't see the connection between business news and bed spreads."

We also found significant differences regarding RQ5; 52% of participants told that ad targeting was based on decisions by an "algorithm" (Algorithm) found this mechanism useful, compared to 29% of those who saw the Control notice ($p = 0.049$). Participants said they preferred to see relevant ads.

Informative
Participants perceived the thoroughness of the explanations differently. Formally, agreement with the Informative reaction statement varied by condition ($p < .001$). Compared to participants who were told that the targeting was based on "computer algorithms", participants told that the targeting was to a specific user found the notification more informative (RQ1). Participants found the You-Interest notification more informative than the Algorithm notification ($p < 0.001$), with 84% of You-Interest participants agreeing the statement was informative, compared to 45% of Algorithm participants.

Participants also found targeting based on a related interest more informative than targeting on unrelated interests (RQ4) for both the You and Visitors conditions ($p < 0.001$ and $p = 0.002$, respectively). 90% of You-Interest participants and 85% of Visitors-Interest participants agreed it was informative, while only 50% of participants in the UnrelatedInterest conditions considered it informative. Participants thought unrelated connections warranted explanation. For example, P-70 wrote, "How y'all making the leap from politics to office chairs?"

Overall, 29 participants did not feel enough information was provided about targeting. P-281 said the notice "does not go into detail about how the specific inferences were made which would have provided more transparency." 13 participants felt unclear about the details and mechanisms of the inferencing. P-72 said, "I would like to know specifically which of my browsing habits (i.e. specific websites visited, search terms used) led to this advertisement being selected."

Fair to Target
The method of targeting also impacted participants' responses to one, but not both, questions we asked about fairness. Agreement with the Fair to Target reaction statement, representing responses to the statement "I think it is fair for a company to target ads for this reason," varied across conditions ($p = .004$).

Interest-based targeting was considered more fair when it was about site visitors as opposed to a specific user ($p = 0.003$). 83% of Visitors-Interest participants considered targeting fair, compared to 47% who saw the You-Interest notice (RQ2). When participants felt the method of targeting was fair, it was often because they saw enough of a possible benefit or felt the company had the right to do as it wanted. Two common sentiments were that companies should use all available data to get the most out of advertising (16.3%) and that targeting is not invasive (21.8%). These feelings are summed up by P-156, who said, "In a group sense, it's not invasive. It's more like commercials on TV. They are targeted to a demographic, but not an absolutely personal one," and P-300, who said, "Companies are free to target their advertising as they wish. I see nothing wrong with this approach." More participants who saw the You-Interest notification felt that the method violated their privacy, such as P-286: "Although it does provide some benefit...I would prefer to actively indicate my interests in product types rather than being tracked and it inferred."

We also found a significant effect for RQ5. The Control condition was considered more fair than the Algorithm condition, 77% to 42% ($p = 0.022$). Participants who considered the control condition fair felt companies were free to buy ad space where they wanted or that there was no bias in the decision. Although many participants who saw the Algorithm notification also felt that it was fair, some expressed privacy concerns, including P-273, who said, "My private information should not be used without explicit permission from me."

Comfortable
The final reaction that differed across conditions was comfort with the method of targeting, measured via the Comfortable reaction ($p = .003$). We found that 80% of participants who saw the Visitors-Interest condition said they were comfortable with targeting, yet only 39% of participants who saw the You-Interest notice agreed ($p < 0.001$). Participants often felt that targeting to aggregated site visitors did not involve personal information. In contrast, participants who did not feel comfortable often made comments about their privacy being invaded and companies monitoring them. 43% of participants who felt uncomfortable with their given method expressed privacy concerns or thought the advertiser knew too much. For instance, P-258 (Visitors-Interest) said, "I am extremely comfortable with being shown an ad for this reason because they are not taking any private information about me, it's just a guess based on general information about other people, that if they like it, I might like it too" while P-91 (You-Interest) said "It makes me feel that my every move on the internet is being watched."

Additional Findings
We did not observe a significant difference in participants' reactions to the fairness of *collecting* the information described in the explanation ($p = .167$). Similarly, we did not observe significant differences across conditions for either the Annoyed outcome ($p = .207$) or the Like to Know outcome ($p = .628$). Nonetheless, 56% of participants would want to know when an ad was targeted to them. As P-190 (Visitors-Demographic condition) said, "I would always like to know specifics of why an ad is being targeted to me, so that I can help control my privacy, for instance to know if my personal data is being leaked."

Reactions to demographic-based inferences were very mixed. Some participants thought that age and gender could be a good predictor of purchases, while others felt the method made too many assumptions or relied on stereotypes. For example, responding to the Useful reaction statement, P-269 said, "I share a lot of interests with people in my age group of the same sex," while P-223 said, "My age and gender is too large of a group to target." Furthermore, P-79 noted, "It

Interested	I am interested in *topic*.
Visit Related Pages	How often do you visit web pages related to *topic*?
Comfort: Inferencing	I would be **comfortable with a company making an inference** about my level of interest in *topic*.
Comfort: Personalizing	I would be **comfortable with a company personalizing my web experience** based on an inference about my level of interest in *topic*.
Useful: Personalizing	I would find it **useful to have my web experience personalized** based on an inference about my level of interest in *topic*, as opposed to my level of interest in other topics.

Table 3: **Reaction Statements** for Study 2. For *Visit Related Pages*, participants chose among five time frequencies. For all other statements, participants rated their agreement on seven-point Likert scales.

feels a bit sexist." We did not observe any results supporting RQ3, that the accuracy of a demographic inference influences perceptions of targeting.

STUDY 2: INFERENCE-TOPIC SENSITIVITY

Having found in Study 1 that the mechanism by which ads are targeted significantly affects privacy attitudes, especially with respect to interest-based targeting, we conducted a follow-up study to understand whether the particular interest topic that an advertiser infers and the accuracy of that inference also affect privacy attitudes. In particular, we asked the following two research questions:

- **RQ6**: Does the particular interest (termed *topic*) that a company infers about a user affect the participant's comfort with that topic as the basis for personalization, as well as perceptions of the utility of such personalization?
- **RQ7**: Does the accuracy of the inference affect the participant's comfort and perceptions of the utility?

Methodology

We again used Mechanical Turk to recruit participants, this time for a "research study about online personalization and targeting." The survey took approximately 30 minutes, and we compensated participants $5.00. The full study instrument is in Appendix B of the online supplementary materials.

In the first part of Study 2, we gave participants a brief explanation that the web content they see can be personalized based on inferences made about their browsing activity. For context, we asked participants general questions about whether they feel like they understand how companies make inferences and personalize their web experience.

In the second part of Study 2, we showed participants a random selection of ten interest topics a company could have inferred about them. To maximize ecological validity, we chose topics verbatim from the list Google uses in AdWords [17]. This list is structured as a tree, containing a hierarchy of topics. When displaying a category name, we displayed the name of each node, and then either the parent at the top level of the tree (for second-level topics), or the parent at the second level (for all other topics).

The full AdWords list contains more than 2,000 topics. To understand how perceptions varied across individuals, we wanted to have at least 10 participants see each topic, so doing so for the full list would have required more than 2,000 participants. We therefore removed from the list the entire top-level category of geographic locations, and then manually removed redundant topics to narrow the list to 160 topics that we used for Study 2. We present the full list of 160 topics in Appendix D of our online supplementary materials.

For each topic, participants were asked about their comfort level and the perceived usefulness of inferencing and web personalization with respect to that specific topic. The text of these reaction statements, and the short names used throughout this paper, are included in Table 3. For *Visit Related Pages*, participants chose among five time frequencies. For all other statements, participants rated their agreement on seven-point Likert scales.

To further examine participants' concerns related to inferences about the ten selected topics, we asked participants to rate their agreement with four questions of the form "I would be comfortable with a *relationship* knowing about my level of interest in *topic*." We asked this question for each of the following four relationships, varying by level of intimacy: significant other; close friend; work supervisor; and acquaintance. The first three relationships were filled into the above sentence verbatim. For clarity and precision, we used "an employee of a store that I see from time to time but have never had a conversation with" in place of acquaintance.

After repeating these sections for ten different topics, participants were asked to provide demographic information, including age, gender, race, ethnicity, and technical expertise.

Analysis Methods and Metrics

To gauge how both topic sensitivity and the user's actual interest in a topic (*accuracy*) affect attitudes, we constructed a series of mixed-effects ordinal regressions. Our model included a random effect for each participant, which models per-participant variation (latent factors) based on the 10 different data points for each participant.

Our models controlled for attribute sensitivity through two separate independent variables: the participant's agreement with Comfort: Personalizing for that topic, and the mean agreement with Comfort: Personalizing across all participants who saw that topic (termed *Mean Comfort: Personalizing*). The latter term serves as a proxy for the per-topic sensitivity by averaging across participants. Note that one model, in which Comfort: Personalizing was the dependent variable, did not also include that same factor as an independent variable. The models also controlled for participants' actual interest (accuracy) through two other independent variables: the participant's stated interest in a topic (the participant's agreement with Interested), as well as the frequency with which they visited pages related to that topic (the response to Visit Related Pages). We included both factors because interests potentially implied by a users' browsing may not align with their actual interests.

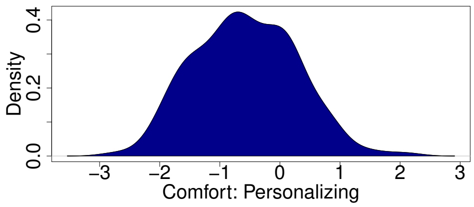

Figure 5: Kernel Density Plot (KDP) of the mean participant response for each topic to "I would be comfortable with a company personalizing my web experience based on an inference about my level of interest in *topic*." On the x-axis, -3 represents "strongly disagree" and 3 represents "strongly agree." Note that values outside the (-3,3) range are because a KDP *estimates* probability density from observations.

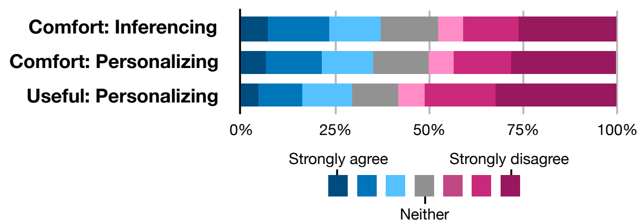

Figure 6: Participants' perceptions of comfort and usefulness regarding personalization based on the topics they saw.

Results

A total of 237 Mechanical Turk workers participated in Study 2. We again used an attention check question and eliminated responses that indicated a lack of attention. Similar to Study 1, our participants were younger than the general population; 57% were between 18 and 34 years old. Among participants, 45% identified as female, 54% identified as male, and 1% identified as non-binary. Again similar to Study 1, 81% of participants identified as white, 8% identified as black or African-American, and 5% identified as Asian. Across groups, 8% of participants identified as Hispanic or Latinx. While 15% of participants held a degree or job in computer science or a related field, 83% did not.

Our key findings pertain to the ten randomly selected topics each participant saw. We first present participants' general responses across all topics before drilling down to understand how topics differ from each other, as well as the extent to which the topic itself and the accuracy of the inference influenced participants' responses. Finally, we discuss participants' general perceptions of online personalization.

Overall Comfort and Accuracy

Two questions we asked focused on participant comfort with a topic. As shown in Figure 6, participants agreed with *Comfort: Inferencing* for 37% of the topics they saw and with *Comfort: Personalizing* for 35% of the topics they saw. Because responses to these statements were highly correlated (Spearman's ρ=0.918, $p < .001$) and our focus is on the use of inferences for personalization, we only use *Comfort: Personalizing* in our models and in the remainder of the paper. For 30% of the topics they saw, participants agreed with *Useful: Personalizing*. Together, these results echo prior work noting that users find topic-based personalization both useful and uncomfortable [43,46].

As we drilled down further, we found that the precise topic mattered a great deal in participants' responses. When considering user privacy, both the ad-interest dashboards provided by companies (e.g., ad-interest dashboards from Google [16], Bluekai [34], Facebook [15], and others) and depictions in the academic literature tend to treat the inferences made about users as bimodal: a small number of topics (e.g., sexual health, alcohol, and gambling) are sensitive, and other topics are not.

We found that this binary model of sensitive and non-sensitive topics did not effectively capture participants' perceptions; instead, comfort with use of inferences for personalization is more accurately viewed as a spectrum (RQ6). Figure 5 depicts a kernel density plot for the mean agreement for Comfort: Personalizing averaged across all participants who were asked about a given topic. This plot shows that the average comfort with personalization for many topics was close to neutral, tending slightly toward discomfort on average. However, participants were nearly universally uncomfortable with certain topics being used for personalization, while they were nearly universally comfortable with others. In Appendix D of our online supplementary materials, we provide the mean participant response for all 160 topics we investigated.

Unsurprisingly, participants were less comfortable with inferencing and personalization on health topics, whereas they were more comfortable with topics like travel. They were also generally uncomfortable with topics both directly and indirectly related to religion. For example, "Christianity," "Christian & Gospel Music," and "Islamic Holidays" were all topics for which participants felt uncomfortable with web personalization. Participants were relatively neutral in their comfort with the topic "Christmas," however, perhaps because it is perceived as more secular.

We observed considerable variation within top-level categories, some of which was surprising. For instance, we found stark differences within the "Social Issues & Advocacy" category. Nearly all participants who saw the "Environmental Issues" topic were comfortable with personalization based on it, whereas participants were uncomfortable on average with most other topics in that category (e.g., "Same-Sex Marriage," yet also "Charity & Philanthropy"). The topics respondents were most comfortable with often included topics for which personalized ads and deals would be helpful, like "Beaches & Islands" or "Video Games." However, that topics closely related to these two ("Air Travel" and "Shooter Games") were perceived as much more sensitive may be related to perceptions of price discrimination in the former case or secondary inferences about violence in the latter case.

Participants agreed that they were interested in 29% of the topics they saw. That said, their reported browsing habits did not always match these expressed interests. Across all topics, participants visited pages related to a topic they saw at least monthly for only 18.6% of topics. Figure 7 illustrates these results about participant interest in the topics they saw.

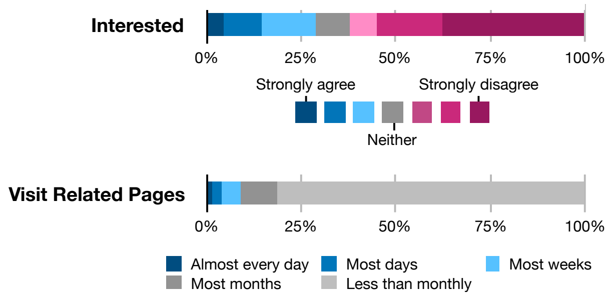

Figure 7: Participants' actual interest in the topics and their frequency visiting page related to those topics.

Table 4: Mixed-model ordinal regression to determine correlations with higher agreement that "I would be comfortable with a company personalizing my web experience based on an inference about my level of interest in *topic*" (*Comfort: Personalizing*). Significant p-values are bolded.

Factor	β	SE	t	p
Mean Comfort: Personalizing	1.03	0.06	17.25	**<.001**
Interest	2.38	0.21	11.49	**<.001**
Visit Related Pages	0.38	0.33	1.17	0.244

Table 5: Mixed-model ordinal regression to determine correlations with higher agreement that "I would find it useful to have my web experience personalized based on an inference about my level of interest in *topic*, as opposed to my level of interest in other topics" (*Useful: Personalizing*).

Factor	β	SE	t	p
Comfort: Personalizing	8.23	0.27	30.48	**<.001**
Mean Comfort: Personalizing	0.09	0.06	1.33	0.184
Interest	1.17	0.21	5.46	**<.001**
Visit Related Pages	1.12	0.37	3.07	**0.002**

Comfort and Usefulness

We found that both the sensitivity of the topic and the accuracy of the inference had a significant impact on participants' perceptions (RQ 7). As shown in Table 4, the comfort participants as a whole had with personalization on a topic (Mean Comfort: Personalization) unsurprisingly was significantly correlated with each participant's comfort with such personalization ($p < .001$). As Mean Comfort: Personalization is a proxy for the (non-)sensitivity of a topic, the overall sensitivity of the topic therefore impacts participants' attitudes.

In addition, participants' actual interest in the topic significantly impacted their comfort with personalization based on that topic ($p < .001$). The less accurate the inference relative to the participant's actual interests, the less comfortable they were with personalization, though we did not observe the frequency with which participants visited pages related to that topic (Visit Related Pages) to be significantly associated with Comfort: Personalizing.

We observed a similar trend for Useful: Personalizing. As shown in Table 5, participants were more likely to agree with Useful: Personalizing if they were more comfortable with personalizing based on that topic ($p < .001$), if they were more interested in that topic ($p < .001$), and if they more frequently visited pages related to that topic ($p = .002$).

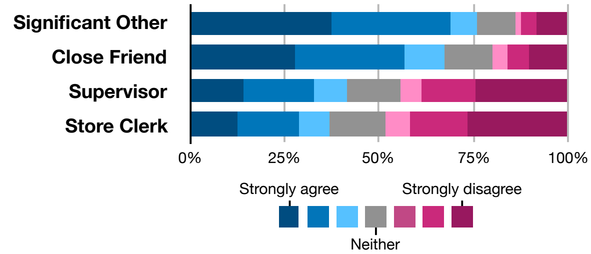

Figure 8: Reponses to questions about other people knowing about participants' interest in the topics shown.

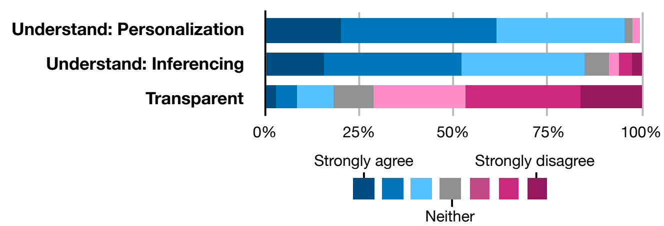

Figure 9: Reponses to Study 2 general perception questions.

Comfort Across Relationships

Our relationship questions asked about participants' comfort with people of four different relationships knowing that a company had made an inference about their interest in each topic they saw. We used these questions both as an additional proxy for topic sensitivity and to calibrate topic sensitivity with offline norms and expectations. Figure 8 depicts participants' responses to the relationship questions, across topics. Participants' responses were highly correlated for significant others and close friends (Spearman's ρ=0.779, $p < .001$), as well as for work supervisors and acquaintances (ρ=0.868, $p < .001$). As might be expected, participants were generally more comfortable with people from the two closer relationships learning about their interest a given topic. We again built regression models of participants' responses; the full models are in Appendix E of the online supplementary materials.

The sensitivity of the topic again correlated with participants' attitudes. Higher agreement with both the participant's own Comfort: Personalizing and the Mean Comfort: Personalizing (averaged across participants) was correlated with higher comfort with each of the four relationships knowing about the topic inference (all eight $p < .001$).

The accuracy of the inference also impacted participants' attitudes, but only for significant others and close friends. Higher levels of interest in a topic were correlated with higher comfort for significant others and close friends (both $p < .001$), and more frequent visits to related pages were correlated with higher comfort for close friends ($p < .001$). For work supervisors and acquaintances, however, neither Interest nor Visit Related Pages was a significant factor. We hypothesize that dealing with sensitive topics with supervisors and acquaintances is typically awkward, regardless of inference accuracy.

General Perceptions

We asked general questions about online personalization analogous to those from Study 1. Whereas we asked these questions at the end of Study 1 to avoid priming participants, we asked these questions at the beginning of Study 2 because possible priming effects were less likely to impact responses about very specific topics. Nonetheless, participants in this study responded to these questions similarly to Study 1 participants. As shown in Figure 9, 95% of participants agreed with the statement, "I feel that I understand how companies personalize my web experience" (*Understand: Personalization*) and 85% agreed "I feel that I understand how companies make inferences about my interests in order to personalize my web experience" (*Understand: Inferencing*). In contrast, only 18% of participants agreed that "I feel that companies are transparent about how they personalize my web experience." Note that these values represent only participants' self-perceptions, not the match between their knowledge and actual practices.

DISCUSSION AND FUTURE WORK

We believe our results will be valuable in designing the next generation of transparency tools and setting best practices for inferencing. Participants in Study 1 were more comfortable with targeting based on all site visitors' aggregate interests, rather than interests inferred about them in particular. Assuming user-specific targeting continues, these results suggest the need for clearer notice about such practices. Explanations about targeting on users' own demographics or interests were considered more informative than the explanation which only mentioned an algorithm. That is, participants found the general explanation of an algorithm doing the personalization opaque. Several participants asked to know what information was being used in the algorithm and where it had been collected. We thus suggest that future transparency tools more clearly identify the method and parameters of targeting, not just that targeting is occurring.

The power of big data rests in its ability to unearth hidden correlations buried in large amounts of data. This can lead to an advertisement being targeted based on interests that, to humans, seem unrelated to the ad, yet accurately capture the algorithm's training data and its inherent biases. In contrast, our participants considered advertising through straightforward inferences to be more useful than advertising through inferences about unrelated interests. Participants also found explanations about straightforward inferences more informative. Many participants felt that the logical jump from an interest in one topic to an interest in an unrelated product necessitated further explanation or justification. Future research could investigate whether greater transparency about the steps from an inference to personalization improves perceptions of the usefulness of targeting or informativeness of privacy notices.

We also compared reactions to an ad explicitly placed on a site and an ad targeted via an algorithm. Our participants found the first method more fair, but the second more useful, in large part because it seemed more likely to show relevant ads. This tension between fairness and usefulness echoes prior work [46]. In response, future transparency tools, both self-regulatory (e.g., AdChoices) and community-developed (e.g., browser plugins), could provide users with step-by-step explanations of inferences made about them and how those inferences are used for targeting. The particular interest categories on which an ad was targeted (rather than just that targeting was based on prior browsing) should be revealed. These recommendations contrast with widely used vague explanations.

In Study 2, inference topics were not equal in participants' eyes. The 160 Google AdWords interest topics that we presented to participants in Study 2 led to a gradient, not a bimodal distribution, of comfort, in contrast to companies' policies that only grant special consideration to a small list of highly sensitive topics [17]. Future work should critically re-examine the bifurcation of topics into sensitive and non-sensitive categories, perhaps creating a targeting-privacy calculus that takes the inference's sensitivity into account both when making an inference and targeting based on that topic.

While some companies are open about what topics they use for personalization [18], our detailed results on topic sensitivity can also inform industry-wide best practices for making inferences about topics. The range of comfort ratings across inferences suggests that users may benefit from more fine-tuned controls over which inferences about them are used to target ads. Additionally, ad buyers should be informed of the sensitivity of different categories so they can make more informed purchasing decisions to avoid alienating customers [39, 41].

Echoing prior work [7], we found that the accuracy of an inferred interest also plays a major role in user comfort. Regardless of the sensitivity of the topic, participants were more comfortable with accurate inferences being used for personalizing their online experience. This finding harkens back to the idea of privacy distortion, which specifies that inaccurate information about an individual is as much of a privacy violation as accurate information. While some companies have created privacy dashboards, our findings about the importance of inferencing accuracy on user attitudes emphasizes the need for improved privacy dashboards and greater access.

Participants also reported different levels of comfort with their interests being known by different social relations. Since ads may reflect a user's private browsing history or sensitive interests [2], different ad settings could be available at work, on shared computers, and at certain times of day.

Limitations

Our studies have a number of limitations. We report on a convenience sample and do not expect that the absolute percentages of participants' perceptions will generalize. As a result, we focus on the differences observed across conditions, rather than absolute numbers. We chose to run a controlled experiment with hypothetical ads and interests shown. This decision maximizes the power of the experiment to explain changes in attitudes, yet limits ecological validity.

ACKNOWLEDGMENTS

We gratefully acknowledge support from Mozilla and the Data Transparency Lab, as well as from a UMIACS contract under the partnership between the University of Maryland and DoD.

REFERENCES

1. Gunes Acar, Marc Juarez, Nick Nikiforakis, Claudia Diaz, Seda Gürses, Frank Piessens, and Bart Preneel. 2013. FPDetective: Dusting the web for fingerprinters. In *Proc. CCS*.

2. Lalit Agarwal, Nisheeth Shrivastava, Sharad Jaiswal, and Saurabh Panjwani. 2013. Do not embarrass: Re-examining user concerns for online tracking and advertising. In *Proc. SOUPS*.

3. Julia Angwin, Madeleine Varner, and Ariana Tobin. September 14, 2017. Facebook Enabled Advertisers to Reach 'Jew Haters'. *ProPublica*. https://www.propublica.org/article/facebook-enabled-advertisers-to-reach-jew-haters.

4. Rebecca Balebako, Pedro Leon, Richard Shay, Blase Ur, Yang Wang, and L. Cranor. 2012. Measuring the effectiveness of privacy tools for limiting behavioral advertising. In *Proc. W2SP*.

5. Farah Chanchary, Yomna Abdelaziz, and Sonia Chiasson. April 22, 2017. Privacy Concerns Amidst OBA and the Need for Alternative Models. To appear in *IEEE Internet Computing*. http://chorus.scs.carleton.ca/wp/wp-content/papercite-data/pdf/chanchary2017sharing-ic.pdf

6. Cliqz. Ghostery. Retrieved September 2017 from http://www.ghostery.com

7. Rena Coen, Emily Paul, Pavel Vanegas, Alethea Lange, and G.S. Hans. 2016. A User-Centered Perspective on Algorithmic Personalization. University of California, Berkeley MIMS Final Project. https://www.ischool.berkeley.edu/projects/2016/user-centered-perspective-algorithmic-personalization.

8. Amit Datta, Michael Carl Tschantz, and Anupam Datta. 2015. Automated experiments on ad privacy settings. *PoPETS* 1 (2015), 92–112.

9. ECUAD. Visualizing Lightbeam. Retrieved September 2017 from http://research.ecuad.ca/lightbeam

10. EFF. Privacy Badger. Retrieved September 2017 from http://www.eff.org/privacybadger

11. Steven Englehardt and Arvind Narayanan. 2016. Online Tracking: A 1-million-site Measurement and Analysis. In *Proc. CCS*.

12. Steven Englehardt, Dillon Reisman, Christian Eubank, Peter Zimmerman, Jonathan Mayer, Arvind Narayanan, and Edward W. Felten. 2015. Cookies that give you away: The surveillance implications of web tracking. In *Proc. WWW*.

13. Motahhare Eslami, Karrie Karahalios, Christian Sandvig, Kristen Vaccaro, Aimee Rickman, Kevin Hamilton, and Alex Kirlik. 2016. First I like it, then I hide it: Folk Theories of Social Feeds. In *Proc. CHI*.

14. Motahhare Eslami, Aimee Rickman, Kristen Vaccaro, Amirhossein Aleyasen, Andy Vuong, Karrie Karahalios, Kevin Hamilton, and Christian Sandvig. 2015. "I always assumed that I wasn't really that close to [her]": Reasoning about Invisible Algorithms in News Feeds. In *Proc. CHI*.

15. Facebook. Your ad preferences. Retrieved September 2017 from https://www.facebook.com/ads/preferences/?entry_product=ad_settings_screen

16. Google. Ad settings. Retrieved September 2017 from https://adssettings.google.com/authenticated

17. Google. Add multiple topics with topic codes. Retrieved September 2017 from https://support.google.com/adwords/answer/156178?authuser=0#topic_codes

18. Google. Advertising Policies Help. Retrieved September 2017 from https://support.google.com/adwordspolicy/answer/143465

19. Google. Topics used for personalized ads. Retrieved September 2017 from https://support.google.com/ads/answer/2842480?hl=en

20. Kashmir Hill. February 16, 2012. How Target figured out a teen girl was pregnant before her father did. *Forbes*. http://www.forbes.com/sites/kashmirhill/2012/02/16/how-target-figured-out-a-teen-girl-was-pregnant-before-her-father-did/#2caffa3a34c6.

21. Alex Kantrowitz. September 15, 2017. Google Allowed Advertisers To Target People Searching Racist Phrases. *BuzzFeed*. https://www.buzzfeed.com/alexkantrowitz/google-allowed-advertisers-to-target-jewish-parasite-black.

22. Balachander Krishnamurthy, Delfina Malandrino, and Craig E. Wills. 2007. Measuring Privacy Loss and the Impact of Privacy Protection in Web Browsing. In *Proc. SOUPS*.

23. Balachander Krishnamurthy and Craig Wills. 2009. Privacy Diffusion on the Web: A Longitudinal Perspective. In *Proc. WWW*.

24. Mathias Lecuyer, Guillaume Ducoffe, Francis Lan, Andrei Papancea, Theofilos Petsios, Riley Spahn, Augustin Chaintreau, and Roxana Geambasu. 2014. XRay: Enhancing the Web's Transparency with Differential Correlation. In *Proc. USENIX Security*.

25. Mathias Lecuyer, Riley Spahn, Yannis Spiliopolous, Augustin Chaintreau, Roxana Geambasu, and Daniel Hsu. 2015. Sunlight: Fine-grained Targeting Detection at Scale with Statistical Confidence. In *Proc. CCS*.

26. Pedro Leon, Blase Ur, Richard Shay, Yang Wang, Rebecca Balebako, and Lorrie Cranor. 2012. Why Johnny can't opt out: A usability evaluation of tools to limit online behavioral advertising. In *Proc. CHI*.

27. Pedro G. Leon, Blase Ur, Yang Wang, Manya Sleeper, Rebecca Balebako, Richard Shay, Lujo Bauer, Mihai Christodorescu, and Lorrie Faith Cranor. 2013. What Matters to Users? Factors that Affect Users' Willingness to Share Information with Online Advertisers. In *Proc. SOUPS*.

28. Adam Lerner, Anna Kornfeld Simpson, Tadayoshi Kohno, and Franziska Roesner. 2016. Internet Jones and the Raiders of the Lost Trackers: An Archaeological Study of Web Tracking from 1996 to 2016. In *Proc. USENIX Security*.
29. Aleksandar Matic, Martin Pielot, and Nuria Oliver. 2017. ”OMG! How did it know that?”: Reactions to Highly-Personalized Ads. In *Proc. UMAP*.
30. Jonathan R. Mayer and John C. Mitchell. 2012. Third-party web tracking: Policy and technology. In *Proc. IEEE S&P*.
31. Aleecia M. McDonald and Lorrie Faith Cranor. 2010. Americans’ Attitudes About Internet Behavioral Advertising Practices. In *Proc. WPES*.
32. William Melicher, Mahmood Sharif, Joshua Tan, Lujo Bauer, Mihai Christodorescu, and Pedro Giovanni Leon. 2016. (Do Not) Track Me Sometimes: Users Contextual Preferences for Web Tracking. *PoPETS* 2 (2016), 135–154.
33. Mozilla. Lightbeam. Retrieved September 2017 from `http://www.mozilla.org/en-US/lightbeam/`
34. Oracle. Oracle Data Cloud Registry. Retrieved September 2017 from `http://www.bluekai.com/registry/`
35. Javier Parra-Arnau, Jagdish Prasad Achara, and Claude Castelluccia. 2017. MyAdChoices: Bringing transparency and control to online advertising. *ACM TWEB* 11, 1 (2017).
36. Angelisa C. Plane, Elissa M. Redmiles, Michelle M. Mazurek, and Michael Carl Tschantz. 2017. Exploring User Perceptions of Discrimination in Online Targeted Advertising. In *Proc. USENIX Security*.
37. Emilee Rader and Rebecca Gray. 2015. Understanding user beliefs about algorithmic curation in the Facebook news feed. In *Proc. CHI*.
38. Franziska Roesner, Tadayoshi Kohno, and David Wetherall. 2012. Detecting and Defending Against Third-party Tracking on the Web. In *Proc. NSDI*.
39. Sonam Samat, Alessandro Acquisti, and Linda Babcock. 2017. Raise the Curtains: The Effect of Awareness About Targeting on Consumer Attitudes and Purchase Intentions. In *Proc. SOUPS*.
40. Florian Schaub, Aditya Marella, Pranshu Kalvani, Blase Ur, Chao Pan, Emily Forney, and Lorrie Faith Cranor. 2016. Watching Them Watching Me: Browser Extensions’ Impact on User Privacy Awareness and Concern. In *Proc. USEC*.
41. Christopher A. Summers, Robert W. Smith, and Rebecca Walker Reczek. 2016. An Audience of One: Behaviorally Targeted Ads as Implied Social Labels. *Journal of Consumer Research* 43 (June 2016), 156–178.
42. Latanya Sweeney. 2013. Discrimination in online ad delivery. *CACM* 56, 5 (2013), 44–54.
43. Omer Tene and Jules Polonetsky. 2014. A Theory of Creepy: Technology, Privacy, and Shifting Social Norms. *Yale Journal of Law and Technology* 16, 1 (2014), 2.
44. Florian Tramer, Vaggelis Atlidakis, Roxana Geambasu, Daniel Hsu, Jean-Pierre Hubaux, Mathias Humbert, Ari Juels, and Huang Lin. 2017. FairTest: Discovering unwarranted associations in data-driven applications. In *Proc. EuroS&P*.
45. Joseph Turow, Jennifer King, Chris Jay Hoofnagle, Amy Bleakley, and Michael Hennessy. 2009. Americans Reject Tailored Advertising and Three Activities that Enable It. *SSRN* (2009).
46. Blase Ur, Pedro Giovanni Leon, Lorrie Faith Cranor, Richard Shay, and Yang Wang. 2012. Smart, useful, scary, creepy: Perceptions of online behavioral advertising. In *Proc. SOUPS*.
47. Jeffrey Warshaw, Nina Taft, and Allison Woodruff. 2016. Intuitions, Analytics, and Killing Ants: Inference Literacy of High School-educated Adults in the US. In *Proc. SOUPS*.
48. Craig E. Wills and Can Tatar. 2012. Understanding what they do with what they know. In *Proc. WPES*.
49. Yaxing Yao, Davide Lo Re, and Yang Wang. 2017. Folk Models of Online Behavioral Advertising. In *Proc. CCSW*.
50. Zhonghao Yu, Sam Macbeth, Konark Modi, and Josep M. Pujol. 2016. Tracking the Trackers. In *Proc. WWW*.

APPENDIX

APPENDIX A: STUDY 1 SURVEY INSTRUMENT

Part 1: Inferencing Perceptions

In the first part of the survey, we will ask some basic questions about your shopping and internet browsing habits.

- Choose one or more types of products that you are likely to shop for, either in physical stores or online:
  - Kitchen & dining
  - Cell phones & accessories
  - Sports & fitness
  - Camera & photo
  - Computer accessories
  - Toys & games
  - Textbooks
  - Gardening & lawn care
  - Pet supplies
  - Luggage & travel gear
  - Office products
  - Bed & bath
- What is your age?
  - 18-24 years old
  - 25-34 years old
  - 35-44 years old
  - 45-54 years old
  - 55-64 years old
  - 65 years or older
  - Prefer not to answer
- What gender do you identify with?
  - Male
  - Female
  - Other (Please state in the text box.)
  - Prefer not to answer
- Choose one or more races that you consider yourself to be:
  - White
  - Black or African American
  - American Indian or Alaska Native
  - Asian
  - Native Hawaiian or Pacific Islander
  - Other
  - Prefer not to answer
- Are you Spanish, Hispanic, or Latino or none of these?
  - Spanish, Hispanic, or Latino
  - None of these
  - Prefer not to answer
- Are you majoring in, or do you have a degree or job in, any of the following fields?*computer science; computer engineering; information technology; or a related field*
  - Yes
  - No
  - Prefer not to answer
- Choose one or more topics that you are likely to read about online: *[Topics which were closely related to the randomly-selected shopping category from the previous question that would be used in the notification were not shown.]*
  - Music & audio
  - Movies
  - Fashion
  - Cooking & recipes
  - Computers & electronics
  - Travel
  - Restaurants
  - Hiking & camping
  - Crafts
  - Job listings
  - Broadcast & network news
  - Gossip & tabloid news
  - Sports news
  - Business news
  - Politics
  - Parenting
  - Home improvement
  - Charity & philanthropy
  - Real estate
  - Pets & animals
  - Books

Advertising companies use different sources of information when deciding which ads to display to a specific user. They can use data about the typical visitors to a site or about a particular visitor, such as yourself. They can also use big-data techniques to predict what people with similar interests or demographics are likely to buy.

Companies then make inferences with this data. That is, they make some general conclusions based on this relevant information in order to decide which ads to show and where to show them.

In the next part of the survey, you will see a simulated ad and an explanation of how the advertiser chose to show the ad to you, which may have involved making inferences.

[shown if "prefer not to answer" for age or gender] You chose not to provide your age or gender. For the purposes of the following questions, you may pretend that you are a [randomly selected demographic]

Imagine you are looking up a word on dictionary.com and see the ad below. Hover over the ad to see an explanation of how the advertiser chose to show the ad to you. Read the explanation closely before moving on to the questions.

[Example ad + explanation shown]

- I would find it useful to have ads targeted to me for this reason, as opposed to any other reasons. *(7-point scale, Strongly agree to Strongly disagree)*
  - Why? Briefly explain. *(Free-response)*
- This notification gives me enough information to understand why an advertiser would show me this ad. *(7-point scale, Strongly agree to Strongly disagree)*
  - Why? Briefly explain. *(Free-response)*
- I would like to know whenever an ad is targeted to me for this reason. *(7-point scale, Strongly agree to Strongly disagree)*
  - Why? Briefly explain. *(Free-response)*
- Overall, how comfortable or uncomfortable are you with companies advertising to you for this reason? *(7-point scale, Extremely comfortable to Extremely uncomfortable)*
  - Why? Briefly explain. *(Free-response)*
- I think it is fair for a company to target ads for this reason. *(7-point scale, Strongly agree to Strongly disagree)*
  - Why? Briefly explain. *(Free-response)*
- I think it is fair for a company to collect the information necessary to target ads for this reason. *(7-point scale, Strongly agree to Strongly disagree)*
  - Why? Briefly explain. *(Free-response)*
- I would be annoyed by this type of ad targeting. *(7-point scale, Strongly agree to Strongly disagree)*
  - Why? Briefly explain. *(Free-response)*

General Perceptions Questions

The following questions are about your general perceptions of how advertising companies target advertisements to you, and not about any specific example or situation.

- I feel that I understand how advertising companies target advertisements to me. *(7-point scale, Strongly agree to Strongly disagree)*
- To your knowledge, how do advertisers choose which advertisements to show you? *(Free-response)*
- I feel that advertising companies are transparent about how they target advertisements to me. *(7-point scale, Strongly agree to Strongly disagree)*
  - Why? Briefly explain. *(Free-response)*
- I have used a web browser to access the internet. *(7-point scale, Strongly agree to Strongly disagree)*
- I feel that I understand how advertisers make inferences about my interests for the purpose of targeting advertisements to me. *(7-point scale, Strongly agree to Strongly disagree)*
- To your knowledge, how do advertisers make inferences about your interests for the purpose of targeting advertisements to you? *(Free-response)*

APPENDIX B: STUDY 2 SURVEY INSTRUMENT

Internet companies keep records of your activity as you browse the internet. Then they make inferences about your interests and demographics from these records to personalize the content you see online. This personalization can include showing you advertisements that are more relevant to you, showing you web search results that might reflect your interests, or choosing which posts to highlight on social media.

Part 1: General Questions About Inferencing

The following questions are about your general perceptions of how companies personalize your web experience, and not about any specific example or situation.

- I feel that I understand how companies personalize my web experience. *(7-point scale, Strongly agree to Strongly disagree)*
- To your knowledge, how do companies personalize your web experience? *(Free-response)*
- I feel that companies are transparent about how they personalize my web experience. *(7-point scale, Strongly agree to Strongly disagree)*
  - Why? Briefly explain. *(Free-response)*
- I feel that I understand how companies make inferences about my interests in order to personalize my web experience. *(7-point scale, Strongly agree to Strongly disagree)*
- To your knowledge, how do companies make inferences about your interests in order to personalize your web experience? *(Free-response)*
- In general, how would you expect your personalized web experience to differ depending on what inferences companies have made about your interests? *(Free-response)*

Part 2: Reactions to Specific Topics

[Looped ten times with different topics]

In the next part of the survey, you will see a series of 10 different topics that companies could make inferences about your level of interest in. Keep in mind that companies often infer that you are ”interested” in a topic based off of your internet activity, so they could think you are interested in something you dislike if you make web searches or look at content related to that topic.

For each topic you see, you will be asked questions about what you think about the topic with respect to inferencing and personalization.

The questions on this page are about the following topic in specific. The first line is the general topic category, and the second line is the specific topic.

[Category] - [Topic]

- How would you expect that someone's personalized web experience might differ depending on whether or not companies inferred that person was interested in *topic*? *(Free-response)*
- How often do you visit web pages related to *topic*?
  - Almost every day
  - Most days

  - Most weeks
  - Most months
  - Less than monthly
- I am interested in *topic*. *(7-point scale, Strongly agree to Strongly disagree)*
- I would be **comfortable with a company making an inference** about my level of interest in *topic*. *(7-point scale, Strongly agree to Strongly disagree)*
- I would be **comfortable with a company personalizing my web experience** based on an inference about my level of interest in *topic*. *(7-point scale, Strongly agree to Strongly disagree)*
- I would find it **useful to have my web experience personalized** based on an inference about my level of interest in *topic*, as opposed to my level of interest in other topics. *(7-point scale, Strongly agree to Strongly disagree)*
- Please briefly explain your answers to the questions above. *(Free-response)*

[The order of the next four questions is randomized]

- I would be comfortable with a **significant other** knowing about my level of interest in *topic*. *(7-point scale, Strongly agree to Strongly disagree)*
- I would be comfortable with a **close friend** knowing about my level of interest in *topic*. *(7-point scale, Strongly agree to Strongly disagree)*
- I would be comfortable with a **work supervisor** knowing about my level of interest in *topic*. *(7-point scale, Strongly agree to Strongly disagree)*
- I would be comfortable with an **employee of a store that I see from time to time but have never had a conversation with** knowing about my level of interest in *topic*. *(7-point scale, Strongly agree to Strongly disagree)*
- Please briefly explain your answers to the questions above. *(Free-response)*

Part 3: Demographics

- What is your age?
  - 18-24 years old
  - 25-34 years old
  - 35-44 years old
  - 45-54 years old
  - 55-64 years old
  - 65 years or older
  - Prefer not to answer
- What gender do you identify with?
  - Male
  - Female
  - Other (Please state in the text box.)
  - Prefer not to answer
- Choose one or more races that you consider yourself to be:
  - White
  - Black or African American
  - American Indian or Alaska Native
  - Asian
  - Native Hawaiian or Pacific Islander
  - Other
  - Prefer not to answer
- Are you Spanish, Hispanic, or Latino or none of these?
  - Spanish, Hispanic, or Latino
  - None of these
  - Prefer not to answer
- I have used a web browser to access the internet. *(7-point scale, Strongly agree to Strongly disagree)*
- Are you majoring in, or do you have a degree or job in, any of the following fields?*computer science; computer engineering; information technology; or a related field*
  - Yes
  - No
  - Prefer not to answer

APPENDIX C: STUDY 1 DETAILED REGRESSION RESULTS

Table 6: Proportional-odds logistic regression to determine correlations with higher agreement that "I would find it useful to have ads targeted to me for this reason, as opposed to any other reasons" (the **Useful** reaction statement). The reaction statement (dependent variable) was treated as ordinal. Among independent variables, participants specified their *age range* as one of the following categories: 18–24; 25–34; 35–44; 45–54; 55–64; 65+. We treated these categories as ordinal, fit linearly. We treated the remaining independent variables as nominal categorical. Omnibus test $p < .001$.

Factor	Baseline	β	SE	t	p
Why: Algorithm	Control	0.84	0.43	1.96	**.049**
Why: Visitors-Demographic	Control	0.24	0.42	0.56	.575
Why: Visitors-WrongDemographic	Control	-0.66	0.52	-1.27	.205
Why: Visitors-Inference	Control	1.46	0.47	3.13	**.002**
Why: Visitors-UnrelatedInterest	Control	0.09	0.45	0.19	.848
Why: You-Demographic	Control	0.16	0.46	0.35	.723
Why: You-WrongDemographic	Control	-0.27	0.47	-0.56	.572
Why: You-Interest	Control	1.07	0.44	2.44	**.015**
Why: You-UnrelatedInterest	Control	-0.11	0.43	-0.27	.789
Age Range	18-24 years old	-1.37	0.94	-1.45	.146
Gender: Male	Female	0.01	0.21	0.05	.957
Gender: Other	Female	-1.92	1.11	-1.73	.084
Technical Expertise: Not Answered	No	-0.43	0.83	-0.51	.607
Technical Expertise: Yes	No	-0.34	0.28	-1.21	.227
Race: White	Non-white	-0.31	0.17	-1.81	.070

Table 7: Proportional-odds logistic regression to determine correlations with higher agreement that "This notification gives me enough information to understand why an advertiser would show me this ad" (the **Informative** reaction statement). Omnibus test $p < .001$.

Factor	Baseline	β	SE	t	p
Why: Algorithm	Control	-0.76	0.43	-1.76	.079
Why: Visitors-Demographic	Control	0.32	0.44	0.73	.463
Why: Visitors-WrongDemographic	Control	0.10	0.50	0.19	.848
Why: Visitors-Inference	Control	1.01	0.47	2.18	**.029**
Why: Visitors-UnrelatedInterest	Control	-0.75	0.46	-1.65	.099
Why: You-Demographic	Control	0.23	0.48	0.49	.626
Why: You-WrongDemographic	Control	0.36	0.47	0.77	.439
Why: You-Interest	Control	1.03	0.44	2.34	**.019**
Why: You-UnrelatedInterest	Control	-0.80	0.45	-1.76	.078
Age Range	18-24 years old	-0.09	0.98	-0.09	.928
Gender: Male	Female	0.05	0.21	0.23	.821
Gender: Other	Female	0.76	1.14	0.67	.502
Technical Expertise: Not Answered	No	0.96	0.94	1.02	.306
Technical Expertise: Yes	No	-0.09	0.28	-0.31	.755
Race: White	Non-white	0.16	0.17	0.92	.357

Table 8: Proportional-odds logistic regression to determine correlations with higher agreement that “I think it is fair for a company to target ads for this reason” (the **Fair to Target** reaction statement). Omnibus test $p = .004$.

Factor	Baseline	β	SE	t	p
Why: Algorithm	Control	-1.02	0.44	-2.29	**.022**
Why: Visitors-Demographic	Control	-0.52	0.44	-1.16	.246
Why: Visitors-WrongDemographic	Control	-0.33	0.52	-0.64	.520
Why: Visitors-Inference	Control	0.18	0.48	0.39	.700
Why: Visitors-UnrelatedInterest	Control	-1.00	0.47	-2.11	**.035**
Why: You-Demographic	Control	-0.57	0.48	-1.20	.231
Why: You-WrongDemographic	Control	-1.07	0.46	-2.29	**.022**
Why: You-Interest	Control	-1.49	0.45	-3.31	$<$**.001**
Why: You-UnrelatedInterest	Control	-1.20	0.45	-2.67	**.007**
Age Range	18-24 years old	-0.12	0.93	-0.13	.895
Gender: Male	Female	0.36	0.21	1.69	.091
Gender: Other	Female	-0.83	1.20	-0.69	.490
Technical Expertise: Not Answered	No	-0.53	1.01	-0.53	.599
Technical Expertise: Yes	No	0.14	0.29	0.48	.629
Race: White	Non-white	0.09	0.17	0.52	.605

Table 9: Proportional-odds logistic regression to determine correlations with higher comfort responding to “Overall, how comfortable or uncomfortable are you with companies advertising to you for this reason?” (the **Comfortable** reaction statement). Omnibus test $p = .003$.

Factor	Baseline	β	SE	t	p
Why: Algorithm	Control	-0.48	0.43	-1.10	.270
Why: Visitors-Demographic	Control	-0.39	0.45	-0.88	.381
Why: Visitors-WrongDemographic	Control	-0.82	0.51	-1.61	.107
Why: Visitors-Inference	Control	0.81	0.47	1.73	.084
Why: Visitors-UnrelatedInterest	Control	-0.32	0.46	-0.69	.492
Why: You-Demographic	Control	-0.36	0.47	-0.75	.451
Why: You-WrongDemographic	Control	-0.96	0.46	-2.07	**.039**
Why: You-Interest	Control	-1.20	0.45	-2.69	**.007**
Why: You-UnrelatedInterest	Control	-0.73	0.44	-1.64	.101
Age Range	18-24 years old	0.45	0.93	0.49	.627
Gender: Male	Female	0.20	0.21	0.92	.356
Gender: Other	Female	-1.09	1.11	-0.98	.327
Technical Expertise: Not Answered	No	-1.35	1.01	-1.34	.181
Technical Expertise: Yes	No	-0.13	0.28	-0.47	.636
Race: White	Non-white	-0.03	0.17	-0.19	.852

APPENDIX D: STUDY 2 CATEGORIES AND DETAILED RESULTS

Inference Category	Comfort: Personalizing	Comfort: Inferencing
Computers & Electronics>Computer Hardware	2.08	2.17
Social Issues & Advocacy>Environmental Issues	1.56	1.56
Tourist Destinations>Beaches & Islands	1.30	1.40
Kitchen & Dining>Cookware & Diningware	0.93	1.00
Games>Computer & Video Games	0.89	1.22
Animal Products & Services>Pet Food & Supplies	0.88	1.00
News>Local News	0.88	0.62
Autos & Vehicles>Bicycles & Accessories	0.83	0.89
TV & Video>TV Comedies	0.83	0.67
Movies>Documentary Films	0.76	0.94
Arts & Entertainment>Comics & Animation	0.75	0.42
Home & Garden>Gardening & Landscaping	0.73	0.73
Food & Drink>Cooking & Recipes	0.60	0.70
Consumer Electronics>Cameras & Camcorders	0.56	0.62
Computer & Video Games>Music & Dance Games	0.53	0.53
Legal>Labor & Employment Law	0.50	0.67
Tourist Destinations>Historical Sites & Buildings	0.47	1.13
Specialty Travel>Vineyards & Wine Tourism	0.44	0.88
Special Occasions>Christmas	0.44	0.38
Consumer Electronics>Headphones	0.42	0.58
Vehicle Parts & Accessories>Engine & Transmission	0.33	0.67
Winter Sports>Skiing & Snowboarding	0.29	0.53
Apparel>Gems & Jewelry	0.29	0.57
Software>Business & Productivity Software	0.25	0.38
Events & Listings>Clubs & Nightlife	0.25	0.17
Jobs>Career Resources & Planning	0.25	0.12
Autos & Vehicles>Vehicle Shopping	0.23	0.69
Team Sports>Cricket	0.21	0.21
Business News>Financial Markets	0.19	0.25
Vision Care>Eyeglasses & Contacts	0.18	0.18
Computer & Video Games>Game Cheats & Hints	0.17	0.33
Arts & Entertainment>Celebrities & Entertainment News	0.17	0.17
Special Occasions>Thanksgiving	0.13	0.27
Apparel>Athletic Apparel	0.12	0.35
Home & Garden>Home Improvement	0.11	0.56
Mental Health>Learning & Developmental Disabilities	0.09	0.00
Movies>Action & Adventure Films	0.08	0.62
Events & Listings>Live Sporting Events	0.08	0.08
Social Issues & Advocacy>Privacy Issues	0.07	0.29
Health Conditions>Allergies	0.07	-0.07
Health>Alternative & Natural Medicine	0.06	0.12
Special Occasions>Valentine's Day	0.06	-0.06
Family & Relationships>Baby & Pet Names	0.05	-0.05
Individual Sports>Running & Walking	0.00	0.80
Energy & Utilities>Recycling	0.00	0.35
Cooking & Recipes>Vegetarian Cuisine	-0.06	-0.06
Travel>Cruises & Charters	-0.06	0.00
Music & Audio>Podcasting	-0.07	0.07
Computer Security>Antivirus & Malware	-0.07	-0.14
Face & Body Care>Make-Up & Cosmetics	-0.08	-0.08
Business & Industrial>Small Business	-0.09	-0.27
Special Occasions>Weddings	-0.11	-0.05
Jobs>Job Listings	-0.11	0.00
Military>Veterans	-0.11	0.00
Family & Relationships>Friendship	-0.11	0.00
Nutrition>Cholesterol Issues	-0.12	-0.12
Autos & Vehicles>Trucks & SUVs	-0.19	0.12

TV & Video>TV Soap Operas	-0.27	-0.27
Education>Special Education	-0.28	-0.28
Real Estate>Apartments & Residential Rentals	-0.29	0.06
Fitness>Bodybuilding	-0.31	-0.38
Home Furnishings>Lamps & Lighting	-0.35	0.18
Movies>Romance Films	-0.36	-0.14
Performing Arts>Broadway & Musical Theater	-0.36	-0.36
Team Sports>American Football	-0.36	0.05
Autos & Vehicles>Hybrid & Alternative Vehicles	-0.38	-0.25
Accounting & Auditing>Tax Preparation & Planning	-0.38	-0.46
Sports>Fantasy Sports	-0.40	0.00
Insurance>Auto Insurance	-0.42	-0.37
Software>Linux & Unix	-0.43	-0.36
Health Conditions>Skin Conditions	-0.44	-0.22
Nutrition>Vitamins & Supplements	-0.45	-0.45
Travel>Air Travel	-0.45	-0.73
Public Safety>Law Enforcement	-0.50	-0.36
Beauty & Fitness>Weight Loss	-0.52	-0.43
Computer & Video Games>Shooter Games	-0.53	-0.13
Insurance>Health Insurance	-0.57	-0.43
Special Occasions>Jewish Holidays	-0.57	-0.57
Apparel>Uniforms & Workwear	-0.60	-0.25
Education>Early Childhood Education	-0.62	-0.62
Education>Homeschooling	-0.65	-0.53
Combat Sports>Boxing	-0.67	-0.44
Health>Public Health	-0.67	-0.50
Health>Oral & Dental Care	-0.67	-0.67
Legal>Criminal Law	-0.67	-1.00
Education>Vocational & Continuing Education	-0.68	-0.77
Family & Relationships>Child Care	-0.69	-0.54
Real Estate>Bank-Owned & Foreclosed Properties	-0.70	-0.65
Religion & Belief>Skeptics & Non-Believers	-0.70	-0.70
Credit & Lending>College Financing	-0.71	-0.47
Health Conditions>Cold & Flu	-0.71	-0.76
Legal>Business & Corporate Law	-0.72	-0.72
Medical Facilities & Services>Surgery	-0.75	-0.75
Social Issues & Advocacy>Immigration Policy & Border Issues	-0.76	-0.52
Face & Body Care>Unwanted Body & Facial Hair Removal	-0.81	-0.56
Mental Health>Anxiety & Stress	-0.81	-1.00
Cosmetic Procedures>Cosmetic Surgery	-0.81	-1.12
Finance>Currencies & Foreign Exchange	-0.83	-0.61
Medical Facilities & Services>Physical Therapy	-0.86	-0.36
Credit & Lending>Credit Cards	-0.87	-0.87
Health Conditions>Heart & Hypertension	-0.87	-1.00
Health Conditions>Sleep Disorders	-0.89	-0.78
Social Issues & Advocacy>Discrimination & Identity Relations	-0.89	-0.78
Credit & Lending>Debt Management	-0.90	-0.80
Credit & Lending>Auto Financing	-0.91	-0.82
Family & Relationships>Romance	-0.91	-1.14
Home & Garden>Pest Control	-0.92	-0.77
Social Issues & Advocacy>Charity & Philanthropy	-0.92	-0.77
Health Conditions>Genetic Disorders	-0.92	-0.92
Autos & Vehicles>Motorcycles	-1.00	-0.47
Apparel>Undergarments	-1.00	-0.69
Law & Government>Military	-1.07	-1.07
Religion & Belief>Judaism	-1.07	-1.07
Health Conditions>Cancer	-1.08	-0.92
Health>Aging & Geriatrics	-1.08	-0.88
People & Society>Social Issues & Advocacy	-1.08	-1.08

Health Conditions>Vaccines & Immunizations	-1.12	-0.94
Legal>Drunk Driving Law	-1.13	-1.00
Reproductive Health>Sex Education & Counseling	-1.15	-0.46
Health Conditions>Obesity	-1.20	-1.07
Health Conditions>Asthma	-1.20	-1.40
Public Safety>Prisons & Corrections	-1.21	-1.16
Politics>Left-Wing Politics	-1.29	-1.10
Health Conditions>Parasites & Parasitic Diseases	-1.29	-1.18
People & Society>Disabled & Special Needs	-1.33	-1.14
Apparel>Swimwear	-1.33	-1.20
Aging & Geriatrics>Alzheimer's Disease	-1.33	-1.44
Mental Health>ADD & ADHD	-1.39	-1.11
Politics>Campaigns & Elections	-1.40	-0.80
Reproductive Health>OBGYN	-1.43	-1.86
Nursing>Assisted Living & Long Term Care	-1.44	-1.62
Family & Relationships>Pregnancy & Maternity	-1.44	-1.39
Music & Audio>Rap & Hip-Hop	-1.45	-0.73
Politics>Right-Wing Politics	-1.46	-1.58
Health Conditions>Thyroid Conditions	-1.50	-1.00
Medical Devices & Equipment>Mobility Equipment & Accessories	-1.50	-1.57
Special Occasions>Islamic Holidays	-1.56	-1.06
Religion & Belief>Hinduism	-1.56	-1.12
Hair Care>Hair Loss	-1.56	-1.25
Ethnic & Identity Groups>Gay-Lesbian-Bisexual-Transgender	-1.57	-1.48
Religion & Belief>Astrology & Divination	-1.58	-1.53
Social Services>Welfare & Unemployment	-1.60	-1.73
Social Issues & Advocacy>Same-Sex Marriage	-1.67	-1.39
Religion & Belief>Islam	-1.69	-1.08
Substance Abuse>Drug & Alcohol Treatment	-1.69	-1.69
Social Issues & Advocacy>Reproductive Rights	-1.71	-1.62
Reproductive Health>Erectile Dysfunction	-1.71	-1.71
Reproductive Health>Infertility	-1.71	-2.00
Music & Audio>Christian & Gospel Music	-1.75	-1.69
Religion & Belief>Christianity	-1.85	-1.55
Health Conditions>Eating Disorders	-1.88	-1.12
Mental Health>Depression	-1.88	-1.88
Reproductive Health>Sexually Transmitted Diseases	-1.88	-2.00
Health Conditions>AIDS & HIV	-1.93	-1.86
Family & Relationships>Divorce & Separation	-1.94	-1.94
Family & Relationships>Adoption	-2.00	-1.85
Legal>Family Law	-2.08	-1.75
Substance Abuse>Smoking & Smoking Cessation	-2.11	-2.33
Reproductive Health>Birth Control	-2.23	-2.46
Online Communities>Dating & Personals	-2.71	-2.71

Table 10: The full list of 160 inference categories we tested in Study 2, as well as the detailed results. The two numerical columns present the mean responses on a 7-point Likert scale ("strongly agree" coded as 3, "strongly disagree" coded as -3), across all participants who saw that category, to the following two statements. **Comfort: Personalizing** was the statement "I would be comfortable with a company personalizing my web experience based on an inference about my level of interest in *topic*." **Comfort: Inferencing** was the statement "I would be comfortable with a company making an inference about my level of interest in *topic*."

APPENDIX E: STUDY 2 DETAILED REGRESSION RESULTS

Table 11: Mixed-model ordinal regression to determine correlations with higher agreement that "I would be comfortable with a **significant other** knowing about my level of interest in *topic*."

Factor	β	**SE**	t	p
Comfort: Personalizing	3.00	0.26	11.40	$<$**.001**
Mean Comfort: Personalizing	0.33	0.07	4.93	$<$**.001**
Interest	0.84	0.25	3.36	$<$**.001**
Visit Related Pages	0.31	0.36	0.86	0.390

Table 12: Mixed-model ordinal regression to determine correlations with higher agreement that "I would be comfortable with a **close friend** knowing about my level of interest in *topic*."

Factor	β	**SE**	t	p
Comfort: Personalizing	3.61	0.00	810.87	$<$**.001**
Mean Comfort: Personalizing	0.50	0.00	127.57	$<$**.001**
Interest	0.88	0.00	210.18	$<$**.001**
Visit Related Pages	0.32	0.00	277.41	$<$**.001**

Table 13: Mixed-model ordinal regression to determine correlations with higher agreement that "I would be comfortable with a **work supervisor** knowing about my level of interest in *topic*."

Factor	β	**SE**	t	p
Comfort: Personalizing	3.71	0.20	18.64	$<$**.001**
Mean Comfort: Personalizing	0.61	0.06	10.20	$<$**.001**
Interest	0.13	0.21	0.62	0.536
Visit Related Pages	0.05	0.33	0.14	0.891

Table 14: Mixed-model ordinal regression to determine correlations with higher agreement that "I would be comfortable with **an employee of a store that I see from time to time but have never had a conversation with** knowing about my level of interest in *topic*."

Factor	β	**SE**	t	p
Comfort: Personalizing	4.04	0.22	18.75	$<$**.001**
Mean Comfort: Personalizing	0.67	0.06	11.07	$<$**.001**
Interest	-0.09	0.21	-0.44	0.663
Visit Related Pages	-0.37	0.32	-1.15	0.252