Taking Data Out of Context to Hyper-Personalize Ads: Crowdworkers' Privacy Perceptions and Decisions to Disclose Private Information

Julia Hanson^{†*}, Miranda Wei^{†◇*}, Sophie Veys[†], Matthew Kugler[▽], Lior Strahilevitz[†], Blase Ur[†]

[†] University of Chicago, [◇] University of Washington, [▽] Northwestern University

{jchanson,weim,sophie,lior,blase}@uchicago.edu, matthew.kugler@law.northwestern.edu

ABSTRACT

Data brokers and advertisers increasingly collect data in one context and use it in another. When users encounter a misuse of their data, do they subsequently disclose less information? We report on human-subjects experiments with 25 in-person and 280 online participants. First, participants provided personal information amidst distractor questions. A week later, while participants completed another survey, they received either a robotext or online banner ad seemingly unrelated to the study. Half of the participants received an ad containing their name, partner's name, preferred cuisine, and location; others received a generic ad. We measured how many of 43 potentially invasive questions participants subsequently chose to answer. Participants reacted negatively to the personalized ad, yet answered nearly all invasive questions accurately. We unpack our results relative to the privacy paradox, contextual integrity, and power dynamics in crowdworker platforms.

Author Keywords

hyper-personalization; targeted advertising; creepy; user study

CCS Concepts

•Security and privacy → Usability in security and privacy;

INTRODUCTION

Much of modern advertising is personalized to its recipients. In past decades, attempts at personalized advertising might have involved choosing a particular location for a billboard or running an ad in a given periodical. The rise of web tracking [15, 32], email tracking [14], and data marketplaces [5, 43], however, enables far greater personalization. Based on data collected about users, advertisers infer their interests and demographics [36] to target ads [51]. The ad ecosystem also incorporates personally identifiable information and offline purchasing data funneled through data brokers [4].

*Co-authors Hanson and Wei contributed equally to this work.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the Owner/Author.

CHI '20, April 25–30, 2020, Honolulu, HI, USA

ACM ISBN 978-1-4503-6708-0/20/04.

© 2020 Copyright is held by the owner/author(s).

<https://doi.org/10.1145/3313831.3376415>

Personalized advertising is complicated from the privacy perspective. While consumers perceive targeted ads as useful, they can also find them creepy [60, 61]. This ecosystem is mostly opaque to consumers [16], yet it sometimes enters the public consciousness, such as when Target inferred a customer's pregnancy before even her family found out [13, 20]. Future ads will undoubtedly be even more targeted [46]. We use the term *hyper-personalization* to describe this logical next step that deeply embeds personal information in ads. While currently rare, the history of increasing personalization [32, 36] suggests such ads are just around the corner.

We used a human-subjects experiment to investigate how such hyper-personalized advertising might impact consumers' behaviors in protecting their own privacy, particularly if this personalization is perceived as a misuse of data collected in one context and used in another. Our core motivation was to capture participants' real-time reactions and responses to hyper-personalized ads. Using an IRB-approved, deception-based study, we surreptitiously collected participants' first names, their romantic partners' names (or that they were single), their preferred type of cuisine, and their town-level location. We randomly assigned some participants to receive an ad with this information, either as a typical online banner ad embedded in our survey software or as a robotext sent to their mobile phone from a short-code number. Robotexts represent an emerging medium for ad delivery and also minimized the chance the participant would associate the ad with the study. Automated calls (robocalls) and text messages (robotexts) have become more frequent in the past year [10], despite questions about their ethics [40] and widespread consumer aversion [58].

We hypothesized participants would be less likely to accurately reveal personal information if they had received a hyper-personalized ad, as opposed to a generic ad. We expected personalized ads would cause feelings of privacy invasion, and that those feelings would lead participants to choose "prefer not to say" to stop the spread of their personal information. Crowdworkers face significant privacy threats through disclosure [66]; personal data given to seemingly trusted sources can be weaponized against users. In fact, the Facebook data in the Cambridge Analytica scandal originated in a research study conducted on Mechanical Turk [38].

After revealing the deception through a debriefing, we asked the participant how they perceived the ad they received, as

well as whether they had seen it in the first place. A total of 25 in-person and 280 online participants completed this protocol. The nuanced manipulation of sending an ad worked as intended. The majority of participants reported in our debriefing that they had received and read the ad at the intended point of the study and had not suspected the ad was study-related.

Participants reacted strongly to hyper-personalization. Roughly half of the participants who received a personalized ad volunteered that they had a strong negative reaction to the personalization, such as feeling creeped out, angry, or otherwise alarmed. Despite these reactions, we found that receiving the hyper-personalized ad, as opposed to the generic ad, had no significant effect on participants' subsequent decisions about disclosing personal information. Nearly all participants accurately answered questions about their religion, finances, and votes in recent elections. Over half of participants chose to disclose where they were born, physical ailments, and habits regarding masturbation and pornography. Some even provided their mother's maiden name and their home address.

After revealing the deception, we probed why participants answered these questions accurately, and at all. We found that an ensemble of previously studied phenomena encouraged disclosure. Participants reported not perceiving risks in disclosing information. They said they trusted the crowdworking platform, our survey software, and research studies broadly. Notably, we did not present ourselves as affiliated with any academic institution. They also reported that financial pressures as crowdworkers and overall feelings of helplessness in the modern data economy encouraged disclosure. These results cast doubt on frameworks that rely on crowdworkers to protect their privacy by self-regulating information disclosure.

RELATED WORK

Data used for ad personalization comes from many sources, including online tracking [15,36] and sale by data brokers [4,43]. However, information collection and online tracking are opaque to consumers [8,36,47]. Prior work has found that many consumers oppose personalized advertising, particularly when told how advertisers gather the information used for personalization [60,61,67]. Recipients of personalized ads often dislike the conclusions that the ads make about them [60]. Ads can also be racially discriminatory [45,54] and involve sensitive categories (contrary to companies' public statements) [30]. Increased ad personalization is associated with increased consumer discomfort [34], though preferences vary based on the data inferred [12,67] and by platform [68]. Additionally, users tend to view ads more favorably if they feel in control of their privacy or that an ad is relevant [28,59].

Practices related to personalized advertising are largely governed by advertising industry self-regulation [55]. Privacy advocates have called, unsuccessfully, for increased regulation of personalized advertising. Advertisers generally oppose such restrictions because of the potential decrease in advertising efficacy. Indeed, advertisements in compliance with EU privacy laws limiting the use of consumer data were found to be less effective [19]. In response, some consumers use privacy-protective browser extensions [29,31,35–37,48], though these tools are only partially effective [5,64].

The privacy calculus model proposes that disclosure decisions reflect an analysis of risks and benefits [11]. When people perceive risk, they are less willing to disclose information [3]. They disclose more when promised a benefit [22] or if privacy protection is emphasized [2]. However, there is a well-known disparity (the “privacy paradox”) between expressed intentions about privacy and actual behaviors [53,65]. People's decisions about disclosing sensitive information are highly context-dependent [3,9,17,24]. People also disclose more when they observe others' disclosures [57]. In one study, a computer successfully encouraged participants to divulge personal information through reciprocity (the computer sharing personal information) and gradually increasing sensitivity [41], although other work [1] contradicts the latter. A study on willingness to unlock smartphones for researchers found participants' actual rate of compliance with such requests was far higher than others' predictions of what people would do [52].

Prior work studying platforms like Amazon's Mechanical Turk has found that workers' privacy concerns and decisions to disclose information are also highly contextual. Although crowdworkers tend to be more privacy conscious than the general population [25], they may share information despite their concerns, motivated by their economic needs or influenced by power imbalances in worker-requester relationships [50]. Assessing the risks for a task and deciding whether to disclose information is “invisible labor” crowdworkers must perform when selecting and completing tasks [50].

The contextual integrity model proposes that privacy harms will occur when *transmission flows* of information (the purpose for disclosure) are broken. Implementing procedural fairness, an organization-level interpretation of CI's transmission principle, can decrease individuals' privacy concerns [7].

METHOD

We designed and conducted a two-part deception study in which participants received a targeted ad that incorporated their own personal information at a specific point during the study. We measured how the receipt of such a hyper-personalized ad, compared with a control group who received a generic ad, impacted participants' responses to potentially invasive questions. We augmented this experiment with qualitative investigations of participants' decisions to disclose information and their perceptions of personalization in advertising.

We conducted our in-person data collection and an initial round of online data collection with our university affiliation on the recruitment script, consent form, and study URL. Our initial data analyses revealed participants' trust in our university affiliation and ethics review process as a reason some chose to disclose information. Therefore, we revised our protocol to perform another round of online data collection in which we claimed to be the “Institute for Interests and Demographics Research” (IIDR) with no further information. We report on this second round of online data collection in this paper, though we found nearly identical results in the first round. We established our fake organization by registering a domain name to host our surveys and posing as IIDR in our consent form and recruitment text. Our IRB approved both our initial protocol and all changes for this additional data collection.

Pre-study

We used a pre-study to identify potentially invasive questions to present in Part 2 of the main study. Ideal questions would have relatively consistent invasiveness across answer choices and a quickly recallable answer. We derived these prospective questions from the American Housing Survey, the US Census, quizzes found in Cosmopolitan magazine and on Facebook, surveys from Pew Research, and group discussions.

To empirically validate each question and down-select to a smaller number of questions, we conducted a pre-study. We recruited 63 pre-study participants on Amazon’s Mechanical Turk, a different recruitment platform than we use for our main study. We are not aware of published, validated scales measuring invasiveness. We gauged invasiveness with yes/no/don’t know responses to “If given the opportunity, I would choose not to answer this question,” which we developed through cognitive interviews. We eliminated questions with high standard deviation for reported comfort in answering the question or low confidence in reported ability to recall the answer. Through this process, we chose 43 potentially invasive questions.

Recruitment for the Main Study

To balance deep insights and large-scale data, we conducted both in-person and online sessions. Across both types of sessions, 305 participants completed our full protocol. We required participants be 18+ years old, live in the USA, and have a cell phone capable of receiving text messages. As is necessary in deception studies, we advertised the study with a title and explanation that differ from the true purpose. We initially told participants the study was titled “The Demographics, Interests, and Experiences of Americans.” We revealed the actual title (“The Impact of Hyper-Personalized Marketing on Information Disclosure”) in the Part 2 debriefing.

We recruited in-person participants on Craigslist. We compensated them with a \$1 Amazon gift card for completing Part 1 online and a \$19 gift card for completing Part 2 in person. 25 participants completed Part 2 in a private room in their choice of our institution’s campus or a local public library.

We recruited a total of 470 online participants for Part 1 through Prolific, which has emerged as a frequently preferred [44] alternative to Amazon’s Mechanical Turk. 280 of these participants completed Part 2. 25 participants began Part 2, but either timed out or returned the study. Prolific participants were also required to have a 95% approval rating.

Main Study Part 1

The purpose of Part 1 was to collect information we could use to personalize ads in Part 2 without arousing suspicion about the true nature of the study. We accomplished this by asking mainly distractor questions in Part 1. To give participants time to forget Part 1, we opened Part 2 over a week later.

Part 1 took the form of a five-minute online survey. As professional survey platforms generally do not show ads, we worried that participants who saw a banner ad in Part 2 might find it suspicious. To habituate participants to seeing ads, we displayed a simulated banner ad for chewing gum, which we labeled as a sponsor ad (see online supplementary materials).

Table 1: Study conditions.

Condition	Recruitment	Location	Delivery	Personalization
<i>Lab-Text-Personalized</i>	Craigslist	In-person	Robotext	Personalized
<i>Lab-Text-Generic</i>	Craigslist	In-person	Robotext	Generic
<i>Online-Text-Personalized</i>	Prolific	Online	Robotext	Personalized
<i>Online-Text-Generic</i>	Prolific	Online	Robotext	Generic
<i>Online-Banner-Personalized</i>	Prolific	Online	Banner	Personalized
<i>Online-Banner-Generic</i>	Prolific	Online	Banner	Generic

The study began with a consent form and a brief overview of Part 1. We then asked the participant to provide their mobile phone number, which we told them we would use to send study-related text messages. This was followed by the standardized Positive Affect and Negative Affect Schedule (PANAS), which has been shown to reliably measure mood [27, 62, 63]. We then asked participants a series of questions designed to elicit the following information needed to personalized an ad in Part 2: their first name, their preferred type of cuisine, their relationship status, and, if in a relationship, their partner’s first name. In addition to a multiple-choice question about the kind of restaurant they might want to visit, we also asked distractor questions about a movie they might want to see and a store they might want to visit. To elicit the participant’s relationship status and, if applicable, partner’s first name, we asked about the tech savviness of a coworker, a family member, and a (current, past, or aspirational) significant other. We approximated the city or town from which they completed Part 1 using Qualtrics’s IP geolocation feature.

Conditions

All participants completed Part 1 of the study online in an identical procedure. However, each participant experienced one of six variants of Part 2 of the study, differing in where the session was conducted, how the ad was delivered, and whether the ad was personalized. We identify these groups using a three-part name indicating these respective differences, as defined below. Table 1 summarizes these conditions.

Location. **Lab** participants were recruited via Craigslist and completed Part 2 in person. **Online** participants were recruited via Prolific and completed Part 2 remotely.

Delivery. Participants received an advertisement in Part 2 as either a robotext (termed **text**) or a **banner** ad. Using information collected in an online survey to personalize a text message makes the flow of personal information less obvious. Because we wanted to ask semi-structured interview questions to better understand the novel medium of robotexted ads, all in-person participants received a robotext. To understand how robotexts, which are becoming increasingly common for political advertising, compare to banner ads, which are already ubiquitous online, we randomly assigned online participants to receive either robotext or banner ads. The text of the ad was identical in both delivery methods.

For banner conditions, the banner ad was displayed as the first page of the survey in Part 2. The ad was formatted to look like it was not part of the survey itself, as shown in the online supplementary materials. For text conditions, we configured our survey software to call an API for sending text messages from commercial five-digit short codes.

(a) Personalized ad (if in a relationship)

(b) Personalized ad (if single)

(c) Generic ad

Figure 1: Example banner ads shown to participants in Part 2.

(a) Personalized

(b) (If single)

(c) Generic

Figure 2: Example robotext ads.

Personalization. Because our goal was to understand the impact of receiving a hyper-personalized ad, we randomly assigned each participant to receive either a hyper-personalized (termed **personalized**) or an analogous, non-personalized (**generic**) ad. The ads promoted a fictional restaurant reservation service called SUPEREats. In comparison to the generic ads (Figure 1c, Figure 2c), the personalized ads (Figure 1a–1b, Figure 2a–2b) invoked the following information from Part 1: the recipient’s first name, their preferred type of restaurant, their relationship status (and, if applicable, partner’s name), and the approximate city or town where they completed Part 1.

Main Study Part 2

In Part 2, we sought to deliver the ad specified by the participant’s condition, measure the participant’s subsequent decisions about answering the potentially invasive questions, and elicit their broad impressions of hyper-personalization. Part 2 took approximately 30 minutes.

When the participant began Part 2, our software automatically triggered the banner ad (Figure 1) or robotext (Figure 2). Participants then completed the PANAS questionnaire. Next, participants answered questions about their general technology usage. Some of these questions were intended to require participants to look at their phone (e.g., “What is the current battery percentage of your phone?”) to increase the likelihood that those who had received a robotext would see the ad.

The 43 potentially invasive questions followed. Because a survey consisting only of very invasive questions might spur participants to drop out or cause similar biases, we selected the 43 questions such that they reflected a range of sensitivity. Reflecting prior work on the impact of sequencing [41], we ordered questions from least to most sensitive in our final battery. This set of questions began with those participants in our pre-study overwhelmingly indicated they would answer, such as

current voter registration status, and increased in sensitivity to questions about financial circumstances, illnesses, experiences with pornography, personal information, and more.

Each of these 43 questions had a conspicuous “prefer not to say” option, as shown in Figure 3. We called attention to this option at the start of the section and explicitly told participants they would not be penalized for selecting it. We hypothesized that participants who received the personalized ad and found it invasive would answer fewer questions in an act of privacy protectiveness. To protect participants’ actual privacy, we wrote custom JavaScript to enable the Qualtrics survey platform to delete participant’s responses to these potentially invasive questions upon survey submission. At no point were these responses available to us. We only recorded whether they responded to the question or chose “prefer not to say.”

Afterwards, we again presented the PANAS questionnaire to test our hypothesis that the potentially invasive questions would negatively affect the participant’s mood. To contextualize our results, we then asked questions about past experiences with, and opinions of, personalized advertising.

At this point, all participants were debriefed about the true nature of the study either by a researcher (in-person participants) or through the survey flow (online participants). We revealed the real purpose of the study, as well as the fact that we, the researchers, had sent the online or robotext ad. In the second round of our data collection, we additionally revealed that we were not conducting this research as the fake IIDR, but as our university. In both rounds, we followed with questions about the extent to which participants thought about our affiliations.

Participants could also have protected their privacy by answering questions inaccurately. Therefore, after the debriefing, we reminded the participant that we were only interested in their decisions about disclosing personal information. We then had them review the accuracy of their answers to the potentially invasive questions before we permanently deleted these raw responses. We made clear that they would not be penalized for having given inaccurate responses. We asked them to select which questions, if any, they had answered inaccurately.

Analysis Method and Metrics

We performed both quantitative and qualitative analyses. Our primary response variable was the number of potentially invasive questions participants answered, and we wanted to know how this number varied based on the personalization

What is the first and last name of your closest friend?

Prefer not to say

Figure 3: Example potentially invasive question in Part 2. All had a conspicuous “prefer not to say” option.

of the ad, the ad’s medium (robotext or banner ad), and between in-person and online participants. We investigated this question using the Mann-Whitney U test pairwise across matched conditions (e.g., *Online-Text-Personalized* vs. *Online-Text-Generic* and *Online-Text-Personalized* vs. *Online-Banner-Personalized*). We conducted this analysis for the total number of questions answered, as well as the total number of questions that participants reported answering accurately.

We gave the PANAS questionnaire three times. As this is a repeated-measures design, we built a Linear Mixed-Effects Model [33] with time, condition, and the interaction thereof as independent variables, and the participant as a random effect.

We analyzed free-response answers (online participants) and interview answers (in-person participants) qualitatively using standard open and axial coding procedures with two independent coders. This process was informed by discussions among members of the research team after they had read the responses and discussed prospective themes they had observed.

Ethics and Protection of Human Subjects

As our study was a deception protocol involving participants’ personal information, we took great care in designing an ethical protocol. We iteratively adjusted our protocol through extensive in-person pilot testing to understand how participants would be likely to react to various elements of the study, including the receipt of the personalized ad. We made every effort to minimize the information we collected in the first place, and to delete it when no longer needed. All protocols were approved by our institution’s IRB.

As our deception protocol involved using participants’ personal information in simulated ads and sending text messages under false pretenses, we obtained consent from all participants to receive study-related text messages. As with all deception protocols, we provided a substantial debriefing toward the end of Part 2. The debriefing explained the true purpose of our study, why deception was necessary, and what data was actually retained. To ensure that any participants who dropped out of the study before seeing the debriefing were informed of the deception, we messaged all Prolific participants who began Part 2 but did not complete it with the relevant parts of the debriefing. We also used Prolific’s deception filter, which means that the study was only shown to prospective participants who had previously indicated to Prolific that they were willing to participate in deception studies.

Typically, crowdworking sites like Prolific do not permit researchers to collect PII. However, we discussed our study with Prolific, and they permitted us to do so for this particular study. To avoid surprising participants, we explained to participants in the brief recruitment text describing our study on Prolific

that we would be requesting personally identifiable information as part of the study. Even though all participants were located in the US, we complied with GDPR as Prolific is based in the UK. We also linked a privacy policy from the consent form and debriefing materials explaining our use of the data.

Limitations

Because we made clear in our recruitment text that we would request personally identifiable information, we may have biased our sample toward people who are already more willing to disclose personal information. Similarly, privacy-sensitive individuals may have dropped out of the study if they felt the questions were too invasive. The vast majority of participants who started Part 2 completed it, tempering this concern.

The high response rate to potentially invasive questions could also be a result of the size of our intervention. One hyper-personalized ad may be sufficient to elicit emotional reactions, but not changes in behavior. Alternatively, participants’ response rate could have been informed by an opposing attitude of resignation, that their data is no longer private anyway and that there is no use withholding information. This would support an argument that broad usage of hyper-personalization inflicts privacy harms. The truthfulness of responses to sensitive questions is partially a function of survey design, which is why we emphasized the “prefer not to answer” option. When asking about the accuracy of answers, we also emphasized there was no penalty for having given inaccurate responses.

We expected that some participants would not receive our ad treatment as we intended. Thus, following the debriefing, we asked participants about whether they saw the ad, as well as whether they suspected it might be connected to the study. We excluded participants who suspected the ad was study-related or who did not see the ad from our main analyses. Nonetheless, as this information was self-reported and people tend to overreport socially desirable behaviors [56], participants could have lied about these aspects, as well as whether they had answered the potentially invasive questions accurately. Participants of particular characteristics might have been more likely to miss the ad, biasing the sub-sample we analyzed.

RESULTS

First, we summarize participants’ demographics. We then review the effectiveness of our deception, observing that most participants saw the ad and did not suspect it to be study-related. We then describe participants’ highly negative reactions to the personalized ads. However, our hypothesis that participants would disclose less information if they saw a personalized ad was ultimately not supported. Instead, participants accurately answered most of the potentially invasive questions. We contextualize these results by analyzing why participants reported choosing to answer these questions. Trust in crowdwork platforms and in research studies played a key role in disclosure decisions, as did a perceived lack of risk.

Initially, we surmised our university affiliation influenced the high response rate to the potentially invasive questions. People might be more likely to trust a university with their information than other parties. This would align with Milgram’s experiments on obedience [39], which found higher compliance

Table 2: The number of participants who *completed* Parts 1 and 2, said they *saw the ad*, and said they both saw the ad and *did not suspect* at the time it was study-related.

Condition	Completed	Saw Ad	Did Not Suspect
<i>Online-Banner-Personalized</i>	66	62 (93.9%)	43 (65.2%)
<i>Online-Banner-Generic</i>	72	67 (93.1%)	42 (58.3%)
<i>Online-Text-Personalized</i>	64	45 (70.3%)	43 (67.2%)
<i>Online-Text-Generic</i>	78	52 (66.6%)	49 (62.8%)
<i>Lab-Text-Personalized</i>	14	10 (71.4%)	10 (71.4%)
<i>Lab-Text-Generic</i>	11	10 (91.0%)	10 (91.0%)

when conducted at Yale (65%) than at a fictitious independent lab (48%). To our surprise, the results from our second round of data collection, branded as the fictitious Institute for Demographics Research, were very similar to the first, university-branded round. For this reason, we report mainly on the latter round in this section, unless otherwise noted.

Participants

A total of 25 in-person and 280 online participants completed both Parts 1 and 2. 52.9% identified as female, 43.9% as male, and 3.2% as non-binary. Our participant sample skewed young, with 28.6% between 18 and 24, 38.9% between 25 and 34, 20.4% between 35 and 44, and 12.1% age 45 or older. Participants' educational attainment skewed slightly more educated than the general U.S. population [49]. Most participants' annual household income was between \$20,000 and \$49,999 (31.4%) or \$50,000 and \$99,999 (35.4%).

Deception Effectiveness

Because we aimed to capture data reflecting participants' naturalistic reactions to their information appearing in a personalized ad, the deception aspect of our study was crucial. To ensure that we obtained authentic reactions and responses, our debriefing evaluated whether participants were truly deceived.

Most participants reported seeing the ad that was robotexted or displayed to them at the intended point in the study. Participants were sent robotexts in both online and in-person conditions: In total, 80.0% of in-person and 68.3% of online robotext participants reported reading at least part of the robotext before answering the potentially invasive questions. Many participants saw these ads right away or during the tech use questions we included to encourage participants to look at their phone. Participants who did not read the text explained they had notifications turned off or were in do-not-disturb mode. 93.4% of participants reported seeing the online banner ad at the start of Part 2 and later recalled some details of the ad. Many of the remaining participants reported seeing what looked like an ad with a countdown timer, but ignored the ad.

Most participants reported that they did not suspect the ad they saw was part of our study (Table 2). Only 9 out of the 117 online and in-person participants who saw a robotext reported suspecting it was study-related when they first received it. Other participants sent robotexts did not suspect it was study-related at any point (77.3%) or only suspected it was study-related when the survey questions became increasingly invasive or asked about advertising (13.4%).

Participants shown a banner ad were slightly more suspicious: 34.1% reported suspecting the ad was study-related when they first saw it, 45.0% said they did not suspect it was study-related at any point, and 17.8% suspected it was study-related only later in the study, often due to questions about advertising.

For the remainder of the results section, we report only on participants who saw at least some of the ad and did not suspect it was study-related, unless otherwise noted. Furthermore, as the sample size for in-person participants was small, we report quantitative results only for online participants. In sum, 63.2% of the online participants across conditions met this threshold.

Reactions to the Ad

As part of the debriefing, we asked participants about their initial reactions to the ad. To avoid priming our participants, we asked about their initial reactions with a free-response question: "What was your initial reaction to the ad?" Many participants who received personalized ads reported a range of strong, negative reactions to the ad. 53.4% of *Online-Banner-Personalized* participants and 44.2% of *Online-Text-Personalized* participants volunteered feeling a combination of either scared, concerned, shocked/surprised, creeped out, or uncomfortable in their initial reaction to the ad. P211 (*Online-Banner-Personalized*)'s response features several of these emotions: "I was quite alarmed by the ad, wondering how it knew my spouse's name and my location. I was disturbed and put off by it." Others explicitly expressed feelings of privacy invasion and fear, like P135 (*Online-Text-Personalized*):

I felt immediately concerned for myself and my girlfriend. Almost threatened by the knowledge that some company had on me. I felt like my privacy was being invaded and that companies were using leaked information. . .

In contrast, no *Online-Banner-Generic* or *Online-Text-Generic* participants reported these emotions in their reactions. Instead, those who saw the generic banner ad most commonly reported indifference (40.5%), confusion about why a survey had an ad (33.3%), or annoyance (14.2%). *Online-Banner-Personalized* participants also reported reacting with indifference (18.6%), confusion about ads on surveys (13.9%), or annoyance (2.3%). Participants who received a generic robotext also reported indifference (16.3%) or annoyance (10.2%), but also wondered how they had been added to a marketing list (22.4%). 39.7% of *Online-Text-Generic* and 16.3% of *Online-Text-Personalized* participants reported thinking the text was spam.

Many participants in *Online-Banner-Personalized* and *Online-Text-Personalized* reported questioning the information flows that underpinned the ad: "[I] immediately wondered how my personal information got out" (P133, *Online-Text-Personalized*). A few participants rationalized how SUPEReats collected their information. Some blamed large tech companies: "I just assumed Google was behind it and they seem to know everything" (P94, *Online-Banner-Personalized*). Others hypothesized that Yelp or GrubHub had shared their information. P49 (*Lab-Text-Personalized*) had just heard about Facebook's new dating feature and thus attributed the robotext to Facebook, citing the ad's inclusion of their relationship status. Others linked it to their online browsing behavior: "I had

Figure 4: Participants' reported comfort with types of information being used to personalize ads to them. Those marked * were actually used in our personalized ads.

a noticeable reaction of both worry and disgust, and then I realized I had been visiting a lot of restaurant websites recently" (P123, *Online-Text-Personalized*). Only P135 (*Online-Text-Personalized*), however, hypothesized that the information flow was related to their participation in online surveys:

I became quite fearful that it was due to the surveys I had been filling out on Prolific and that a "researcher" broke some sort of policy and sold data to some food reservation company...

Many participants reported relief when the debriefing revealed the personalization was study-related. Through direct messages on Prolific Academic, we also contacted the < 30 participants who did not complete the study, whom we debriefed. After the debriefing, we asked participants how comfortable or uncomfortable they would be with personalization elements used in our ads. By far, participants were the most uncomfortable with the use of a partner's name; 80.8% responded they would be uncomfortable or very uncomfortable. Participants also felt uncomfortable or very uncomfortable with the use of their relationship status (56.5%) and first name (55.4%) for ad personalization. Participants were more comfortable with city and favorite cuisine, perhaps because these are much more commonly used in advertising. Figure 4 summarizes these results alongside other potential personalization elements.

During in-person interviews, we had the opportunity to ask participants to elaborate on their perceptions of comfort. Participants were more likely to feel comfortable about the use of an element if it was relatively coarse. "If it was [a food] specific to me, [I'd be] uncomfortable... [but] a lot of people like Mexican food" (P6, *Lab-Text-Personalized*). This logic aligns with participants' comfort with age and gender, as these categories are often grouped in large buckets. This could also explain the discrepancy between comfort with "city or town" and "location of cell phone." The former evoked a coarse-grained category; the latter implied fine-grained tracking.

Participants' varying levels of comfort were also informed by their mental models of how the information could be gathered. For some, information that was relatively easy to figure out made them more comfortable with its use. Regarding first name, P16 (*Lab-Text-Generic*) explained, "when you buy stuff online, they get your name and shipping address." However,

there was not a consensus among participants about how easily their first name could be obtained. Nevertheless, the majority of participants felt uncomfortable with personalization based on relationship status and partner's name because they felt it would be difficult for an advertiser to find. P3 (*Lab-Text-Generic*) went so far as to say about the use of partner's first name: "my social relationships? That's pretty fascist information to have. All for the greater good of what, selling bullshit?"

Information Disclosure

To measure the impact of hyper-personalization, we asked 43 questions of varying levels of potential invasiveness, as measured in our pre-study. We recorded how many they answered, how many they reported answering accurately, and how long they spent doing so. We hypothesized these values would vary by condition. That is, we expected that seeing a personalized ad would lead participants to answer fewer questions. However, this hypothesis was not supported by the evidence.

Overall, participants answered a mean of 37.1 of the 43 potentially invasive questions (86.3%). Figure 5 shows the breakdown by condition. Despite their negative reactions, participants who received personalized ads did not differ significantly from those who received generic ads in how many questions they answered. A Mann-Whitney U test revealed no significant differences between *Online-Banner-Personalized* and *Online-Banner-Generic* in the number of questions answered ($U = 878.5$, $p = 0.83$) or answered accurately ($U = 864$, $p = 0.73$), nor between *Online-Text-Personalized* and *Online-Text-Generic* ($U = 770$, $p = 0.17$) in the number of questions answered. We observed a marginally significant difference between *Online-Text-Personalized* and *Online-Text-Generic* ($U = 829$, $p = 0.08$) in the number of accurate answers.

Following the debriefing, we asked participants to review the accuracy of their answers to the invasive questions. We reiterated that there was no penalty for reporting inaccurate answers and that doing so was important for the integrity of the research. Nonetheless, participants indicated they had given an accurate answer for a mean of 36.6 (83.7%) of the potentially invasive questions, just shy of the 86.3% of questions they answered at all, as shown in Table 3.

Reasons for Disclosure

Overall, participants answered more invasive questions than we expected. Following the debriefing, we asked a series of questions to better understand decisions to answer the invasive questions. Participants gave many reasons for deciding to answer questions, versus selecting "prefer not to say."

Because money is a primary motivation for many crowdworkers when selecting and completing tasks [26], we made sure to clearly explain to participants that selecting "prefer not to say" would not affect their compensation. When asked post-debrief about their agreement that "I felt I would be paid the same regardless of selecting prefer not to say," 71.1% of participants chose "strongly agree" and 21.3% chose "somewhat agree."

Next, we directly asked participants in a free-response field how they decided whether to select "prefer not to say" for the invasive questions. 41.5% of participants mentioned selecting

Figure 5: Box plots showing how many questions participants answered and reported answering accurately by condition.

“prefer not to say” for questions they perceived as too personal or embarrassing. Other common reasons for selecting “prefer not to say” were for questions perceived as involving personally identifying information (27.0%) or information that could be used to compromise an account, such as answers to security questions (9.6%).

Participants also disclosed information — their first name and phone number — in Part 1. When we asked how they decided to do so, participants reported they did not think there were any risks (19.2%) or that they did not think or care about them (12.3%). Surprisingly, 12.3% of participants did not remember giving us this information at all. 12.9% of participants reported comfort in sharing their name and phone number because they trusted research studies generally, and 17.5% mentioned giving their information because they trusted Prolific (and, by proxy, any study hosted on the site). 9.6% of participants reported giving their information because they wanted compensation, underscoring the financial pressures on crowdworkers.

Prior work has suggested that when evaluating privacy risks, crowdworkers evaluate requesters’ characteristics, perceiving academic requesters as more legitimate than others [50]. We asked participants to explain how trustworthy they judged IIDR, a non-academic requester, to be. One-fifth of participants reported that they did not know or had not thought about it. Some participants expressed strong beliefs about IIDR’s trustworthiness (16.3%) or untrustworthiness (14.6%). Overall, though, 47.8% participants rated IIDR as just somewhat trustworthy, and 19.7% reported becoming more concerned about IIDR’s trustworthiness as they answered more questions.

When asked, only six participants reported looking up external information about IIDR. Crowdworkers often attempt to complete studies as efficiently as possible to maximize their earnings [26]. Consistent with this, 11.8% of participants cited not having time to do research on who conducts the studies they complete, or that it would be faster or simpler to lie or select a “prefer not to say” option. Interestingly, however, these same participants generally did not use the “prefer not to say” option. Some explicitly mentioned their reliance on crowdwork: “I am doing these to supplement my non income at the moment while I wait and hope for my disability to go through... I mean, at this point, why check? I am still going

to do the studies” (P70, *Online-Text-Generic*). A majority of participants reported that either they did not feel it was necessary to look up crowdwork requesters (31.5%) or that they did not think about doing so at all (21.9%). Overall, 15.2% mentioned trusting Prolific as a platform, with one explaining, “I assume prolific [v]ets its studies” (P121, *Online-Banner-Personalized*). To our knowledge, Prolific does not. Another participant wrote, “Quite honestly had it been a survey outside of a platform such as Prolific or Amazon Turk I would have absolutely researched it” (P207, *Online-Text-Generic*).

When asked explicitly, even more participants reported that trusting Prolific impacted their trust of our fictional institution; 82.5% of participants either agreed or strongly agreed with the statement, “I trusted the Institute for Interests and Demographic Research’s study because I found it using Prolific.” We also hypothesized that our survey’s use of the expensive Qualtrics survey software could impact participant trust, but found only 27.2% of participants agreed or strongly agreed that they trusted IIDR because the survey used Qualtrics.

Perceptions of Disclosure Risks

To further investigate factors impacting perceptions of privacy and risk when taking the survey, we asked participants how they thought IIDR would use the data collected in the study. Among participants, 28.1% wrote that they did not know or consider for what purpose IIDR would use the data. Some dismissed the question as irrelevant. For example, P95 (*Online-Banner-Personalized*) wrote, “I feel like it is more my job to respond to inquiries and to not think about the reasoning behind the inquiries.” However, of participants who did report having some expectations, 38.8% reported believing the data was for research purposes generally, with 17.4% reporting that they expected the data would be securely stored, properly anonymized, or not commercially sold. Only 11.2% reported believing that the data would end up being used for advertising or marketing purposes, while 8.4% (correctly) surmised that the data would not be used at all and was actually a test of what they would answer. Finally, 3.9% reported worrying that the data would be used for spam, identity theft, or other harms.

Prior research has shown that some participants volunteer for research to improve scientific knowledge [6]. However, the reality of information disclosure for research can be more complex. Facebook’s Cambridge Analytica scandal centered on a research study conducted on Mechanical Turk [23], narrowing the gap between research and ad personalization. We asked our participants about their familiarity with the Cambridge Analytica scandal regarding Facebook data and its involvement of crowdworkers. We hypothesized that participants’ knowledge of crowdworkers’ role in the scandal could affect their information disclosure habits on surveys. While 64.4% of participants reported at least a little familiarity with the scandal, only 2 participants (1.1%) said they thought crowdworkers were definitely involved in the scandal, with the majority of participants (72.4%) answering that they didn’t know whether or not crowdworkers were involved. After these questions, we summarized the Cambridge Analytica scandal and the involvement of crowdworkers. In their final comments, several participants thanked us for this explanation.

Despite limited knowledge of crowdworker involvement in the Cambridge Analytica scandal, most participants acknowledged that there were risks to disclosing information in research studies online. Most participants (89.8%) agreed there were some risks or said it depended on the context. Of those who mentioned specific risks, 13.0% mentioned their data being shared or sold, 11.3% mentioned data breaches or hacking, and 9.0% mentioned being personally identified. Only 9.0% reported feeling there were no risks. Those who did not report risks mentioned trust (“legitimate survey sites use bullet proof confidentiality,” P49, *Online-Banner-Generic*) or self-regulation of disclosure (“there is very little risk because I won’t share anything too personal,” P215, *Online-Banner-Generic*). Among participants, 16.4% said that there were fewer or no risks on Prolific compared to other survey sites. At the end of the study, some participants chided themselves. For example, P5 (*Online-Banner-Personalized*) wrote, “Now I do [think there are risks], [but] because of an inert trust with Prolific I acted dumb and revealed a bit too much.”

A few participants mentioned their economic status, defending their participation despite their awareness of risks. P153 (*Online-Text-Personalized*) said, “I think there is a lot of risk, and if I wasn’t flat broke I wouldn’t be doing this.” P207 (*Online-Text-Generic*) directly compared information disclosed in crowdwork to information disclosed elsewhere:

It’s impossible to really know where the data goes after you leave it. This is VERY unrestricted territory but the fact I get compensated? Makes me feel justified in taking any risks. google takes all day.. so does amazon. so do so many others. i dont get paid for that.

Changes in Mood During the Protocol

Participants responded to the PANAS questionnaire at three points. Both positive and negative scores range from 10 to 50. Larger numbers represent stronger affect. Participants had a mean positive score of 25.7 and a mean negative score of 16.0 at the first measurement. Answering the potentially invasive questions decreased participants’ mood. Positive affect decreased 2.4 points on average ($p < 0.001$), while the negative score increased 1.4 points ($p = 0.074$). The condition, and thus the personalization, was not a significant factor.

DISCUSSION

Our deception study investigated hyper-personalized advertising’s impact on information disclosure and risk perception. We hypothesized participants would be surprised and alarmed by our hyper-personalized ads. This was the case. Half of participants who saw the personalized ad reported strong, negative emotions in response to the inclusion of personal information. In contrast, no participants who saw the generic ad reported such strong, negative emotions. We also expected participants who saw a hyper-personalized ad to be less willing to answer the potentially invasive questions than their generic counterparts. To our surprise, this was not the case.

Risk Perception and Crowdwork

That participants continued to disclose personal information while feeling shocked, angered, or scared in the wake of their information being misused could be seen as an example of

Table 3: The percentage of participants who *answered*, and self-reported *accurately* answering, the 43 invasive questions. These percentages only include participants who saw the ad and did not immediately suspect it was study-related.

Question	Answered	Accurate
How many years of experience do you have with using the Internet?	100.0	100.0
What is the highest degree or level of school that you have completed?	100.0	100.0
What brand of phone do you have?	100.0	100.0
Does your residence have both hot and cold running water?	100.0	100.0
Have you ever traveled out of the country?	100.0	99.4
How do you feel about allowing refugees into the United States?	100.0	99.4
How many hours of sleep you get on the average night?	100.0	99.4
Do you have any children?	99.4	99.4
Do you have a currently valid driver’s license?	99.4	99.4
Did you receive an allowance as a child?	98.9	98.9
Which racial categories best describe you?	98.9	98.9
Which of the following qualities is most important to you in looking for a romantic partner?	98.9	98.3
Who did you vote for in the 2016 presidential election?	98.3	98.3
What is your zodiac sign?	98.3	98.3
Do you or anyone in your household own a car?	98.3	98.3
Did you vote in the 2018 midterm election?	97.7	97.7
Are you registered to vote at your current place of residence?	97.2	97.2
What was the last big purchase you made or considered making?	97.2	97.2
What did you eat for dinner last night?	96.6	96.0
How likely is it that you will fall behind in paying your housing costs during the next 6 months?	96.6	96.0
Do you have a disability?	96.0	96.0
Have you ever purchased alcohol for someone under 21 years of age?	96.0	95.5
Have you ever used marijuana?	95.5	95.5
In what religion were you raised, if any?	95.5	94.4
What color is the underwear you are wearing right now?	93.2	93.2
Are you sexually active?	93.2	93.2
Have you ever stolen from a person or store?	93.2	92.7
Do you believe in God or any deities?	92.1	91.0
What is your annual salary?	91.5	91.0
How much do you weigh?	89.3	88.7
For what reason did you last cry?	86.4	86.4
If applicable, at what age did you lose your virginity?	86.4	85.9
What was the last ailment that took you to the doctor’s office or ER?	85.9	85.3
How frequently do you view pornography?	85.3	84.7
How frequently do you masturbate?	81.9	80.2
In what city or town were you born?	78.0	77.4
When is your birthday?	75.7	71.2
When is your mother’s birthday?	63.8	57.1
When is your father’s birthday?	42.9	38.4
What is the first and last name of your closest friend?	40.7	34.5
What is your mother’s maiden name?	25.4	23.7
What is your work address?	22.6	19.8
What is your home address?	14.7	14.1

the privacy paradox, a well-studied phenomenon in which individuals report highly valuing their privacy, yet fail to take privacy-protective actions [18]. A common explanation for this paradox is that people engage in privacy calculus, weighing privacy risks against perceived benefits.

Prior work on crowdworkers has shown that their privacy calculus reflects the unique power dynamics and economic concerns of crowdwork [50]. Indeed, various aspects of participants’ perceptions of risk and decisions to disclose information were specific to the conditions of crowdwork. Participants reported concern, but did not engage in privacy-protecting behaviors for economic reasons, such as wanting to earn the study compensation and avoiding rejections on their account. However, their assessments of risk were also shaped by trust in the Prolific platform and research studies more broadly. Faced with ambiguity about our identity as a requester (the fictitious IIDR) and with no information about how the data we collected would be used, many participants adopted a trusting optimism. They trusted that Prolific vetted requesters and that information they disclosed for a study would be anonymized.

Unfortunately, Prolific does not vet requesters, and anonymization is nearly impossible to guarantee. For example, combining basic demographic information like birth date, gender,

and ZIP code can uniquely identify 87% of the U.S. population [54]. Further, anonymization is not sufficient for privacy protection. The privacy harms at play in our study were not a result of de-anonymization, but rather the unexpected reuse of data collected in another context. For academic researchers, institutional review boards (IRBs) oversee the responsible use of data, requiring researchers to clearly state their motivations and management practices for data collections. IRBs, however, do not directly supervise non-academic researchers or academics at institutions outside the purview of IRBs. Crowdwork platforms are a partially unsupervised ground where unverified requesters can collect data for undisclosed purposes.

Crowdworkers' trust of research combined with a lack of protection puts them at risk for significant privacy harm, motivating the need for better privacy by design on crowdwork platforms. Our results suggest that authenticating requesters is an important step. We join Sannon et al. [50] in recommending that crowdwork platforms communicate verified requester identities to workers. Additionally, IRBs could require studies to outline mitigations for concerns unique to crowdworkers and facilitate direct communication between crowdworkers and IRBs on the platform itself. Authenticated requesters and IRB oversight could improve transparency and accountability. Xia et al. [66] suggest that platforms require requesters to specify the purpose for, and types of, data requested in a task before asking a worker to complete it. Still, such transparency requires requester verification for privacy benefits.

Information Disclosure and Data Use

Beyond the crowdworking context, participants' highly negative reactions to the hyper-personalized ad are consistent with Nissenbaum's theory of contextual integrity (CI), which states that inappropriate data flows cause privacy concern [42]. Our deception manipulated CI's transmission principle for the data collected in Part 1. Instead of being used for demographics research (the apparent purpose), it was used to personalize an ad. We had hypothesized that participants might experience a cautionary lesson effect, in which exposure to a violation of CI would cause them to be more cautious sharing in the future. Participants' willingness to disclose information despite their discomfort suggests that a single recent violation of contextual integrity is not enough to change individuals' disclosure of personal information. Future work could explore whether multiple violations have a different effect. Future work should also repeat the protocol with non-crowdworkers.

Participants potentially compartmentalizing the privacy violation might further explain the absence of significant differences in disclosure across conditions. Many participants did not attribute the source of the hyper-personalized ad's information to our Part 1 survey, so they may have considered the ad's privacy invasion unrelated to the Part 2 survey. As a result, they may not have considered adopting privacy-protecting behaviors in the survey context. Future work should better unpack the potential for privacy compartmentalization. Future work could also seek to explore the lack of differences in disclosure. Two effects could have cancelled each other out: a cautionary lesson effect and a demoralization effect in which exposure caused users to feel their data is already out of their control.

When our participants were exposed to hyper-personalized ads, we observed an unexpected dynamic: attribution errors. As our participants generally did not know the mechanism by which the ads were personalized, some were quick to blame Google, Facebook, and various apps. Participants' inability to identify the source of personalization presents problems for online commerce generally and suggests the existence of negative externalities in targeted advertising. This finding, in turn, raises two intriguing possibilities. First, large digital platforms might suffer a disproportionate share of the consumer frustration associated with targeted advertising that they do not directly cause. This potential may incentivize large companies to favor policies restricting the most invasive third-party behaviors. Second, recent legal efforts like GDPR and CCPA have aimed to help consumers better understand and control downstream transfers of personal information. These efforts might be strengthened if consumers were better able to attribute the provenance of data used in ways they find privacy-violating.

Finally, participants' disclosure decisions in our study also raise the issue of whether the means to protect their privacy are readily accessible. Participants are likely habituated to the fact that disclosing personal information is often an unavoidable precondition for using the modern Internet:

The internet and various apps require me to sign away some of my privacy in order to use them, and this includes personal information that I don't want advertisers to see. I don't want this personal, private information in a database. I don't want the government or major corporations to know so much about me. They know more about me than I do at this point. . .¹

In the current data economy, people are encouraged to give up vast amounts of personal information. Given the ubiquity of downstream transfers of personal information between nameless third parties, as a recent newspaper article about secret consumer scores emphasized [21], and the corresponding breaches of contextual integrity, the kind of depressed acceptance expressed by the above participant should not be surprising. Consumers' autonomy is sharply limited when they do not, and cannot, know why their data is being collected, by whom, and for what purpose. Though sometimes outraged, they are unable to attribute privacy violations to their initial source and are unsure of where to direct their anger. In light of this, we should be slow to interpret the disconnect between individuals' actions and beliefs — their privacy surrender — as the result of willing privacy calculus, with people freely choosing disclosure. Instead, we should recognize that the framework created by technology and advertising companies leaves upset consumers without a viable way of protecting themselves and productively expressing their disquiet.

ACKNOWLEDGMENTS

We thank Juliette Hainline, Emmi Russo, Chris Eidsmoe, Soyoun Eom, Wendy Li, Alex Mueller, Elizabeth Woolridge Grant, Maitreyi Nabar, and Michael Vetter. This material is based upon work supported by the National Science Foundation in collaboration with Amazon under Grant No. 1939728.

¹Quote from a participant in our first round of data collection.

REFERENCES

- [1] Alessandro Acquisti, Leslie K. John, and George Loewenstein. 2012. The Impact of Relative Standards on the Propensity to Disclose. *Journal of Marketing Research* (2012).
- [2] Idris Adjerid, Alessandro Acquisti, Laura Brandimarte, and George Loewenstein. 2013. Sleights of Privacy: Framing, Disclosures, and the Limits of Transparency. In *Proc. SOUPS*.
- [3] Adam L. Alter and Daniel M. Oppenheimer. 2009. Suppressing Secrecy Through Metacognitive Ease Cognitive Fluency Encourages Self-Disclosure. *Psychological Science* 20, 11 (October 2009), 1414–1420.
- [4] Muhammad Ahmad Bashir, Umar Farooq, Maryam Shahid, Muhammad Fareed Zaffar, and Christo Wilson. 2019. Quantity vs. Quality: Evaluating User Interest Profiles Using Ad Preference Managers. In *Proc. NDSS*.
- [5] Muhammad Ahmad Bashir and Christo Wilson. 2018. Diffusion of User Tracking Data in the Online Advertising Ecosystem. *PoPETS* 2018, 4 (2018), 85–103.
- [6] Tom Clark. 2010. On ‘Being Researched’: Why Do People Engage With Qualitative Research? *Qualitative Research* (2010).
- [7] Mary J. Culnan and Pamela K. Armstrong. 1999. Information Privacy Concerns, Procedural Fairness, and Impersonal Trust: An Empirical Investigation. *Organization Science* 10, 1 (1999).
- [8] Amit Datta, Michael Carl Tschantz, and Anupam Datta. 2015. Automated Experiments on Ad Privacy Settings. *PoPETS* 2015, 1 (2015), 92–112.
- [9] Deborah Davis, Assaf Soref, J. Guillermo Villalobos, and Mario Mikulincer. 2016. Priming States of Mind Can Affect Disclosure of Threatening Self-Information: Effects of Self-Affirmation, Mortality Salience, and Attachment Orientations. *Law and Human Behavior* 40, 4 (2016), 351–361.
- [10] AJ Dellinger. 2019. Americans Received A Record 5.7 Billion Robocalls In October. *Forbes*. (November 2019).
- [11] Tamara Dinev and Paul Hart. 2006. An Extended Privacy Calculus Model for E-Commerce Transactions. *Information Systems Research* 17, 1 (2006).
- [12] Claire Dolin, Ben Weinshel, Shawn Shan, Chang Min Hahn, Euirim Choi, Michelle L. Mazurek, and Blase Ur. 2018. Unpacking Perceptions of Data-Driven Inferences Underlying Online Targeting and Personalization. In *Proc. CHI*.
- [13] Charles Duhigg. 2012. How Companies Learn Your Secrets. *New York Times*. (February 2012). <https://www.nytimes.com/2012/02/19/magazine/shopping-habits.html>.
- [14] Steven Englehardt, Jeffrey Han, and Arvind Narayanan. 2018. I Never Signed Up For This! Privacy Implications of Email Tracking. *PoPETS* 2018, 1 (2018), 109–126.
- [15] Steven Englehardt and Arvind Narayanan. 2016. Online Tracking: A 1-Million-Site Measurement and Analysis. In *Proc. CCS*.
- [16] Motahhare Eslami, Sneha R. Krishna Kumaran, Christian Sandvig, and Karrie Karahalios. 2018. Communicating Algorithmic Process in Online Behavioral Advertising. In *Proc. CHI*.
- [17] Alisa Frik and Alexia Gaudeul. 2016. The Relation Between Privacy Protection and Risk Attitudes, with a New Experimental Method to Elicit the Explicit Monetary Value of Privacy. *CEGE Discussion Papers* 296 (November 2016).
- [18] Nina Gerber, Paul Gerber, and Melanie Volkamer. 2018. Explaining the Privacy Paradox: A Systematic Review of Literature Investigating Privacy Attitude and Behavior. *Computers & Security* 77 (2018), 226–261.
- [19] Avi Goldfarb and Catherine E. Tucker. 2011. Privacy Regulation and Online Advertising. *Management Science* 57, 1 (2011), 57–71.
- [20] Kashmir Hill. 2012. How Target Figured Out A Teen Girl Was Pregnant Before Her Father Did. *Forbes*. (February 2012). <https://www.forbes.com/sites/kashmirhill/2012/02/16/how-target-figured-out-a-teen-girl-was-pregnant-before-her-father-did>.
- [21] Kashmir Hill. 2019. I Got Access to My Secret Consumer Score. Now You Can Get Yours, Too. *New York Times*. (November 2019). <https://www.nytimes.com/2019/11/04/business/secret-consumer-score-access.html>.
- [22] Kai-Lung Hui, Bernard C.Y. Tan, and Chyan-Yee Goh. 2006. Online Information Disclosure: Motivators and Measurements. *ACM TOIT* 6, 4 (January 2006), 415–441.
- [23] Jim Isaak and Mina J Hanna. 2018. User Data Privacy: Facebook, Cambridge Analytica, and Privacy Protection. *Computer* 51, 8 (2018), 56–59.
- [24] Leslie K. John, Alessandro Acquisti, and George Loewenstein. 2011. Strangers on a Plane: Context-Dependent Willingness to Divulge Sensitive Information. *Journal of Consumer Research* 37, 5 (February 2011), 858–873.
- [25] Ruogu Kang, Stephanie Brown, Laura Dabbish, and Sara Kiesler. 2014. Privacy Attitudes of Mechanical Turk Workers and the U.S. Public. In *Proc. SOUPS*.
- [26] N. Kaufmann, Thimo Schulze, and Daniel Veit. 2011. More Than Fun and Money. Worker Motivation in Crowdsourcing—A Study on Mechanical Turk. In *Proc. AMCIS*.
- [27] Kyle Kercher. 1992. Assessing Subjective Well-Being in the Old-Old: The PANAS as a Measure of Orthogonal Dimensions of Positive and Negative Affect. *Research on Aging* 14, 2 (June 1992), 131–168.

- [28] Freya De Keyser, Nathalie Dens, and Patrick De Pelsmacker. 2015. Is This For Me? How Consumers Respond to Personalized Advertising on Social Network Sites. *Journal of Interactive Advertising* 15, 2 (2015), 1–11.
- [29] Balachander Krishnamurthy, Delfina Malandrino, and Craig E. Wills. 2007. Measuring Privacy Loss and the Impact of Privacy Protection in Web Browsing. In *Proc. SOUPS*.
- [30] Mathias Lecuyer, Riley Spahn, Yannis Spiliopoulos, Augustin Chaintreau, Roxana Geambasu, and Daniel Hsu. 2015. Sunlight: Fine-Grained Targeting Detection at Scale With Statistical Confidence. In *Proc. CCS*.
- [31] Pedro G. Leon, Blase Ur, Rebecca Balebako, Lorrie Faith Cranor, Richard Shay, and Yang Wang. 2012. Why Johnny Can't Opt Out: A Usability Evaluation of Tools to Limit Online Behavioral Advertising. In *Proc. CHI*.
- [32] 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*.
- [33] John Maindonald and John Braun. 2010. *Data Analysis and Graphics Using R – An Example-Based Approach* (3 ed.). Cambridge University Press. <https://maths-people.anu.edu.au/~johnm/r-book/2edn/xtras/mlm-lme.pdf>.
- [34] Miguel Malheiros, Charlene Jennett, Sneha Patel, Sacha Brostoff, and M. Angela Sasse. 2012. Too Close for Comfort: A Study of the Effectiveness and Acceptability of Rich-Media Personalized Advertising. In *Proc. CHI*.
- [35] Arunesh Mathur, Jessica Vitak, Arvind Narayanan, and Marshini Chetty. 2018. Characterizing the Use of Browser-Based Blocking Extensions to Prevent Online Tracking. In *Proc. SOUPS*.
- [36] Jonathan R. Mayer and John C. Mitchell. 2012. Third-Party Web Tracking: Policy and Technology. In *Proc. IEEE S&P*.
- [37] Georg Merzdovnik, Markus Huber, Damjan Buhov, Nick Nikiforakis, Sebastian Neuner, Martin Schmiedecker, and Edgar Weippl. 2017. Block Me If You Can: A Large-Scale Study of Tracker-Blocking Tools. In *Proc. EuroS&P*.
- [38] Rachel Metz. 2018. The Scientist Who Gave Cambridge Analytica Its Facebook Data Got Lousy Reviews Online. MIT Technology Review. (March 2018). <https://www.technologyreview.com/f/610598/the-scientist-who-gave-cambridge-analytica-its-facebook-data-got-lousy-reviews/>.
- [39] Stanley Milgram. 1965. Some Conditions of Obedience and Disobedience to Authority. *Human Relations* (1965).
- [40] Jason C. Miller. 2009. Regulating Robocalls: Are Automated Calls the Sound of, or a Threat to, Democracy. *Michigan Technology Law Review* 16, 1 (2009), 213–253.
- [41] Youngme Moon. 2000. Intimate Exchanges: Using Computers to Elicit Self-Disclosure from Consumers. *Journal of Consumer Research* 26, 4 (March 2000), 323–339.
- [42] Helen Nissenbaum. 2004. Privacy as Contextual Integrity. *Washington Law Review* 79, 1 (2 2004), 119–157.
- [43] Lukasz Olejnik, Minh-Dung Tran, and Claude Castelluccia. 2014. Selling Off Privacy at Auction. In *Proc. NDSS*.
- [44] Stefan Palan and Christian Schitter. 2018. Prolific.ac—A Subject Pool for Online Experiments. *Journal of Behavioral and Experimental Finance* 17 (December 2018), 22–27.
- [45] Angelisa C. Plane, Elissa M. Redmiles, Michelle L. Mazurek, and Michael Carl Tschantz. 2017. Exploring User Perceptions of Discrimination in Online Targeted Advertising. In *Proc. USENIX Security*.
- [46] Tyson Quick. 2017. The Advertising Personalization Classification System. Medium. (January 2017). <https://medium.com/@tysonquick/the-advertising-personalization-classification-system-4ea1dd794d19>.
- [47] Edith Ramirez, Julie Brill, Maureen K. Ohlhausen, Joshua D. Wright, and Terrell McSweeney. 2015. *Data Brokers A Call for Transparency and Accountability*. Technical Report. FTC.
- [48] Franziska Roesner, Tadayoshi Kohno, and David Wetherall. 2012. Detecting and Defending Against Third-Party Tracking on the Web. In *Proc. NSDI*.
- [49] Camille L. Ryan and Kurt Bauman. 2016. *Educational Attainment in the United States: 2015*. Technical Report. US Census.
- [50] Shruti Sannon and Dan Cosley. 2019. Privacy, Power, and Invisible Labor on Amazon Mechanical Turk. In *Proc. CHI*.
- [51] Tyler Shanahan, Trang P. Tran, and Erik C. Taylor. 2019. Getting to Know You: Social Media Personalization as a Means of Enhancing Brand Loyalty and Perceived Quality. *Journal of Retailing and Consumer Services* 47 (March 2019), 57–65.
- [52] Roseanna Sommers and Vanessa K. Bohns. 2019. The Voluntariness of Voluntary Consent: Consent Searches and the Psychology of Compliance. *Yale Law Journal* 128, 7 (2019).
- [53] Sarah Spiekermann, Jens Grossklags, and Bettina Berendt. 2001. E-privacy in 2nd Generation E-Commerce: Privacy Preferences versus Actual Behavior. In *Proc. EC*.

- [54] Latanya Sweeney. 2013. Discrimination in Online Ad Delivery. *ACM Queue* 11, 3 (April 2013).
- [55] The AAAA and ANA and BBB and DMA and IAB. 2009. *Self-Regulatory Principles for Online Behavioral Advertising*. Technical Report.
- [56] Roger Tourangeau and Ting Yan. 2007. Sensitive Questions in Surveys. *Psychological Bulletin* 133, 5 (2007), 859–883.
- [57] Sabine Trepte, Michael Scharkow, and Tobias Dienlin. 2019. The Privacy Calculus Contextualized: The Influence of Affordances. *Computers in Human Behavior* (2019).
- [58] Huahong Tu, Adam Doupe, Ziming Zhao, and Gail-Joon Ahn. 2016. SoK: Everyone Hates Robocalls: A Survey of Techniques against Telephone Spam. In *Proc. IEEE S&P*.
- [59] Catherine Tucker. 2014. Social Networks, Personalized Advertising and Privacy Controls. *Journal of Marketing Research* 51, 5 (October 2014), 546–562.
- [60] Joseph Turow, Jennifer King, Chris Jay Hoofnagle, Amy Bleakley, and Michael Hennessy. 2009. Americans Reject Tailored Advertising and Three Activities That Enable It. *SSRN* (September 2009).
- [61] 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*.
- [62] David Watson and Lee Anna Clark. 1999. *The PANAS-X: Manual for the Positive and Negative Affect Schedule - Expanded Form*.
- [63] David Watson, Lee Anna Clark, and Auke Tellegen. 1988. Development and Validation of Brief Measures of Positive and Negative Affect: The PANAS Scales. *Journal of Personality and Social Psychology* 54, 6 (June 1988), 1063–1070.
- [64] Ben Weinshel, Miranda Wei, Mainack Mondal, Euirim Choi, Shawn Shan, Claire Dolin, Michelle L. Mazurek, and Blase Ur. 2019. Oh, the Places You’ve Been! User Reactions to Longitudinal Transparency About Third-Party Web Tracking and Inferencing. In *Proc. CCS*.
- [65] Allison Woodruff, Vasyi Pihur, Sunny Consolvo, Lauren Schmidt, Laura Brandimarte, and Alessandro Acquisti. 2014. Would a Privacy Fundamentalist Sell Their DNA for \$1000... If Nothing Bad Happened as a Result? The Westin Categories, Behavioral Intentions, and Consequences. In *Proc. SOUPS*.
- [66] Huichuan Xia, Yang Wang, Yun Huang, and Anuj Shah. 2017. “Our Privacy Needs to Be Protected at All Costs”: Crowd Workers’ Privacy Experiences on Amazon Mechanical Turk. *PACM HCI* 1, Article 113 (Dec. 2017).
- [67] Yaxing Yao, Davide Lo Re, and Yang Wang. 2017. Folk Models of Online Behavioral Advertising. In *Proc. CSCW*.
- [68] Jay (Hyunjae) Yu and Brenda Cude. 2009. Hello, Mrs. Sarah Jones! We Recommend This Product! Consumers’ Perceptions About Personalized Advertising: Comparisons Across Advertisements Delivered via Three Different Types of Media. *International Journal of Consumer Studies* 33, 4 (June 2009), 503–514.