What Does It Mean to Be Creepy? Responses to Visualizations of Personal Browsing Activity, Online Tracking, and Targeted Ads

Nathan Reitinger
University of Maryland

Bruce Wen
University of Chicago

Michelle L. Mazurek
University of Maryland

Blase Ur
University of Chicago

ABSTRACT

Internet companies routinely follow users around the web, building profiles for ad targeting based on inferred attributes. Prior work has shown that these practices, generally, are creepy—but what does that mean? To help answer this question, we substantially revised an open-source browser extension built to observe a user's browsing behavior and present them with a tracker's perspective of that behavior. Our updated extension models possible interest inferences far more accurately, integrates data scraped from the user's Google ad dashboard, and summarizes ads the user was shown. Most critically, it introduces ten novel visualizations that show implications of the collected data, both the mundane (e.g., total number of ads you've been served) and the provocative (e.g., your interest in reproductive health, a potentially sensitive topic). We use our extension as a design probe in a week-long field study with 200 participants. We find that users do perceive online tracking as creepy—but that the meaning of creepiness is far from universal. Participants felt differently about creepiness even when their data generated similar visualizations, and even when responding to the most potentially provocative visualizations—in no case did more than 66% of participants agree that any one visualization was creepy.

KEYWORDS

Web Tracking, Transparency, Usable Privacy, User Study

1 INTRODUCTION

As people browse the web, third-party trackers follow them and build profiles of their demographic and psychographic attributes. This practice is known as Online Behavioral Advertising (*OBA*) [16, 85]. The goal is to serve consumers with more precisely targeted ads, often at the cost of privacy. For instance, labeling someone as "Parent of Infants (0-1 years)" after they search Google for "size two diapers" has become a normalized experience, but it can still feel invasive and creepy [32, 37, 56, 60, 138].

For years, researchers have explored how users perceive OBA. From *Smart, Useful, Scary, Creepy* [129] over a decade ago to the more recent *Whispering with Voice Assistants* [89], researchers have found consistent, but nuanced, discomfort with online tracking. Users find OBA "scary or creepy" due to privacy concerns, but also "smart" and "useful" due to the increased relevance of ads [129].

"Creepy" has become a common, almost default description for this nuanced discomfort, but the term remains ambiguous and difficult to parse takeaways from. Is tracking, a now familiar staple of the web, creepy? If tracking is creepy, are all types of tracking practices equally creepy? Are any tracking practices universally understood to be creepy? Can we quantitatively say that certain factors, like accuracy, increase or decrease feelings of creepiness? Does the sensitivity of data affect perceptions of creepiness?

To better define creepiness in the online tracking setting, we set out to answer the following research questions:

RQ-1: When tracking practices (from mundane to provocative) are made visible and intelligible to users, do users find tracking to be creepy? How does this affect their attitudes and behavioral intentions toward tracking?

RQ-2: What makes specific instances of tracking creepy or not creepy? Are there certain factors (e.g., accuracy) that increase or decrease feelings of creepiness?

To answer these questions, we created a design probe to explore which aspects of OBA users find creepy, and why. We started by conducting feature-testing interviews ($n = 13$) to see which types of tracking-focused visualization mock-ups participants might deem creepy. Based on the results, we selected ten of the creepiest prototypes to develop. We implemented visualizations (i.e., graphical representations of the information trackers may collect) based on the selected mock-ups, as part of a radical re-envisioning of an existing open-source browser extension created by Weinshel et al. [132]. These visualizations were not created to optimize design choices, but rather to invite participant reaction. We term the new browser extension Tracking Transparency v2 (*TT2*).

We made significant improvements to the original extension. We implemented a new web page classification model (stored locally) that is significantly more accurate and much more fine-grained in the labels it assigns. Additionally, we added new data sources by scraping and recording both the ads a user sees when browsing the web and the user's Google Ad Settings dashboard, which lists attributes (e.g., demographics and interests) Google has inferred for that user. We detail all of our improvements in Section 3. We have open-sourced the code for TT2.[1]

These new features allowed us to visualize tracking in far more fine-grained and potentially provocative ways. For instance, our extension highlights the most sensitive interests inferred about the user, visualizes the user's sleeping habits, details exactly how the user divides their time among interests, summarizes the ads the

Proceedings on Privacy Enhancing Technologies 2024(3), 715–743

https://doi.org/10.56553/popets-2024-0101

[1] Information on the project is available at https://osf.io/45kgc/ and the code is available at https://github.com/UChicagoSUPERgroup/TrackingTransparencyPETS2024.

user has been shown, and displays guesses of why the user might have been shown particular ads.

We then conducted a longitudinal field study ($n = 200$) in which participants downloaded the redesigned extension and used it for one week. Participants were randomly assigned either to use our fully featured new extension or more limited variants as controls. We conducted two surveys, one before participants used the extension and one after, to gauge changes in participants' attitudes and behavioral intentions toward tracking. In the second survey, we also asked participants to load the extension and answer per-visualization questions on specific factors—drawn from the literature about technology in general, not specific to tracking—that affect creepiness (i.e., Creepiness Factors).

In summary, we use our design probe to unpack nuances of creepiness as it relates to OBA. We report the following key findings:

- Decades after the advent of OBA, users (still) find the tracking necessitated by OBA to be creepy. More than 80% of participants agreed or strongly agreed that at least one visualization in our extension was creepy.
- Agreement on which tracking practices are creepy, as visualized in TT2, is far from universal: "more sensitive" does not mean "more creepy," and neither does data itself correlate with opinions on creepiness (e.g., later inferred bedtimes did not mean "more creepy").
- Participants found *more accurate* visualizations to be creepier than inaccurate ones, contradicting findings in prior work that *inaccurate* inferences were especially creepy or problematic [23, 30, 129, 131].
- Contextual violations, privacy invasions, and willingness to take action correlate with perceptions of creepiness.

2 BACKGROUND AND RELATED WORK

Here we discuss prior work on behavioral advertising, creepiness, and transparency-enhancing technologies.

2.1 Online Behavioral Advertising

From the first third-party tracker in 1996 [33, 68, 72] to Target predicting the due dates of pregnant women a decade ago [14, 31], being a consumer today means being watched at every turn [2, 27, 101]. One goal of this surveillance is OBA: collecting data about consumers and using that data to personalize ads.

Today, OBA works through ad exchanges [86]. When a user visits a web page, each ad the user sees is determined by a proprietary, real-time bidding mechanism involving *publishers* (i.e., the website being visited), *ad exchanges* (i.e., the company eventually serving the ad), *bidders* (i.e., the advertising agencies), and *advertisers* (i.e., the company that wants to make a sale). Advertisers want to serve ads to interested users and bidders facilitate this with cookie matching to single out users. Once singled out, the user's profile is checked for matching interests using demographic and psychographic attributes. The entire process happens in less than 100 milliseconds and is opaque to users.

Researchers have studied OBA for many years, often concluding that users are uncomfortable with online tracking. For example, researchers found that users generally reject OBA [128], find OBA invasive [74], and are "not okay" with OBA because of the tracking it requires [96]. On the other hand, researchers have also shown that users find OBA "smart" and "useful" when providing relevant content [30, 129] and prefer personalization over "vanilla" search results [88]. In short, OBA can be creepy, but in a nuanced way.

2.2 Creepiness

Researchers have made attempts at disambiguating technology creepiness, but not in the OBA setting and not with participant data. Factors of creepiness—what makes something creepy—have generally fallen into the following buckets: context, personal privacy invasions, accuracy, and the willingness to take action.

In *Theory of Creepy*, Tene and Polonetsky looked at the disconnect between technical capability and social values, finding that a technology is considered creepy when it violates an existing norm [124]. This finding has been supported by later studies looking at practices like whispering to voice assistants [89], unpredictable features [137], ambiguity in expected behavior [66, 73, 78], measuring creepiness [136], and a lack of transparency [122].

Perceived invasions of privacy, like feeling watched [87, 111, 113], have also been linked to creepiness [129]. In the OBA context, the concept, originated by Altman [7], is most similar to re-targeting, the practice of showing users ads for items they have previously viewed [11]. Trackers learning potentially sensitive attributes [81], like inferring your interest in a particular medication based on searches for that medication, also fit into this category.

Accuracy has been shown to play an important, but mixed, role for creepiness. Researchers who have measured the accuracy of trackers inferring interests using real participant data have found, somewhat surprisingly, that more accuracy leads to more comfort [23]. For example, when showing participants their own Twitter ad inferences, Wei et al. found that participants who saw accurate inferences were more likely to be comfortable with these inferences, find them fair, and want to see more of them [131]. Dolin et al. found similar results with Google Ad Words data [30]. However, researchers in [126] found the opposite result; participants were less comfortable with algorithmic recommendations when the recommendations were more accurate.

Willingness to take action has also been identified as an aspect of creepiness. Researchers have shown that there can be an apparent disconnect (sometimes called the "privacy paradox") between users perceiving a technology as invasive, but not taking privacy-protective actions [24, 62, 84, 92, 119]. One line of work suggests that willingness to take a privacy-protective action, like sharing less data, is a poor proxy for privacy concern [12, 49, 118]. Other researchers have looked more generally at how feelings impact perceived privacy risk, privacy choices, levels of concern, and coping strategies [22, 25, 35, 58, 61, 69, 120]. This work demonstrates that some emotions are more closely tied to a willingness to take privacy-protective actions than others [26]. Specific to OBA, researchers have found that users may be uncomfortable with certain digital services, but nonetheless continue using them [111]. Normalized discomfort, perceived helplessness, and lack of self-efficacy are common explanations [9, 21, 113]. In our work, willingness to act is most applicable in terms of finding a correlation between identifying a visualization as creepy and measuring the participant's willingness to act based on the feeling of creepiness.

Table	Key	Value (example)	S	A
❶ Pages	protocol	https	Yes	No
	hostname	www.nytimes.com	Yes	Yes
	domain	nytimes.com	Yes	Yes
	path	new-york-snow-rain-winter.html	Yes	Yes
	pageId (primary key, timestamp)	1674525348324 (page visit time)	Yes	No
	time spent on page	[type: start, value: 1661375368994, time on page: 0...]	Yes	No
	title	"Why Hasn't it Snowed Yet in New York City?"	Yes	Yes
	life activities (Google search)	"work from home" label applied from query "work from home if snow NY"	Yes	No
❷ Trackers	trackers	[Ensighten, Google, Quantcast ...]	Yes	No
❸ Inferences	inference	Weather	Yes	No
	full inference path	News ⟶ Weather	Yes	No
	word cloud (if sensitive)	Example website interest—weather—is not sensitive, no entry	Yes	No
❹ Ads	DOM	`<head><meta charset=UTF-8><script>var jscVe...`	No	–
	domain	"nytimes.com"	Yes	Yes
	inference	Footwear	Yes	No
	full inference path	Shopping ⟶ Apparel ⟶ Footwear	Yes	No
	explanation	["Your visit to the advertiser's website or app",...]	Yes	No
	URL: provided ad explanation	`https://adssettings.google.com/whythisad...`	Yes	No
	URL: ad hyperlink	`https://adclick.g.doubleclick.net...`	Yes	No
	URL: inferred final destination	`https://thursdayboots.com...`	Yes	No
❺ Ads Settings	inferences	[type: demographic, value: 18-24 years old...]	Yes	No
	pageId	1660429179601 (Ads Settings fetch time)	Yes	No
❻ IP Address	current IP	{IP: 2601:...:9535:d672, city: [redacted], state: [redacted] ... }	No	–

Figure 1: Example data extraction from a web page, used by TT2 for its visualizations (author-generated values here, tables stored in client-side storage). Some data is shared (S) with researchers, either anonymized (A) or not (always encrypted in transit, see Section 4.5).

2.3 Transparency Enhancing Technologies

Transparency enhancing technologies, which help visualize the hidden ways data is used, have increased in popularity in recent years due to the proliferation of laws like the GDPR in the European Union [34] and the CCPA/CPRA in California [19]. Research focusing on participant reactions to transparency typically uses either company-provided or researcher-created transparency tools. Results differ for each method.

Users' reactions to company-provided resources are often positive. Participants express feelings of control [8, 10] and knowledge [36, 109]. At the same time, company-provided tools may leave something to be desired. Data downloads or ad preference managers can be ambiguous, confusing, lack definitions or data important to users, and make it difficult to understand why certain interest or demographic inferences were made [13, 50, 97, 98, 130, 130, 131].

On the other hand, studies using researcher-provided tools have found that participants are generally, but not universally, uncomfortable with tracking. Wills and Zeljkovic queried a user's browser history and displayed trackers from actual websites visited [134]. A total of 63% of participants reported concern with tracking based on their own browsing history, but that number dropped to half when looking at tracker-inferred information. Likewise, a majority of Weinshel et al.'s participants agreed that tracking was creepy, yet were also *comfortable* with companies inferring their interests [132].

Our study helps explain and contextualize these nuanced results—creepiness is far from a universal concept, see Section 5.3—in part thanks to our novel approach of using both company-provided data from the Google Ads Settings dashboard and our own inferences based on user browsing. Our design enables a more contextualized perspective on the tracking ecosystem.

3 TT2

As a design probe, we developed TT2, which substantially revises Weinshel et al.'s browser extension [132]. In this section, we discuss how we updated the extension's interest inference engine, added new data sources to the extension, and developed ten new participant-defined "creepy" visualizations. An overview of the data TT2 collects for its visualizations, per web page, is found in Figure 1.

3.1 Improving Interest Inferences

Like Weinshel et al.'s extension, TT2 infers potential ad-interest categories from web pages visited by users to demonstrate the potential impacts of tracking. As with Weinshel et al.'s extension, our classifier runs locally to maximize participant privacy—in order to avoid requiring participants to share their browsing history with us or an external service's API. Assigning an ad-interest category based on text found on a web page is a difficult problem, as suggested by research showing that behavioral profiling is largely inaccurate [13, 98]. Our goal was to create a useful design probe, which does not require perfect accuracy.

To enable our new visualizations, TT2 significantly improves the interest inference engine used by Weinshel et al. [132]. Their extension simulated how companies infer users' interests by matching Wikipedia-classified keywords on Google AdWord topics with keywords found on web pages visited by participants. The implementation of their model also used post-processing to improve accuracy (~60% accurate) at the cost of inference granularity—truncating most classifications to only a handful of top-level categories.

We designed a new model that improves both accuracy (80% on top-level categories and 71.4% overall) and, more importantly, diversity by assigning labels from nearly 1,000 Google Cloud Natural Language Content Classification categories. We started by collecting a large amount of inference-topic training data using the Google Cloud Natural Language Content Classification API [44] as an oracle for labeling websites' topics. Labels from Google's API included up to three levels of depth: 27 top-level categories, 245 second-level categories, and 620 third-level categories [43] (see Figure 2 ❷ for an example). Given input text, the Google API returns a list of ad-interest categories and confidence scores. We used a supervised learning approach, taking output labels from the API as ground truth and training our own shadow model [114].

We created a corpus of web pages for training and testing. Using Tranco's [93] one million top domains,[2] we selected the top 100,000 domains and a random sample of 2% of the remaining 900,000 for a total of 118,000 domains. We added ten random sub-pages per domain, filtering to exclude auxiliary pages like privacy policies or contact pages. We also removed any pages not written predominantly in English using the `langdetect` library [28]. We scraped web page text using Selenium in early 2021 and extracted web page text using Mozilla's Readability tool [79]. We found from this initial dataset (see Figure 12 in Appendix E) that some ground-truth labels only had a few examples associated with them. We manually added targeted web pages for labels with fewer than 50 examples by keyword searching using the Google Search API [6]. We also limited our final corpus to labels with a $\geq$ 90% confidence score from the Google API. With these modifications, the final corpus contained examples for 96% of all possible labels in Google's content classification, excluding world localities.

We trained and compared several models using our corpus and the Google-API-provided labels (see Table 6 in Appendix D). A bag-of-words model with a single layer perceptron [104] was most effective in terms of accuracy and efficiency (200 predictions per second). We achieved a test accuracy of 71.4% (train-test split at 9:1). We considered this accuracy sufficient for measuring participant reactions to interest inferences. The model performed best for second-level or first-level categories (74.2% and 80.1%, respectively) and varied per-category. For additional analysis, see Appendix E.

3.2 New Data Sources: Google Inferences & Ads

TT2 bridges a gap in the literature (see Section 2.3) by integrating both its own estimates of inferences about users and data scraped from company-provided transparency dashboards, specifically, the Google Ads Settings dashboard [40]. The Google Ads Settings page[3] presents an unordered list of attributes (i.e., inferences) associated with that user. These attributes may be *interests* (e.g., "Arts→Entertainment→Movies"), *demographics* (e.g., "Homeowner"), *companies* (e.g., "Nordstrom"), *videos* (e.g., YouTube channels like "9 videos from BetterHelp"), or *locations* (e.g., Latin America). We fetch Ads Settings data automatically if a user is already logged in to a Google account and has personalized ads turned on. If not, we provide instructions on how to log in to import this data. To measure changes over time, we re-fetch the Ads Settings page on every fifth web page visit. We picked this number based on observed updates to the Ads Settings page, which occurred more frequently during browsing. We reported on results exclusive to Google Ads Settings data separately [102].

We also captured ads served to users and corresponding "why did I get this ad" explanations [51]. Because most ads are encased in `iframes`, we analyze all `iframes` found on a web page. If the `iframe` includes an outgoing hyperlink matching a known ad server [3, 4, 70, 125] (e.g., `ssp.yahoo.com`) then TT2 stores the entire `iframe`, logs the outgoing links triggering the capture, and looks for an ad explanation hyperlink (e.g., `google.com/whythisad`). We then generate interest labels for ads by fetching outgoing hyperlinks to the ad's final destination (i.e., where the user would go if they clicked the ad). To avoid click fraud [55, 65, 133]—the extension fetching a full link could charge advertisers as a higher-cost "click"—we do not fetch full hyperlinks directly, but instead use regular expressions to infer final destinations (i.e., identify the last-most URL in the full hyperlink by parsing on "`/http/g`" and then keeping, from that URL, only the domain to be visited). Appendix A gives a detailed example. We fetch these inferred links and use our inference classifier to categorize the resulting web page. Using this method is more ethical than fetching ad links directly, but limits us to analyzing a little less than $\leq$ 50% of the total ads a user is served.

3.3 Intentionally Provocative Visualizations

To explore user reactions to salient facets of online tracking, we developed visualizations that were intentionally more provocative when compared to Weinshel et al.'s visualizations. We brainstormed potentially creepy visualizations by conducting a literature review on "creepy web tracking" and created 23 mock-up visualizations from what we learned. We then conducted, after receiving Internal Revenue Board (IRB) approval, exploratory interviews to investigate which prototypes drew the strongest reactions. Participants were recruited via Prolific and were required to be at least 18 years old, located in the United States, able to use video conferencing software, and have a 95%+ approval rating [94]. We continued recruiting until reaching theoretical saturation (i.e., no longer hearing substantially new comments [106]). In total, we conducted 13 interviews.

Each interviewee viewed a subset of the 23 prototypes [53]. We used open-ended questions to elicit initial reactions (e.g., what are your general thoughts on this visualization) and asked follow-up questions on any provided feelings of comfort or discomfort. After viewing the prototypes, participants explained which visualizations they felt provoked the strongest reaction overall and which they were most surprised by. Throughout the interview, we refrained from using the word "creepy" to insulate against bias [45, 90] and encourage honest answers [63, 64]. To reduce demand effects [80] and social desirability [45], we told participants we had been hired by a third party to evaluate the prototypes and had not made them ourselves. Appendix B.1 contains the interview script and Appendix F gives examples of the prototypes we picked for development.

Based on these interviews, we selected mock-ups for development using a variety of factors. For one, we picked visualizations participants reacted most strongly to, taking note of the visualizations participants identified as creepy Statements indicating creepiness were found by reviewing our interview transcripts, like this comment from P-12: "Incognito mode, my best friend." Additionally, we considered the practical ability to develop each mock-up: the feasibility with existing in-browser visualization tools, the difficulty or ease of developing the visualization, and whether we felt that participants were likely to have enough of a certain type of data to make a visualization possible (e.g., if health data were uncommon, then Figure 20 would not appear). In the end, we selected ten visualization prototypes from these interviews to develop.

We next discuss the ten new visualizations we developed. Figure 2 provides an example of each. The extension's dashboard page

[2] Although all website popularity rankings contain biases [107, 108], we used the Tranco list to gather content for Google's content classification API, a use case less affected by these biases. We also applied a post-hoc process to gather less popular content categories, further reducing the impact of bias from website rank.

[3] Ads Settings [39, 40] has since been updated to My Ad Center [42]. All data collected for this study occurred when Ads Settings was in place.

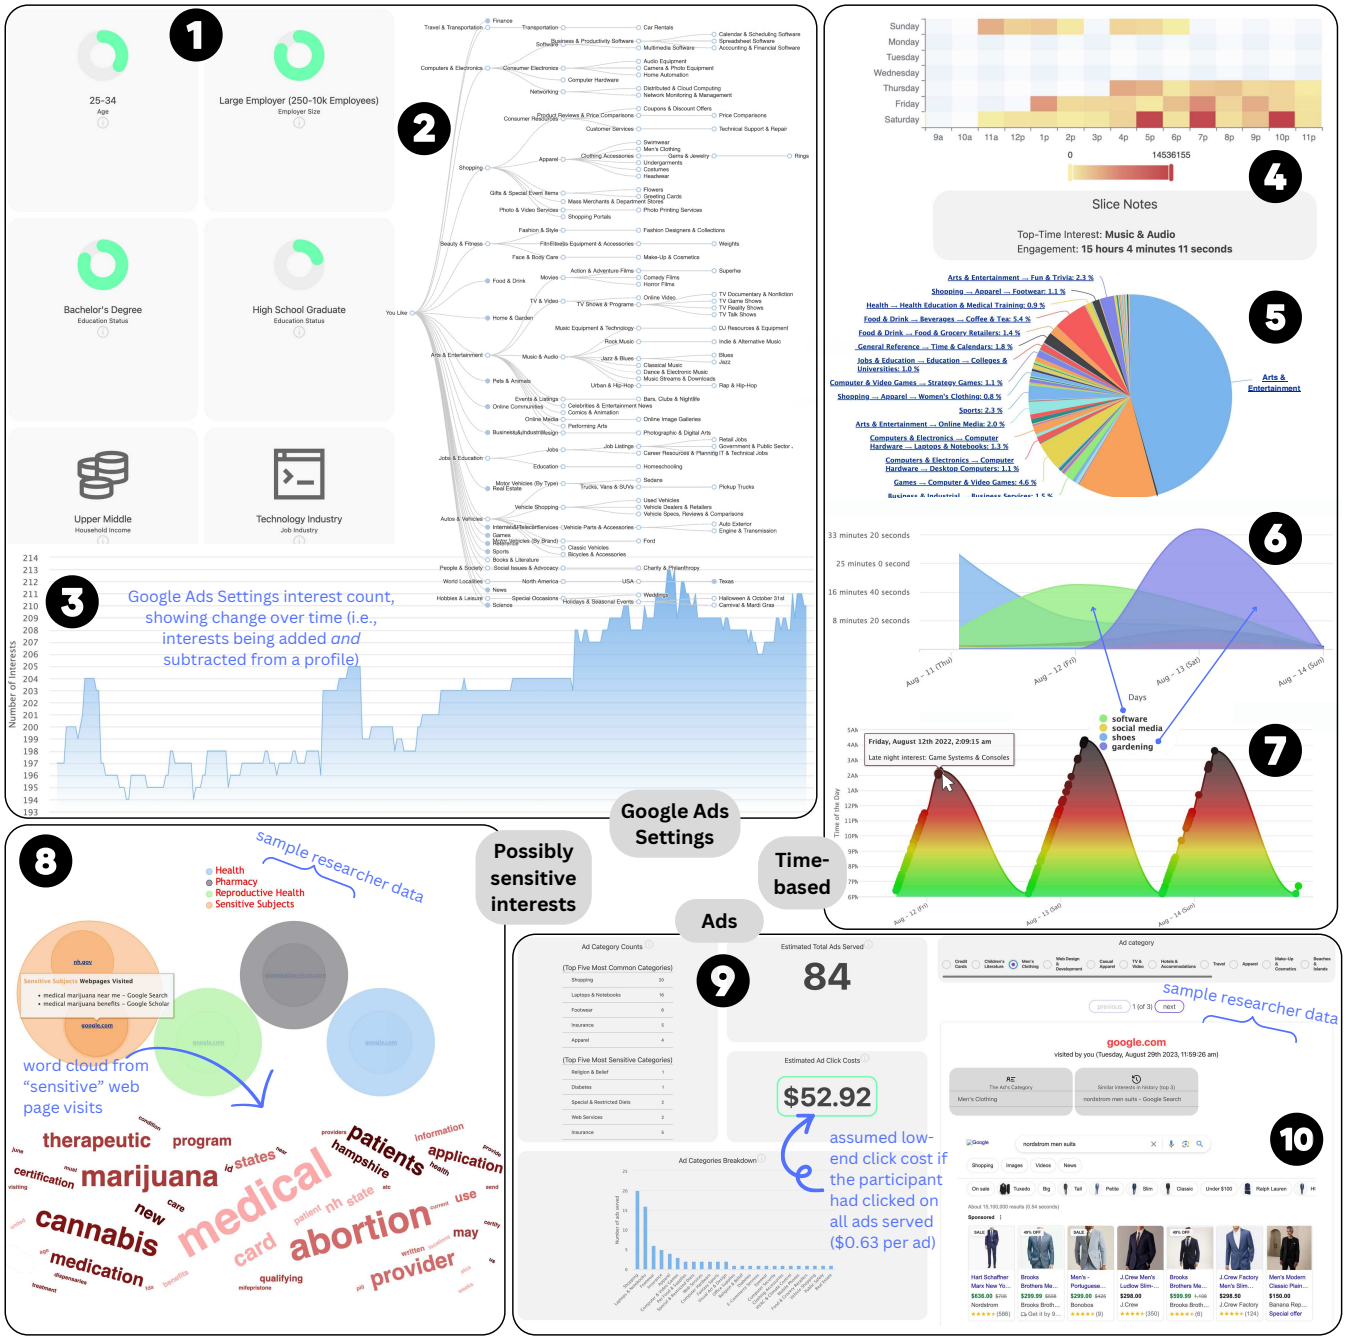

Figure 2: TT2 visualizations. Unless noted, examples come from anonymized participant data. Google Ads Settings visualizations: ❶ *Google Demographics* shows tiles of demographic data from Google Ads Settings; ❷ *Google Interests* shows all attributes associated with the participant (since extension install); and ❸ *Google Interests Dynamic* shows increase or decrease in Ads Settings attributes over time. Time-based visualizations: ❹ *When You Are Engaged* shows a heatmap of browsing activity; ❺ *Time Spent per Interest*, shows the total amount of time spent engaging with an interest (inferred from time spent on web pages); ❻ *Late-Night Engagement* shows late-night browsing activity; and ❼ *Search Habits* shows Google searches clustered into marketing keywords. Possibly sensitive interests visualizations: ❽ *Possible Sensitive Interests* shows potential sensitive interests, domains associated with these interests, and a word cloud from web page text associated with the interests. Ads visualizations: ❾ *Ads Served Overview* shows an overview of the ads served to the participant; and ❿ *Ad Explanations* shows inferred explanations for why the participant received the ad.

opened with an explainer providing details on the tracking ecosystem. Scrolling down the page would reveal a table of contents with hyperlinked sections to individual visualizations. From any page in the extension, the user could have also clicked on the left-most red button which would send them to a "take action" page providing examples of tracker-blocker technologies.

Google Ads Settings visualizations: Three of the new visualizations portray Google Ads Settings data. The underlying tracking mechanism here concerns Google's cookies, found on many web pages participants visit [102]. To display this data, users must be logged into Google and have "personalized ads" turned on. These visualizations focus on demographics, interests, and total attribute counts over time. In the *Google Demographics* visualization ❶, we show tiles related to a participant's Google-inferred demographics like age, gender, income, and marital status. In *Google Interests* ❷, all interests associated with the user since installing the extension are displayed in a tree structure (levels indicate finer-grained targeting). This visualization includes typical interests like "hockey," but also interest types not as widely discussed in the literature, including videos (e.g., "a video from T.J.Maxx"), companies (e.g., USAA, Nordstrom), and locations (e.g., "Grand Rapids-West Michigan"). Lastly, in *Google Interests Dynamic* ❸, we show how attributes are updated by Google as users browse the web: the total count of attributes associated with the user increases or decreases over time.

Time-based visualizations: Visualizations here focus on time, and involve trackers logging when and for how long web pages are visited by users (i.e., focus events [76]). In *When You Are Engaged* ❹, we display a heatmap of weekly, per-hour engagement, with engagement measured by time spent per web page (i.e., using `focus events` to assess time per page [76]). High periods of engagement are noted with darker colors. We also use a pie chart to display the participant's most common interests in terms of time spent per interest (*Time Spent per Interest* ❺). Participants are able to click on an interest and view a bar chart of aggregated time spent per domain. In *Late-Night Engagement* ❻, we highlight when a user is *not* engaged, inferring when the user has gone to bed and what their late-night browsing habits are (tooltips reveal late-night interests). We consider "late-night" to be web page visits occurring in the 6 PM to 4 AM range. In *Search Habits* ❼, we group similar Google searches together by time, focusing on periods of heavy engagement and inferring life events like job-seeking or marriage. *Search Habits* identifies life activities by matching Google search queries against a list of Mondovo keywords (e.g., the life-event "wedding" is associated with Google search keywords "wedding songs" and "wedding cakes") [77].

Possibly sensitive interests visualization: In this visualization ❽, we attempt to highlight the potentially "sensitive" websites participants visit. We define sensitivity based on data from related work [30] categorizing ad interests by comfort level, as rated in a user study. The underlying tracking behavior for this visualization comes from categorizing interests by sensitivity, as, for example, Google's Ads Settings does on topics like pregnancy, alcohol, or weight loss [102]. The visualization provides a list of sensitive interests followed by a bubble chart mapping domains to sensitive categories. Tooltips on each domain highlight the different trackers

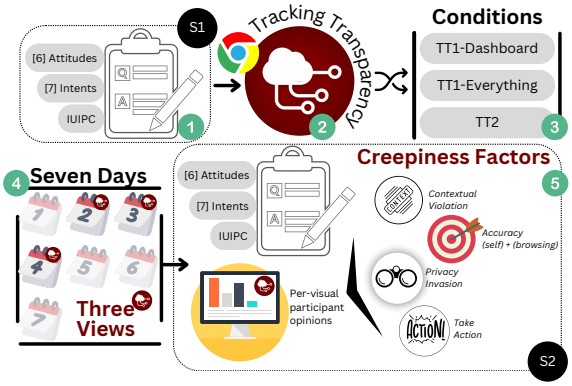

Figure 3: Participants began by answering questions in *S1* ⓐ, downloading the extension ⓑ, and were then put into one of three separate conditions ⓒ (assigned randomly, balanced per condition). After using the extension for seven days, with at least three days of extension dashboard views ⓓ, participants were invited to *S2*, which included reactions to the participant's own visualizations as outlined by the Creepiness Factors ⓔ.

on those domains, and clicking on a bubble reveals a domain-specific word cloud (i.e., a TF-IDF list of words found on web pages under this domain [5]). A word cloud on text from all sensitive web pages is shown below the bubbles.

Ads visualizations: The last two visualizations concern ads, and involve tracking practices related to ad exchanges [86]. The first, *Ads Served Overview* ❾, provides an overview of ads served, including the total number of ads served, most-served and most-sensitive ad interest categories, and an estimated click cost (i.e., if the user were to click on all ads served, how much would these clicks cost an advertiser, assuming a low-end click cost of $0.63 per click [18, 54, 112]). The second, *Ad Explanations* ❿, replays captured ads paired with "why this ad" information—taken both from ad-provided sources and from matching ad interest information with our own browser-history inferences (e.g., you may have seen this ad because Google thinks you are interested in "Home→Garden→Home Furnishings," you've visited three other websites about "Home Furnishings," and this ad is about "Home Furnishings").

4 FIELD STUDY METHOD

We conducted a two-part, IRB-approved field study to investigate participant reactions to TT2. An overview is provided in Figure 3. Participants were recruited through Prolific, were 18 years of age or older, located in the United States, proficient in English, and had a 95% or higher approval rating. We asked that participants use Google Chrome as their primary web browser, refrain from using "private mode" (the extension is disabled in private mode), and disable any ad blockers while enrolled in the study. These unenforced requests were designed to provide participants with a more realistic perspective of the tracking ecosystem they would be commenting on. Participants were instructed to view the extension's dashboard page on three separate days over the duration of the study (seven days). The extension provided a "viewing count" in its pop-up to inform participants of their progress. This requirement was enforced by logging a new view if it had been at least five hours

Table 1: The three conditions used in the study. Conditions differ by which visualizations are visible to participants, although the underlying architecture of the extension (e.g., the interest inference engine) is the same across all conditions.

TT1-Dashboard (See Figure 9 in Appendix D)	A short explanation of "trackers" and "interests," followed by snapshots of top five interests, trackers, and recent sites, and aggregate statistics on the total trackers encountered, pages visited, and potential interests (dashboard-only version of "Longitudinal:Interests" [132]).
TT1-Everything (See Figure 10 in Appendix D)	All visualizations from TT1-Dashboard plus additional tabs to include visualizations from TT1: trackers, interests, activities, and sites (mirroring "Longitudinal:Interests" [132]).
TT2 (See Figure 2)	Revised explanation of "trackers" and "interests," followed by ten new "creepy" visualizations.

since the previously-counted view (i.e., we did not strictly enforce viewing "days" because we wanted to encourage participants to view the dashboard naturally, and felt the five-hour limit would strike a balance). Data was collected from July to September 2022.

4.1 Study Conditions

Participants were assigned to one of three conditions with different visualizations, all using the updated interest inference engine (Table 1). The main TT2 condition included our ten new visualizations and none of the older visualizations. For comparisons, we also included TT1-Everything, which mirrored the full-featured, original extension from Weinshel et al.'s original study [132], and TT1-Dashboard, a stripped-down version of TT1-Everything designed to display a minimal amount of aggregate data about tracking. These conditions were designed to enable further comparison among different visualization types with different levels of detail.

4.2 Pre-Usage Survey: *S1*

In the first survey, ***S1***, participants answered general questions about their awareness of tracking and familiarity with transparency tools like Google Ads Settings. Participants were given an invite-only link to the extension on the Chrome Web Store and asked to install the extension. Once the installation was complete, participants could finish the survey. We opted to have participants download the extension prior to finishing the survey to allow those who did not wish to install the extension to exit early. The recruitment text warned participants about the requirement to download software. The full text of ***S1*** may be found in Appendix B.2.

Participants then answered six questions regarding their attitudes toward tracking and seven questions regarding their intent to use privacy-protective tools. These questions were reused from Weinshel et al.'s study [132]. We included the eight-question IUIPC [46] and concluded ***S1*** with demographic questions and a reminder about the requirement to view the extension's dashboard page on at least three separate days. To remind participants to check the dashboard and reduce dropout between ***S1*** and ***S2***, we messaged participants through Prolific with occasional reminders. We reached out to any participant with zero reported dashboard views on every fourth working day (up to five times). Once participants qualified for ***S2*** (i.e., used extension for seven or more days and viewed the dashboard on at least three separate days), within 24 hours, we invited them to ***S2*** via a Prolific message. If the participant had met the install-day requirement but not the viewing requirement, we sent a reminder message each fourth working day, at most four times.

Participants were compensated $3.00 for successfully completing ***S1***, which was estimated to take 20 minutes (including extension installation, estimated to take five minutes). Participants were rejected from ***S1*** if they failed an attention check (i.e., indicated that they had never heard of Facebook or Gmail) *and* their self-reported time zone did not match system logs (time-zone information used for analysis, see Figure 8 (B)), which might occur due to VPN use.

4.3 Post-Usage Survey: *S2*

The second survey, only available after seven days of extension use and at least three days of dashboard views (full text in Appendix B.3) started by verifying the participant's Prolific ID. The first section of ***S2*** included repeat questions about attitudes, intentions, and the IUIPC. Next, participants were asked to open the extension and answer questions (Likert, five point) about each visualization found in their extension (up to four visualizations each for TT1-Dashboard and TT1-Everything, and up to ten visualizations for TT2). Participants were asked six Likert questions per visualization. These included one question about overall creepiness (*General Creepiness*) and five questions about literature-based factors indicative of creepiness (Creepiness Factors):

- f_1 *Contextual violation*—It is creepy that data brokers could sell this information [66, 89, 124, 129, 136].
- f_2 *Accuracy, me*—Visualization accurately reflects me as a person [23, 30, 131].
- f_3 *Accuracy, browsing*[4]—Visualization accurately reflects my web browsing (novel question).
- f_4 *Personal privacy invasion*—Visualization increases my privacy concern [10, 59, 78, 136].
- f_5 *Motivating to action*—Visualization makes me want to take privacy-protective actions [12, 49, 111].

We also included a free-text question about why each visualization was or was not creepy and a validation question about some aspect of the visualization, to ensure that participants were looking at the right visualization when answering the questions (e.g., "in the 'Your Top Trackers' visualization, what is the #2 tracker listed?").

After commenting on each visualization, we asked participants how much they would pay (free text) to prevent the types of tracking the extension visualized, as well as what actions they might take to lessen tracking (multiple choice selector, with options like "use a privacy-focused browser"). We asked how participants felt creepiness was related to accuracy, personal privacy invasion, and willingness to take action, reminding them how they answered these questions on their most creepy visualization. Participants were asked a yes–no question on whether creepiness is related to each of these three concepts, and a follow-up free text question as to why. Finally, we prompted participants to uninstall the extension.

S2 was designed to take approximately 30 minutes to complete. Participants were compensated $7.00 for successful completion. Participants were deemed unsuccessful for providing unreasonable free text responses. Participants who reported being unable to see

[4] In pilot testing, some participants reported that inferences were relevant to their browsing, in some cases as part of a different Prolific study, but not to their personal interests or situation.

one or more of the visualizations were asked to return the task via Prolific, but were compensated for their time.

4.4 Telemetry

During our field study, we collected pseudonymous telemetry data about users' activity and their interaction with the extension, allowing us to analyze participant opinions with regard to their data (Figure 1, columns S and A). All users were mapped to a string (saved in client-side storage) concatenating their Prolific ID with their condition, allowing us to coordinate between ***S1*** and ***S2*** and ensure that participants had met the required three-visit threshold for ***S2***. Unless otherwise noted, all telemetry data shared with researchers was hashed using SHA-256 (plus a per-participant salt) prior to transmission and encrypted in transit.

The telemetry data we collected concerned: (1) activity data describing interactions with the extension; and (2) tracking data describing the user's profile. Activity data included information like extension page visits, clicks on various parts of the extension, and whether the participant had other tracking-focused extensions installed. Tracking data included both external data sources (e.g., Google Ads Settings data) and internal data sources (e.g., web page timing based on visibility events [76] and the interests we inferred). The only time titles and domains of web pages were shared with researchers without hashing (though still encrypted in transit) was when: (1) the user was logged in to Google Ads Settings; (2) Ads Settings information was updated; and (3) the user was visiting web pages within a five-minute window from the time when the extension logged a change in Ads Settings data. This information was shared in order to assess more fully the way Ads Settings works.

4.5 Ethics and Consent

The two-part field study was approved by the University of Maryland's IRB. Participants were informed of the data collected by the extension, both when consenting to the study and in the extension's privacy policy (Appendix C). The extension was approved by the Chrome Web Store and published as "unlisted." Participants were informed they could opt out of the study at any time. The extension's "settings" page included an opt-out button, which notified researchers and automatically uninstalled the extension, and the extension's pop-up included a tab for a one-click uninstall of the extension. We acknowledge that the browser where the extension was installed may have been used by other individuals. Although we attempted to control for this by informing participants that telemetry data would be shared with researchers (both in the privacy policy via examples and in the survey consent) and reducing the privacy risk by anonymizing and encrypting telemetry data in transit (Section 4.4), we urge future work to highlight this point more clearly during consent.

4.6 Analysis Methods

We use a variety of statistical tests to analyze our data, focusing on: (a) differences between conditions, (b) changes between ***S1*** and ***S2***, (c) per-visualization opinion differences, (d) underlying data differences, and (e) the Creepiness Factors. For (a) differences between conditions, we compare each of the two control conditions against the new visualizations (i.e., not fully pairwise) on attitudes, intents, and IUIPC scores. We first use an ombnibus Kruskal-Wallis (*K-W*) test, and, if significant ($\alpha \leq .05$), a two-tailed Mann-Whitney U (*MWU*) [75]. We do not correct results for MWU tests as these are limited to planned comparisons. For (b) survey-to-survey differences, we only analyze TT2 data, as TT2 is our main condition of interest. We analyze participants' attitudes and intents toward tracking and IUIPC scores using a two-tailed MWU.

To assess (c) per-visualization differences, we likewise analyze TT2 data only. We compare all visualizations, per Creepiness Factors, in pairwise fashion (e.g., *Late-Night Engagement* versus *Ad Explanations*) using a two-tailed MWU. We correct p-values for multiple testing using Benjamini-Hochberg [15]. For (d) data differences, we look to see if participants with similar data answered Likert questions similarly (e.g., to see if later bedtimes correlate with higher *General Creepiness* scores). We use both Spearman's ρ and the multivariate T-test, Hotelling [71, 110]. The Hotelling test looks for equality of mean vectors between two groups, which would help us determine if participants who agreed or strongly agreed that a visualization was creepy had a similar number of sensitive interests in their Ads Settings data as participants who did not find the visualization creepy. For this test, we bucket strongly disagree, disagree, and neither as zero and strongly agree and agree as one [82]). The null hypothesis for the Hotelling test is that the vectors of the groups are the same. We do not use correction for Spearman's ρ or Hotelling; we anticipated these tests to show similarity.

For (e) Creepiness Factors, we use an ordinal logistic regression with *General Creepiness* as the dependent variable (Likert scale, five points). As potential covariates, we included the five Creepiness Factors as well as age, education, and technology experience (Table 2). To avoid overfitting, we used model selection (minimum AIC [17]) while always retaining the Creepiness Factors.

4.7 Limitations

Our extension collects significant information about users' web activities. As such, our participants may be less privacy-conscious on average than the general population. We attempted to mitigate this by ensuring that data was stored locally, limiting the collection of study data, hashing it for privacy, and carefully explaining to participants how their data would be used and protected. Relatedly, participants may have self-selected in part based on an interest in learning more about tracking. Because the extension targets people who wish to better understand how they are tracked, we considered these limitations acceptable, but urge future work to explore these questions among other population samples.

We recruited from Prolific, which provides high-quality and reasonably representative data [91, 123]. However, our sample has demographic limitations typical of crowdsourced studies, including that participants are younger, more educated, and more technically savvy than the U.S. population as a whole [29, 99, 105, 123, 127]. We limited recruitment to U.S. participants in order to study a culture for which we had context. We urge future work to apply a similar approach using similar tools in other cultural contexts.

After recruiting the first set of participants, we discovered that the extension demonstrated lag on computers using certain Apple CPUs. We compensated the few participants who experienced this issue. Going forward, we excluded people with this hardware from

Table 2: Variables used in regression models. We used model selection to choose independent variables, with Creepiness Factors ($f_{1\text{-}5}$) always retained. Likert scores were on a five-point scale, from strongly agree to strongly disagree. The baseline is the first value in the "values" column.

Independent Variables	Theory	Description	Values
f_1 Creepy to sell	Contextual violation	Creepy for data brokers to sell	no || yes
f_2 Accurate, me	Self-descriptive accuracy	Accurately reflects me	no || yes
f_3 Accurate, browsing	*Novel*	Accurately reflects web browsing	no || yes
f_4 Privacy concerning	Personal privacy invasion	Increases my privacy concern	no || yes
f_5 Take action	Motivating to action	Want to take privacy-protective actions	no || yes
Age	Demographic	How old are you	35+ || 34−
Education	Demographic	Highest level of education achieved	+college || −college
Technology Experience	Demographic	Educational background or job field in IT	no || yes
Dependent Variables			
General Creepiness	Creepy this information is associated with me		5-point Likert

Table 3: Participant demographics (rounded).

	Age		Education
18-24	17%	Trade school, Associate's, or less	29%
25-34	36%	Bachelor's or some college	58%
35-44	24%	Master's or more	12%
45-54	12%	Prefer not to say	< 1%
55-64	7%		
65+	4%		
	Gender		**Technology Experience**
Female	59%	No experience in tech. field	72%
Male	39%	Yes experience in tech. field	24%
Non-Binary	2%	Prefer not to say	4%

the study, potentially biasing our sample. Similarly, the extension was only available for desktop users of the Chrome browser, though Chrome has the largest share of the browser market [121].

It is hard to say how well our inferences match large companies' inferences, algorithms for which are closely guarded. Nonetheless, we believe our inferences provide a reasonable example of inferences that trackers *could* make, meaning they are useful for teaching users about tracking. At the same time, some of these inferences, and our visualizations more generally, may have been difficult for users to understand, affecting their opinions. To protect against this, we had participants view the extension's dashboard page on three separate occasions, giving them a greater chance of understanding how the visualizations worked prior to sharing their opinions.

Finally, our study shares common limitations with other online surveys. For example, answering somewhat repetitive questions about different visualizations can lead to fatigue [1]. To mitigate this, we limited the survey length as well as the number of repetitive questions asked. Participants could also respond negatively to tracking and inferencing if they perceived that as the researchers' position (demand effects [80]) or if they felt social pressure to value privacy (social desirability [45, 83]). We attempted to mitigate this by using neutral language and open-ended questions.

5 RESULTS

In this section, we report the results of our longitudinal field study. We start by describing our participants, their use of the extension, and what we learned about web tracking from our telemetry data. Next, we answer our research questions: which tracking practices are creepy, and what makes something creepy.

5.1 Participants

A total of 223 participants successfully completed both ***S1*** and ***S2*** between late August and mid-September 2022 (322 completed ***S1***; 458 returned the task or timed out). From these participants, we excluded 23 who visited fewer than 100 web pages throughout the study (following Weinshel et al. [132]), leaving a total of 200 participants. Due to drop-out from completing ***S1*** but not qualifying for or completing ***S2***, the distribution of final participants per condition varied (although demographics remained roughly the same per condition): 62 in TT1-Dashboard, 65 in TT1-Everything, and 73 in TT2. It took participants an average of 12 minutes to complete ***S1*** and an average of 27 minutes to complete ***S2***. Participant demographics are summarized in Table 3. As is common among crowdsourcing platforms, our participants are younger, more educated, and more tech-savvy than the general population [29, 99, 105, 123, 127].

Browser usage: Participants estimated an average of 74% of their online activity occurred on the browser where our extension was installed. Nearly half (49%) of participants said they make an online (non-app) purchase weekly or more frequently, and 83% said they make this type of purchase monthly or more. Fewer than half (40%) of participants reported having an ad or tracker blocker *currently* installed (a larger portion installed a blocker at any point in the past, 88%). These numbers are representative of the general population [95]. A dedicated tracker blocker (e.g., Disconnect, Firefox tracking protection, Ghostery, and Privacy Badger) was reported to be currently or previously installed by 8% and 10% of participants, respectively. Most participants (59%) reported not seeing the AdChoices icon [67] while browsing the web (18% did, 23% did not know), while nearly half (47%) reported looking at their Google Ads Settings dashboard page [40] at some point in the past (45% had never looked, 8% did not know).

Dashboard engagement: Participants on average visited the extension's dashboard page five times. A quarter of participants clicked the extension's "take action" button, which offered tips to improve online privacy (21% of participants in TT1-Dashboard, 40% in TT1-Everything, and 14% in TT2). Participants were most likely to open the extension for the first time less than ten minutes after completing ***S1*** (48%). No participant opened the extension's dashboard page prior to finishing ***S1***. Some participants opened the extension for the first time *hours* (22%) or *days* (30%) after finishing ***S1***.

Web and tracker activity: During the study, participants visited 231,550 web pages (13,129 unique domains), and encountered 492

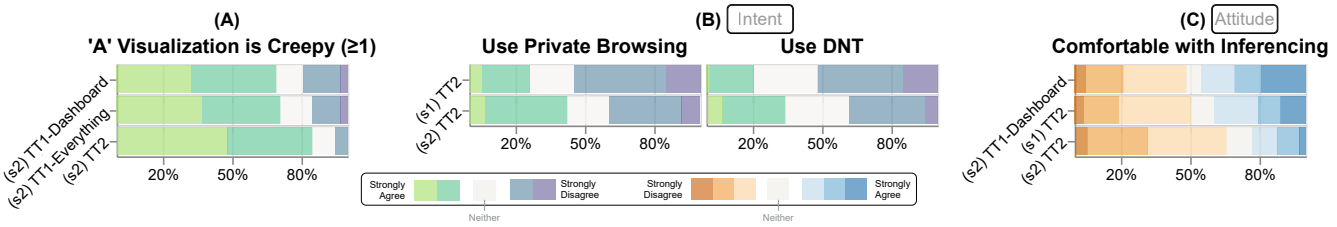

Figure 4: (A) The majority of participants agree or strongly agree that one or more visualizations are creepy—over 80% in TT2. (B) A significant difference, in TT2, between *S1* and *S2* on the intent to use privacy-protective tools like private browsing and Do Not Track (DNT). (C) General discomfort with inferencing, significant between TT1-Dashboard and TT2 and between *S1* and *S2* for TT2.

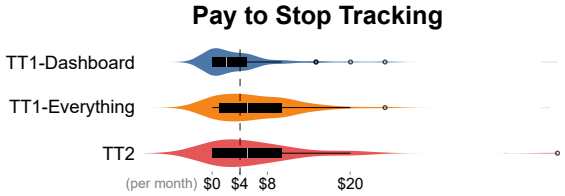

Figure 5: What participants said they would pay per month to stop the tracking shown in the extension. Amounts were significantly lower in TT1-Dashboard than the other conditions (K-W, $p = 0.004$).

unique trackers. Participants, on average, visited 1,159 pages related to 66 domains and encountered three trackers per page. Trackers were present on 51% of web pages visited; if trackers were present, the average number of trackers per page was five. These measurements are similar to those found by Weinshel et al. [132].

5.2 Is Tracking Creepy? Yes

In ***S2***, the vast majority of participants identified at least one visualization as creepy (Figure 4(A)), across all conditions, and most said they would pay (Figure 5) to stop the tracking the extension visualized. We found significant, but small, differences between conditions (Figure 4(C)), as well as before versus after using TT2 (Figure 4(B)). We found no significant difference in IUIPC scores.

Most found something creepy: Most participants (76%) found at least one visualization creepy or very creepy, and that number rose with the increasing amount of information presented in each condition: 69% in TT1-Dashboard, 71% in TT1-Everything, and 85% in TT2 (Figure 4(A)). Most participants also said they would pay to stop the tracking information represented by the visualizations (70%). Participants on average reported a willingness to spend $5.41 per month to stop the types of tracking our extension highlighted (Figure 5). Although the averages in TT1-Everything and TT2 were similar ($6.11 and $6.33 respectively), participants in TT1-Dashboard were only willing to spend about half as much ($3.60). This may be because almost half of the participants in TT1-Dashboard (42%) would not spend any money per month, compared to 22% in TT1-Everything and 27% in TT2. Differences between per-condition pay-to-stop means were significant (K-W, $p = .004$). Pay per month rates for TT1-Everything and TT2 are similar to what prior research has found [20, 116, 135].

TT2 changed participants' attitudes and intents: More than 65% of participants in TT2, after using the extension for one week, were not comfortable with the idea of trackers inferring their interests (Figure 4(C)). This is a significant change compared to before using the extension (MWU $p = 0.024$, small effect size 0.2), when only 50% of participants were uncomfortable. Participants in TT2 also reported greater intent to use privacy-protective tools after using the extension (Figure 4(B)). For intent to use DNT, described as "a browser setting to indicate to web pages you visit that you do not want to be tracked online," participants agreeing or strongly agreeing increased from ~21% in ***S1*** to ~34% in ***S2*** (MWU $p = 0.021$, small effect size 0.2). We found similar results for the use of private browsing mode (26% to 39%, MWU $p = 0.034$, small effect size 0.2).

Participants in TT1-Dashboard were more comfortable with inferencing than those in TT2: Comparing conditions, participants in TT1-Dashboard were statistically significantly more likely to be comfortable with trackers inferring their interests than participants in TT2 (MWU $p = 0.011$, small effect size 0.2). In fact, almost half of TT1-Dashboard participants found inferencing comfortable, compared to less than a quarter in TT2 (Figure 4(C)).

Small effect sizes: Our effect sizes for statistical comparisons are small. We hypothesized that increasing the "creepiness" of TT2 (based on feature testing) would change the way participants felt about online tracking when compared to TT1-Dashboard and TT1-Everything. Although we did find that TT2 was creepy overall, we found smaller-than-expected differences compared to our control conditions. We hypothesize this occurred because all three conditions used the updated interest inference engine and enforced the same viewing requirements. A power analysis, using the point biserial model [47], strengthens this hypothesis, as all comparisons between conditions had ≥80% power.

Comparisons: We make several comparisons between Weinshel et al.'s original study and TT2. These comparisons are inexact. Several years, and intervening events (e.g., the increasing popularity of data protection regulations [52]) have elapsed since Weinshel et al.'s study took place. Additionally, our methods varied somewhat: we used Prolific, versus Amazon Mechanical Turk, and we required participants to view our extension at least three times during the survey. Condition differences also exist. All of our conditions use the updated interest inference engine, and our TT1-Dashboard condition is a stripped-down, aggregated version of TT1-Everything, not found in Weinshel et al.'s study. Because all of our conditions display at least some tracking information, it is perhaps easier for us to show that participants in all conditions found something creepy, but this means that differences among our conditions may be smaller than in the original study.

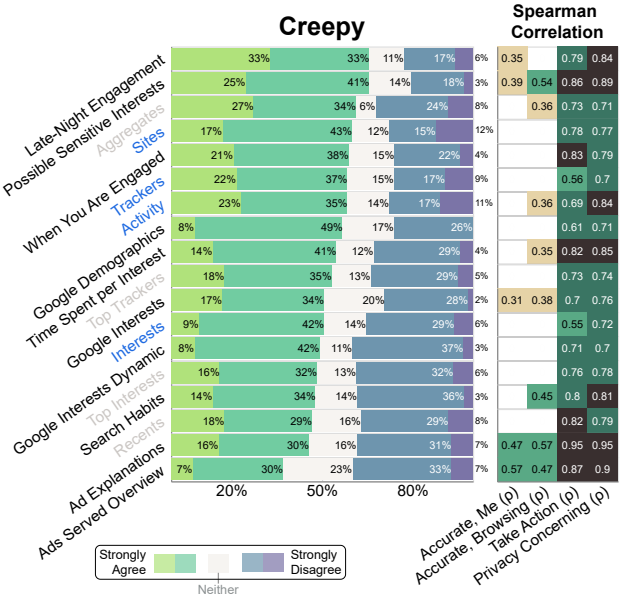

Figure 6: Likert responses to *General Creepiness*, colored as TT1-Dashboard, TT1-Everything, and TT2. Spearman correlation (e.g., *General Creepiness* versus Accurate, Me) shown on the right, p-values corrected with Benjamini-Hochberg [15]. For all Creepiness Factors, see Figure 22.

Figure 7: BH-corrected pairwise MWU tests. The y axis lists visualizations perceived to be creepier (e.g., *Late-Night Engagement*'s creepiness ratings were significantly higher than *Google Interests Dynamic*'s creepiness ratings).

5.3 What Makes Tracking Creepy: Subjectivity

We find that creepiness is far from universal. Although a vast majority of participants found *something* in the extension creepy (Section 5.2), no single visualization was deemed creepy by more than two-thirds of participants. We further analyze this result by looking at whether the sensitivity of the data underlying visualizations or other aspects of the data correlates with creepiness and how the Creepiness Factors affect opinions.

Most found something creepy, but not the same thing: We hypothesized that certain visualizations would be nearly universally perceived as creepy. This proved incorrect. Although the vast majority of participants found at least one visualization in our extension creepy, opinions on which visualizations were creepy did not coalesce to more than two-thirds of the participants (Figure 6).

The *Possible Sensitive Interests* visualization had the most participants agreeing or strongly agreeing it was creepy (66%), but 21% of participants felt otherwise. Likewise, the visualization with the most participants strongly agreeing it was creepy, *Late-Night Engagement* (33%), had 23% of participants disagree or strongly disagree. The same was true when mentioning how data brokers might sell the information visualized. Although 82% of participants who viewed *Google Demographics* deemed it creepy to sell this information, 18% were either unsure or disagreed. Conversely, although *Ads Served Overview* and *Google Interests Dynamic* were viewed as the least creepy (i.e., the highest rate of strongly disagree plus disagree, 40% each), for each of these visualizations, at least 35% of participants agreed or strongly agreed they were creepy.

Next, we consider whether participants found some TT2 visualizations creepier than others (Figure 6). In Figure 7, we compare visualization Likert responses against each other (pairwise, grouped per question) using an MWU corrected for multiple testing with Benjamini-Hochberg [15]. We confirm that *Late-Night Engagement* and *Possible Sensitive Interests* were statistically significantly creepier than many other visualizations. Further, adding context (i.e., informing participants that a data broker could sell this information), highlights creepiness differences among visualizations, with *Google Demographics* and *Late-Night Engagement* being creepier than most other visualizations.

Similar visualizations, differing opinions: To further assess why some participants found something creepy, but others did not, we looked to the data underlying visualizations (Table 4). Our hypothesis was that participants with specific types of data might feel more strongly about a visualization's creepiness. For example, participants with very late bedtimes might find the *Late-Night Engagement* visualization creepier than those with earlier bedtimes (i.e., social norms around appropriate bedtime [48]). To our surprise, this was not the case: participants who viewed substantially similar visualizations felt very differently about their creepiness.

To test this hypothesis, we compare a summarized version of each participant's data (e.g., average bedtime, see Figure 8) with the participant's Likert response to the *General Creepiness* question. We use two similarity metrics, Spearman's ρ [110] (five-point Likert) and a multivariate T-test, Hotelling [71] (two buckets), depending on the type of data being analyzed. We assessed six visualizations.

We find little correlation between Likert responses and underlying data. All visualizations analyzed using Spearman's ρ show near-zero correlation and high p-values, meaning there is insufficient evidence to show a linear relationship between the summarized data and creepiness. Likewise, all of the Hotelling tests returned high p-values as well, allowing us to accept the null hypothesis that the means are similar between groups [57]. In short, creepiness was not strongly related to the data itself: participants' perceptions of creepiness were not grouped by certain types of data (e.g., more sensitive data or data that is more likely to violate a social norm).

Table 4: Does data predict creepiness? No. Spearman correlation is low (i.e., similarity between data and five-point Likert responses) and the null hypothesis for Hotelling is accepted (i.e., similar mean vectors between grouped Likert responses into two buckets).

Visualization	Data Summary (per participant)	Why Possibly Creepy	Similarity Metric	
			Spearman's ρ	*p-value*
Late-Night Engagement	Inferred average bedtime	Later bedtimes	-0.041	0.750
Google Interests	Total number of interests inferred	More interest being inferred	0.065	0.606
Google Interests Dynamic	Count of times attributes change	Frequent updates, more surveillance	-0.069	0.583
				Hotelling p-value
Ads Served Overview	Sensitive ad interests	More sensitive ads served		0.465
Possible Sensitive Interests	Sensitive page interests	More sensitive web page visits		0.262
Google Interests	Sensitive google interests	More sensitive interests in Ads Settings		0.173

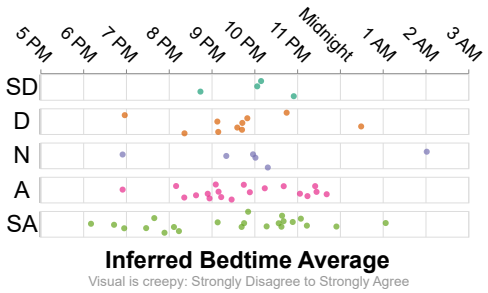

Figure 8: Per-participant responses to *General Creepiness* when viewing the *Late-Night Engagement* visualization, only found in TT2, grouped by response (strongly disagree, SD, to strongly agree, SA). There is no clear trend between Likert scores and inferred bedtimes. See Appendix D, Figure 11, for another example.

Table 5: Regression results showing how much more likely participants are to find something creepy if they agree with a Creep Factor. Coefficients exponentiated to create Odds Ratios (OR); confidence intervals are [2.5%, 97.5%]; statistically significant *p*-values are noted in bold. Pseudo-R^2 (Aldrich-Nelson) is 0.77.

	Odds Ratio	CI	*p*
Privacy Factors			
$f1$ Creepy to sell	9.0	[6.3, 13.1]	**<0.001**
$f4$ Privacy concerning	5.4	[3.7, 8.0]	**<0.001**
$f5$ Take action	3.1	[2.2, 4.4]	**<0.001**
$f3$ Accurate, browsing	1.0	[0.8, 1.3]	0.94
$f2$ Accurate, me	0.8	[0.6, 1.0]	**0.04**
Demographic Covariates			
Age	1.3	[1.1, 1.7]	**0.018**
Tech. Experience	0.7	[0.5, 0.9]	**<0.01**

Creepiness is related to context, privacy invasion, and willingness to take action: We consider how factors of general technology creepiness from the literature (Table 2) apply in our OBA transparency setting. We use a regression to analyze how each of the Creepiness Factors (e.g., the visualization *accurately* describes my browsing) relate to *General Creepiness* (dependent variable).

Table 5 shows how participants were 5.4× and 9.0× more likely to perceive a visualization as creepy (*General Creepiness*) if they also found the visualization invasive of their personal privacy $f4$ or a violation of a social norm $f1$, respectively. This confirms prior work showing how these two factors are indicative of creepiness in the general technology setting (see Table 2). Likewise, participants who reported wanting to take privacy-protective action were three times more likely to increase a step in reported *General Creepiness*. These results mirror responses to Likert-scale questions in ***S2***. Most participants agreed or strongly agreed that privacy concerns (87%) and willingness to take action (80%) were related to creepiness.

Accuracy *somewhat* increases creepiness: Notably, the regression (Table 5) showed some correlation between self-descriptive accuracy and *General Creepiness*, and no relationship between browsing accuracy and *General Creepiness*. To explore this in more detail, we calculated Spearman's ρ between *General Creepiness* and the Creepiness Factors per visualization (Figure 6, *p*-values corrected with Benjamini-Hochberg [15]).

We find that accuracy can decrease comfort, contrary to prior work suggesting that accuracy increases comfort [23, 30, 131]. Self-descriptive accuracy was correlated with *General Creepiness* for visualizations like *Ads Served Overview* or *Ad Explanations*. And browsing accuracy was correlated with *General Creepiness* for visualizations like *Search Habits* and *Possible Sensitive Interests*. In contrast, willingness to take action and personal privacy invasion were significantly correlated for every visualization.

Accuracy decreasing comfort was also supported by participants' survey answers. When asked directly, 61% of participants agreed or strongly agreed that accuracy was related to creepiness. We hypothesize that the 39–61 split occurs because some tracking practices may be creepy regardless of accuracy: "The fact that the advertising companies are *attempting* to get data off me is creepy. The data collected doesn't have to be accurate to be creepy" (P-28).

6 DISCUSSION

We built a design probe to explore creepiness in online tracking and used it in a week-long field study (n = 200). Some participants saw visualizations we designed and feature-tested to be creepy, while others saw basic aggregate information not designed to maximize creepiness. Regardless of condition, the vast majority of participants found one or more visualizations creepy, providing insights on what creepiness means in the OBA context.

Social norms are in flux: Although many people viewed *something* in our extension as creepy (over 80% in TT2), no single visualization was considered creepy by more than 66% of participants. This suggests that norms surrounding OBA are, even several decades after its advent, still forming [124]. Everyone knows it is inappropriate to peek through your neighbors' windows [124], but, as we find, not everyone feels it is inappropriate to collect and target users for having an interest in potentially sensitive topics. This lack of cohesion likely means there is little social pressure on companies to change their tracking practices—i.e., 34% of our

participants were comfortable or unsure with even the visualization most frequently perceived as creepy.

Transparency matters: Participants using our updated extension identified a previously unknown discomfort with companies inferring their interests, suggesting that transparency can improve understanding. However, the specifics of transparency designs matter. For example, we found that adding context strongly affected comfort (e.g., *Google Demographics* was deemed creepy by only 57% of participants, but "creepy to sell" by 82%). Likewise, our presentation of opinionated visualizations (e.g., using red colors to connote late bedtimes in Figure 2 ❻) seems to have affected attitudes more strongly than Weinshel et al.'s original study, which did not find significant differences in comfort pre- to post-extension use. On the one hand, these findings highlight a potential conflict between companies and users: If increases in meaningful transparency negatively affect attitudes toward tracking, it may discourage company-provided transparency, or encourage potentially misleading claims about "not sell[ing] your data" [38] that elide context. On the other hand, the findings point to the importance of design, something we urge future work to assess more deeply (e.g., analyzing how per-visualization design could affect opinions).

Surprising findings on accuracy and sensitivity: We found evidence to suggest that more accurate targeting is creepier, although this contradicts prior work and seems to depend on the specific visualization being presented. We hypothesize that our study was able to uncover this nuance because our extension provided information on real-time browsing data. Accuracy is difficult to measure given the dynamic nature of human interests [103, 115], and unless the measurement tool operates on a real-time basis, like ours did and others did not [30, 131], perceptions of accuracy are likely clouded by time. We also note that our results add empirical evidence to the legal argument that protecting "sensitive" data is fraught with error. Demarcating sensitive from non-sensitive data is not actionable and in my ways counterproductive to privacy protections [117], protecting data is not useful as a binary, sensitive/not-sensitive determination [100], and, as we show, whether data is labeled "sensitive" does not necessarily correlate with perceptions of creepiness. We urge future work to look more closely both at accuracy and sensitivity in this context.

Recommendations: Our findings have several takeaways for regulators, developers, and designers. On the regulatory side, it is clear that the online tracking environment is ripe for regulation—over 80% of participants who viewed our new visualizations were uncomfortable with tracking, but few could agree on what specific tracking practices were creepy. Without being able to unite around a specific discomfort, consumers' privacy is at risk of being invaded by companies who have little incentive to change—an effect that disproportionately affected a small but emphatic group of participants (7%) who *strongly* agreed that every visualization observed was creepy (increasing to 32% when also considering agree). We suggest that regulators consider implementing guardrails around tracking practices that might protect a plurality of users, even if this would be considered unnecessary by some subset of users.

Turning to developers, we note that efforts to demarcate "sensitive," and therefore more protected, attributes (e.g., inferred interests or demographics) may be misleading. For example, Google's Ads Settings page allows consumers to "opt out"[5] of certain types of targeted advertising thought to be sensitive: alcohol, weight loss, parenting, dating, or gambling. While these categories are deserving of the ability to opt-out, our study shows that users may feel the same way about parenting or gambling as they do about an interest in "Massage Therapy." We suggest that developers and designers consider the variance in how users define sensitivity. Conducting context-dependent research to see how users define sensitivity for themselves, in particular circumstances or in particular data use cases, would be a much better route than building tools that have built-in assumptions on what users consider as sensitive or not.

For designers, we note how the design of transparency tools may alter user opinions. Although many of our visualizations used similar underlying data as the visualizations created by Weinshel et al. [132], our study found stronger evidence of pre- to post-survey changes in intents to use privacy-enhancing tools and condition-based differences in attitudes toward tracking. In other words, transparency dashboard design may impact consumer perceptions toward tracking, either heightening or dampening concern. Design decisions, therefore, should be transparent, tested, and documented.

ACKNOWLEDGMENTS

We gratefully acknowledge support from a UMIACS contract under the partnership between the University of Maryland and the Department of Defense. This material is also based upon work supported by the National Science Foundation under Grants No. CNS-2047827, CNS-2149680, and CNS-2151290.

REFERENCES

[1] Lauren S. Aaronson, Cynthia S. Teel, Virginia Cassmeyer, Geri B. Neuberger, Leonie Pallikkathayil, Janet Pierce, Allan N. Press, Phoebe D. Williams, and Anita Wingate. 1999. Defining and measuring fatigue. *The Journal of Nursing Scholarship* 31 (1999).
[2] Gunes Acar, Christian Eubank, Steven Englehardt, Marc Juarez, Arvind Narayanan, and Claudia Diaz. 2014. The web never forgets: Persistent tracking mechanisms in the wild. In *Proc. CCS*.
[3] AdAway Default Blocklist. 2022. AdAway default blocklist: Blocking mobile ad providers and some analytics providers. https://adaway.org/hosts.txt.
[4] Adblock Plus 2.0. 2022. EasyList. https://easylist-downloads.adblockplus.org/easylist.txt.
[5] Akiko Aizawa. 2003. An information-theoretic perspective of TF-IDF measures. *Information Processing & Management* 39 (2003).
[6] Alphabet. 2022. Custom search JSON API: Introduction. https://developers.google.com/custom-search/v1/introduction.
[7] Irwin Altman. 1975. *The environment and social behavior: Privacy, personal space, territory, and crowding.* ERIC.
[8] Patricia Arias-Cabarcos, Saina Kjalili, and Thorsten Strufe. 2023. "Surprised, shocked, worried": User reactions to Facebook data collection from third parties. In *Proc. PETS*.
[9] Ruwan Bandara, Mario Fernando, and Shahriar Akter. 2020. Explicating the privacy paradox: A qualitative inquiry of online shopping consumers. *Journal of Retailing and Consumer Services* 52 (2020).
[10] Natã M. Barbosa, Gang Wang, Blase Ur, and Yang Wang. 2021. Who am I? A design probe exploring real-time transparency about online and offline user profiling underlying targeted ads. In *Proc. IMWUT*.
[11] Lisa Barnard. 2014. The cost of creepiness: How online behavioral advertising affects consumer purchase intention. https://core.ac.uk/download/pdf/210603295.pdf.
[12] Susanne Barth and Menno D.T. de Jong. 2017. The privacy paradox—Investigating discrepancies between expressed privacy concerns and actual online behavior—A systematic literature review. *Telematics and Informatics* 34, 7 (2017).

[5] Google is clear that opting out will only reduce—not eliminate—the appearance of certain types of ads [41].

[13] 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*.

[14] Steven M. Bellovin, Preetam K. Dutta, and Nathan Reitinger. 2019. Privacy and synthetic datasets. *Stanford Technology Law Review* 22 (2019).

[15] Yoav Benjamini and Yosef Hochberg. 1995. Controlling the false discovery rate: A practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society* 57 (1995).

[16] Steven C. Bennett. 2011. Regulating online behavioral advertising. *John Marshall Law Review* 44 (2011).

[17] Hamparsum Bozdogan. 1987. Model selection and Akaike's Information Criterion (AIC): The general theory and its analytical extensions. *Psychometrika* 52, 3 (1987).

[18] Business of Apps. 2022. Cost per click (CPC) rates 2022. https://www.businessofapps.com/ads/cpc/research/cpc-rates/.

[19] California. 2018. California Consumer Protection Act (CCPA). *California Civil Code §§ 1798.100–1798.199.100* (2018).

[20] Farah Chanchary and Sonia Chiasson. 2015. User perceptions of sharing, advertising, and tracking. In *Proc. SOUPS*.

[21] Hsuan-Ting Chen. 2018. Revisiting the privacy paradox on social media with an extended privacy calculus model: The effect of privacy concerns, privacy self-efficacy, and social capital on privacy management. *American Behavioral Scientist* 62, 10 (2018).

[22] Hichang Cho, Pengxiang Li, and Zhang Hao Goh. 2020. Privacy risks, emotions, and social media: A coping model of online privacy. *ACM Transactions on Computer-Human Interaction* 27, 6 (2020).

[23] Rena Coen, Emily Paul, Pavel Vanegas, Alethea Lange, and G.S. Hans. 2016. A user-centered perspective on algorithmic personalization. https://www.ischool.berkeley.edu/sites/default/files/projects/algorithmic-personalization-coen-paul-vanegas.pdf.

[24] Jessica Colnago, Lorrie Faith Cranor, and Alessandro Acquisti. 2023. Is there a reverse privacy paradox? An exploratory analysis of gaps between privacy perspectives and privacy-seeking behaviors. In *Proc. PETS*.

[25] Kovila P.L. Coopamootoo and Thomas Groß. 2017. Why privacy is all but forgotten: An empirical study of privacy & sharing attitude. *Proc. PETS* (2017).

[26] Kovila P.L. Coopamootoo and Maryam Mehrnezhad. 2022. "I feel invaded, annoyed, anxious and I may protect myself": Individuals' feelings about online tracking and their protective behaviour across gender and country. In *Proc. USENIX Security*.

[27] Savino Dambra, Iskander Sanchez-Rola, Leyla Bilge, and Davide Balzarotti. 2022. When Sally met trackers: Web tracking from the users' perspective. In *Proc. USENIX Security*.

[28] Michal Mimino Danilak. 2022. Language detection library ported from Google's language-detection. https://pypi.org/project/langdetect/.

[29] Djellel Difallah, Elena Filatova, and Panos Ipeirotis. 2018. Demographics and dynamics of Mechanical Turk workers. In *Proc. WSDM*.

[30] 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*.

[31] Charles Duhigg. 2013. How companies learn your secrets. In *The Best Business Writing*. Columbia University Press.

[32] Mohan J. Dutta-Bergman. 2006. The demographic and psychographic antecedents of attitude toward advertising. *Journal of Advertising Research* 46 (2006).

[33] 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*.

[34] European Parliament. 2016. Regulation (EU) 2016/679 of the European Parliament and the Council of 27 April 2016 on the Protection of Natural Persons with Regard to the Processing of Personal Data and on the Free Movement of Such Data, and Repealing Directive 95/46/EC (General Data Protection Regulation). *L 119, Official Journal of the European Union* (2016).

[35] Fariborz Farahmand and Firoozeh Farahmand. 2019. Privacy decision making: The brain approach. *Computer* 52, 4 (2019).

[36] Florian M. Farke, David G. Balash, Maximilian Golla, Markus Dürmuth, and Adam J. Aviv. 2021. Are privacy dashboards good for end users? Evaluating user perceptions and reactions to Google's My Activity. In *Proc. USENIX Security*.

[37] Lisa Farman, Maria Leonora Comello, and Jeffrey R. Edwards. 2020. Are consumers put off by retargeted ads on social media? Evidence for perceptions of marketing surveillance and decreased ad effectiveness. *Journal of Broadcasting & Electronic Media* 64 (2020).

[38] Christian Fuchs. 2012. The political economy of privacy on Facebook. *Television & New Media* 13 (2012).

[39] Google. 2018. Greater transparency and control over your Google ad experience. https://blog.google/technology/ads/greater-transparency-and-control-over-your-google-ad-experience/.

[40] Google. 2019. Google ad settings. https://adssettings.google.com.

[41] Google. 2023. Frequently asked questions. https://support.google.com/My-Ad-Center-Help/answer/12155964?hl=en.

[42] Google. 2023. Your ads, your choice. https://myadcenter.google.com.

[43] Google Cloud. 2022. Content categories. https://cloud.google.com/natural-language/docs/categories#version_2.

[44] Google Cloud. 2022. Natural language API. https://cloud.google.com/natural-language.

[45] Pamela Grimm. 2010. Social desirability bias. In *Wiley International Encyclopedia of Marketing*. Wiley Online Library.

[46] Thomas Groß. 2021. Validity and reliability of the scale Internet Users' Information Privacy Concerns (IUIPC). In *Proc. PETS*.

[47] Das S. Gupta. 1960. Point biserial correlation coefficient and its generalization. *Psychometrika* 25 (1960).

[48] Rachel Hall. 2020. Extreme night owls: 'I can't tell anyone what time I go to bed'. https://www.theguardian.com/lifeandstyle/2020/may/31/extreme-night-owls-i-cant-tell-anyone-what-time-i-go-to-bed.

[49] Julia Hanson, Miranda Wei, Sophie Veys, Matthew Kugler, Lior Strahilevitz, and Blase Ur. 2020. Taking data out of context to hyper-personalized ads: Crowdworkers' privacy perceptions and decisions to disclose private information. In *Proc. CHI*.

[50] Samantha Hautea, Anjali Munasinghe, and Emilee Rader. 2020. "That's not me": Surprising algorithmic inferences. In *Proc. CHI*.

[51] Daniel C. Howe and Helen Nissenbaum. 2017. Engineering privacy and protest: A case study of AdNauseam. In *Proc. IWPE*.

[52] IAPP. 2022. US state privacy legislation tracker. https://iapp.org/media/pdf/resource_center/State_Comp_Privacy_Law_Chart.pdf.

[53] InVision. 2022. Online whiteboard meets productivity platform. https://www.invisionapp.com.

[54] Mark Irvine. 2022. Google Ads benchmarks for your industry [updated!]. https://www.wordstream.com/blog/ws/2016/02/29/google-adwords-industry-benchmarks.

[55] Bernard J. Jansen. 2007. Click fraud. *Computer* 40, 7 (2007).

[56] Bernard J. Jansen, Kathleen Moore, and Stephen Carman. 2013. Evaluating the performance of demographic targeting using gender in sponsored search. *Information Processing & Management* 49 (2013).

[57] Francisco Juretig. 2019. *R statistics cookbook*. Packt Publishing.

[58] Flavius Kehr, Tobias Kowatsch, Daniel Wentzel, and Elgar Fleisch. 2015. Blissfully ignorant: The effects of general privacy concerns, general institutional trust, and affect in the privacy calculus. *Information Systems Journal* 25, 6 (2015).

[59] Vera Khovanskaya, Eric P.S. Baumer, Dan Cosley, Stephen Voida, and Geri Gay. 2013. "Everybody knows what you're doing": A critical design approach to personal informatics. In *Proc. CHI*.

[60] Minji Kim, Sarah Olson, Jeffrey W. Jordan, and Pamela M. Ling. 2020. Peer crowd-based targeting in E-cigarette advertisements: A qualitative study to inform counter-marketing. *BMC Public Health* 20 (2020).

[61] Agnieszka Kitkowska, Mark Warner, Yefim Shulman, Erik Wästlund, and Leonardo A. Martucci. 2020. Enhancing privacy through the visual design of privacy notices: Exploring the interplay of curiosity, control and affect. In *Proc. SOUPS*.

[62] Spyros Kokolakis. 2017. Privacy attitudes and privacy behaviour: A review of current research on the privacy paradox phenomenon. *Computers & Security* 64 (2017).

[63] Jon A. Krosnick, Sowmya Narayan, and Wendy R. Smith. 1996. Satisficing in surveys: Initial evidence. *New Directions for Evaluation* 1996, 70 (1996).

[64] Ivar Krumpal. 2013. Determinants of social desirability bias in sensitive surveys: A literature review. *Quality & Quantity* 47, 4 (2013).

[65] Nir Kshetri. 2010. The economics of click fraud. In *Proc. IEEE S&P*.

[66] Markus Langer and Cornelius J. König. 2018. Introducing and testing the creepiness of situation scale (CROSS). *Frontiers in Psychology* 9 (2018).

[67] Pedro Giovanni Leon, Justin Cranshaw, Lorrie Faith Cranor, Jim Graves, Manoj Hastak, Blase Ur, and Guzi Xu. 2012. What do online behavioral advertising privacy disclosures communicate to users?. In *Proc. WPES*.

[68] 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*.

[69] Han Li, Rathindra Sarathy, and Heng Xu. 2011. The role of affect and cognition on online consumers' decision to disclose personal information to unfamiliar online vendors. *Decision Support Systems* 51, 3 (2011).

[70] Peter Lowe. 2022. Ad servers. https://pgl.yoyo.org/adservers/serverlist.php?hostformat=webclean.

[71] Yan Lu, Peng-Yuan Liu, Peng Xiao, and Hong-Wen Deng. 2005. Hotelling's T2 multivariate profiling for detecting differential expression in microarrays. *Bioinformatics* 21, 14 (2005).

[72] Jonathan R. Mayer and John C. Mitchell. 2012. Third-party web tracking: Policy and technology. In *Proc. IEEE S&P*.

[73] Francis T. McAndrew and Sara S. Koehnke. 2016. On the nature of creepiness. *New Ideas in Psychology* 43 (2016).

[74] Aleecia McDonald and Lorrie Faith Cranor. 2010. Beliefs and behaviors: Internet users' understanding of behavioral advertising. In *Proc. TPRC*.

[75] Patrick E. McKnight and Julius Najab. 2010. Mann–Whitney U Test. In *The Corsini encyclopedia of psychology*. Wiley Online Library.

[76] MDN web docs. 2022. Document: visibilitychange event. https://developer.mozilla.org/en-US/docs/Web/API/Document/visibilitychange_event.

[77] Mondovo. 2022. Top searched keywords: Lists of the most popular Google search terms across categories. https://www.mondovo.com/keywords/.

[78] Robert S. Moore, Melissa L. Moore, Kevin J. Shanahan, and Britney Mack. 2015. Creepy marketing: Three dimensions of perceived excessive online privacy violation. *Marketing Management Journal* 25, 1 (2015).

[79] Mozilla. 2022. Readability node package. https://www.npmjs.com/package/@mozilla/readability.

[80] Jonathan Mummolo and Erik Peterson. 2019. Demand effects in survey experiments: An empirical assessment. *American Political Science Review* 113, 2 (2019).

[81] Sigrun Myhrvold and Mari-Ann Sekkenes Hamre. 2018. *Too creepy for comfort? A study of personalized online advertising effects on attitude towards the ad and the advertised brand across high/low involvement and socially sensitive products, and the mediating role of the creepiness factor.* Master's thesis. BI Norwegian Business School.

[82] Michael J. Nanna. 2002. Hotelling's T^2 vs. the rank transform with real Likert data. *Journal of Modern Applied Statistical Methods* 1, 1 (2002).

[83] Anton J. Nederhof. 1985. Methods of coping with social desirability bias: A review. *European Journal of Social Psychology* 15 (1985).

[84] Patricia A. Norberg, Daniel R. Horne, and David A. Horne. 2007. The privacy paradox: Personal information disclosure intentions versus behaviors. *Journal of Consumer Affairs* 41 (2007).

[85] OKO. 2019. The history of online advertising. https://oko.uk/blog/the-history-of-online-advertising.

[86] Lukasz Olejnik, Tran Minh-Dung, and Claude Castelluccia. 2013. Selling off privacy at auction. In *Proc. NDSS*.

[87] Leysia Palen and Paul Dourish. 2003. Unpacking privacy for a networked world. In *Proc. CHI*.

[88] Saurabh Panjwani, Nisheeth Shrivastava, Saurabh Shukla, and Sharad Jaiswal. 2013. Understanding the privacy-personalization dilemma for web search: A user perspective. In *Proc. CHI*.

[89] Emmi Parviainen and Marie Louise Juul Søndergaard. 2020. Experiential qualities of whispering with voice assistants. In *Proc. CHI*.

[90] Delroy L. Paulhus. 1991. Measurement and control of response bias. *Measures of Personality and Social Psychology Attitudes* (1991).

[91] Eyal Peer, David Rothschild, Andrew Gordon, Zak Evernden, and Ekaterina Damer. 2022. Data quality of platforms and panels for online behavioral research. *Behavior Research Methods* 54 (2022).

[92] Chanda Phelan, Cliff Lampe, and Paul Resnick. 2016. It's creepy, but it doesn't bother me. In *Proc. CHI*.

[93] Victor Le Pochat, Tom Van Goethem, Samaneh Tajalizadehkhoob, Maciej Korczyński, and Wouter Joosen. 2019. Tranco: A research-oriented top sites ranking hardened against manipulation. In *Proc. NDSS*.

[94] Prolific. 2021. https://www.prolific.co.

[95] Enric Pujol, Oliver Hohlfeld, and Anja Feldmann. 2015. Annoyed users: Ads and ad-block usage in the wild. In *Proc. IMC*.

[96] Kristen Purcell, Lee Rainie, and Joanna Brenner. 2012. Search engine use 2012. https://www.ris.org/uploadi/editor/1341041853PIP_Search_Engine_Use_2012.pdf.

[97] Emilee J. Rader. 2014. Awareness of behavioral tracking and information privacy concern in Facebook and Google. In *Proc. SOUPS*.

[98] Ashwini Rao, Florian Schaub, and Norman Sadeh. 2014. What do they know about me? Contents and concerns of online behavioral profiles. In *Proc. ASE*.

[99] Elissa M. Redmiles, Sean Kross, and Michelle L. Mazurek. 2019. How well do my results generalize? Comparing security and privacy survey results from Mturk, web, and telephone samples. In *Proc. IEEE S&P*.

[100] Nathan Reitinger and Amol Deshpande. 2023. Epsilon-differential privacy, and a two-step test for quantifying reidentification risk. *Jurimetrics: The Journal of Law, Science & Technology* 63, 3 (2023).

[101] Nathan Reitinger and Michelle L. Mazurek. 2021. ML-CB: Machine learning canvas block. In *Proc. PETS*.

[102] Nathan Reitinger, Bruce Wen, Michelle L. Mazurek, and Blase Ur. 2023. Analysis of Google Ads Settings over time: Updated, individualized, accurate, and filtered. In *Proc. WPES*.

[103] Ann K. Renninger and Rose K. Pozos-Brewer. 2015. Psychology of interest. *International Encyclopedia of the Social & Behavioral Sciences* 12 (2015).

[104] Frank Rosenblatt. 1958. The perceptron: A probabilistic model for information storage and organization in the brain. *Psychological Review* 65, 6 (1958).

[105] Joel Ross, Lilly Irani, M. Six Silberman, Andrew Zaldivar, and Bill Tomlinson. 2010. Who are the Turkers? Worker demographics in Amazon Mechanical Turk. In *Proc. CHI EA*.

[106] Benjamin Saunders, Julius Sim, Tom Kingstone, Shula Baker, Jackie Waterfield, Bernadette Bartlam, Heather Burroughs, and Clare Jinks. 2018. Saturation in qualitative research: Exploring its conceptualization and operationalization. *Quality & Quantity* 52, 4 (2018).

[107] Quirin Scheitle, Oliver Hohlfeld, Julien Gamba, Jonas Jelten, Torsten Zimmermann, Stephen D Strowes, and Narseo Vallina-Rodriguez. 2018. A long way to the top: Significance, structure, and stability of Internet top lists. In *Proceedings of the Internet Measurement Conference 2018*.

[108] Quirin Scheitle, Oliver Hohlfeld, Julien Gamba, Jonas Jelten, Torsten Zimmermann, Stephen D. Strowes, and Narseo Vallina-Rodriguez. 2022. Toppling top lists: Evaluating the accuracy of popular website lists. In *Proc. IMC*.

[109] Sebastian Schnorf, Martin Ortlieb, and Nikhil Sharma. 2014. Trust, transparency & control in inferred user interest models. In *Proc. CHI*.

[110] Patrick Schober, Christa Boer, and Lothar A. Schwarte. 2018. Correlation coefficients: Appropriate use and interpretation. *Anesthesia & Analgesia* 126, 5 (2018).

[111] John S. Seberger, Irina Shklovski, Emily Swiatek, and Sameer Patil. 2022. Still creepy after all these years: The normalization of affective discomfort in app use. In *Proc. CHI*.

[112] Dan Shewan. 2022. The comprehensive guide to online advertising costs. https://www.wordstream.com/blog/ws/2017/07/05/online-advertising-costs.

[113] Irina Shklovski, Scott D. Mainwaring, Halla Hrund Skúladóttir, and Höskuldur Borgthorsson. 2014. Leakiness and creepiness in app space: Perceptions of privacy and mobile app use. In *Proc. CHI*.

[114] Reza Shokri, Marco Stronati, Congzheng Song, and Vitaly Shmatikov. 2017. Membership inference attacks against machine learning models. In *Proc. IEEE S&P*.

[115] Paul J. Silvia. 2006. *Exploring the psychology of interest.* Psychology of Human Motivation.

[116] Daniel J. Solove. 2021. The myth of the privacy paradox. *The George Washington Law Review* 89, 1 (2021).

[117] Daniel J. Solove. 2024. Data is what data does: Regulating based on harm and risk instead of sensitive data. *Northwestern University Law Review* 118 (2024).

[118] Roseanna Sommers and Vanessa K. Bohns. 2019. The voluntariness of voluntary consent: Consent searches and the psychology of compliance. *Yale Law Journal* 128 (2019).

[119] Sarah Spiekermann, Jens Grossklags, and Bettina Berendt. 2001. E-privacy in 2nd generation E-commerce: Privacy preferences versus actual behavior. In *Proc. EC*.

[120] Luke Stark. 2016. The emotional context of information privacy. *The Information Society* 32, 1 (2016).

[121] Statcounter. 2022. Browser market share worldwide. https://gs.statcounter.com/browser-market-share/desktop/united-states-of-america.

[122] Arlonda M. Stevens. 2016. *Antecedents and outcomes of perceived creepiness in online personalized communications.* Ph. D. Dissertation. Case Western Reserve University.

[123] Jenny Tang, Eleanor Birrell, and Ada Lerner. 2022. Replication: How well do my results generalize now? The external validity of online privacy and security surveys. In *Proc. SOUPS*.

[124] Omer Tene and Jules Polonetsky. 2013. A theory of creepy: Technology, privacy and shifting social norms. *Yale Journal of Law & Technology* 16 (2013).

[125] tgc. 2021. Host database. https://sos-ch-dk-2.exo.io/noblt/RPZ/Hosts-database/full-alive.txt.

[126] Helma Torkamaan, Catalin-Mihai Barbu, and Jürgen Ziegler. 2019. How can they know that? A study of factors affecting the creepiness of recommendations. In *Proc. RecSys*.

[127] Anne M. Turner, Thomas Engelsma, Jean O. Taylor, Rashmi K. Sharma, and George Demiris. 2020. Recruiting older adult participants through crowdsourcing platforms: Mechanical Turk versus Prolific Academic. In *Proc. AMIA*.

[128] Joseph Turow, Jennifer King, Chris Jay Hoofnagle, Amy Bleakley, and Michael Hennessy. 2009. Americans reject tailored advertising and three activities that enable it. https://repository.upenn.edu/cgi/viewcontent.cgi?article=1551&context=asc_papers.

[129] 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*.

[130] Sophie Veys, Daniel Serrano, Madison Stamos, Margot Herman, Nathan Reitinger, Michelle L. Mazurek, and Blase Ur. 2021. Pursuing usable and useful data downloads under GDPR/CCPA access rights via co-design. In *Proc. SOUPS*.

[131] Miranda Wei, Madison Stamos, Sophie Veys, Nathan Reitinger, Justin Goodman, Margot Herman, Dorata Filipczuk, Ben Weinshel, Michelle L. Mazurek, and Blase Ur. 2020. What Twitter knows: Characterizing ad targeting practices, user perceptions, and ad explanations through users' own Twitter data. In *Proc. USENIX Security*.

[132] 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*.

[133] Kenneth C. Wilbur and Yi Zhu. 2009. Click fraud. *Marketing Science* 28 (2009).
[134] Craig E. Wills and Mihajlo Zeljkovic. 2011. A personalized approach to web privacy: awareness, attitudes and actions. *Information Management & Computer Security* 19, 1 (2011).
[135] Angela G. Winegar and Cass R. Sunstein. 2019. How much is data privacy worth? A preliminary investigation. *Journal of Consumer Policy* 42 (2019).
[136] Paweł W. Woźniak, Jakob Karolus, Florian Lang, Caroline Eckerth, Johannes Schöning, Yvonne Rogers, and Jasmin Niess. 2021. Creepy technology: What is it and how do you measure it?. In *Proc. CHI*.
[137] Jason C. Yip, Kiley Sobel, Xin Gao, Allison Marie Hishikawa, Alexis Lim, Laura Meng, Romaine Flor Ofiana, Justin Park, and Alexis Hiniker. 2019. Laughing is scary, but farting is cute: A conceptual model of children's perspectives of creepy technologies. In *Proc. CHI*.
[138] Hui Zhang, Munmun De Choudhury, and Jonathan Grudin. 2014. Creepy but inevitable? The evolution of social networking. In *Proc. CSCW*.

A AD URL PARSING

Example full URL flagged by TT2 as originating from an ad server. This URL has an inferred final destination of `hbomax.com`.

https://adclick.g.doubleclick.net/pcs/click?xai=AKAOjsvMMrslT6dHWYrdeiUhecuOlEm5TM8QsYqirhqjzeMhW8SCiMoYC4SAvAoEfL-fitrU0l-PB-m2IaFF-hdJyze_zcYptqWH2auEgBEcUPHvB2X7yC1_uNc7RM9UhgX_bGFf4tcc3VV5FrD7dNt1RQCMrN5q6Djxo3PD1wA86Zo2dJ-7H8Ybm9BOxVIaEmApq7gSEiBH_SqsUk9z0ZtVylaXWQpCM_egu3ENsgjwYM_quF8BNK5TASsH24Q7U5z3pUXdMW2mhfoACU2SVESMprbAmOMsp_Vi3DsHOuyqSufF0qwUfBDCNiBjLksGCpJMTYcAaDshFb6_7RnkEp6iGpuX9eLoQ6_A3nseqTlgKSbaAtyOEft_dHWm1BNop-q_qbyRM_I29NoDjoNQz_09OWqz-0rXGbwecmwJHp8FQND4tjlRbep511NQVZ0mI4FHezn4x4jNHybqzVEqp-mNR8-CqwscG8twn7jSPG8oM4VNHDGcSJ-8iHFe59_p8AarjYBOFNIn8JqfZipCydEmoNykO2b2TFxvvynC_zi8y504Z9Oj0ku9_14vdf3YuD4w8ZSRr5gkDg0ZJgQxuzghiEDABS0_5E4wbYGy2FwI445_fbbrg0GEkdfOBdapOM9gjgIR9isqIkShbeV1ozlsQZ997ODv1CSG4dRiCB1IAMxU9XPau1BQf-IUbUlbN5RehLYZCmcQARIgQLCAAUXZBGbVv4f3YGxyKhEJfv6KuL_mAHWIJsp1sqgrGO4QhKjxMqyPAszx-qHPwgNBDjwf3JJdSzzb7HpoFWMImjdP5tkGFwrhtxZvxPJMq5Lk91eLsgCot0XZqLeWN02aiQii138aA8Gt0x8qA1qff0W3seciVdpgmrmXn5J7Z9HeOI0F8RnEy5xTaxew9IkPTUiDGoxKMJZPmVx18yyaDZmZvQ9tTdQdp7FOP3EByJiaUHtYD5Pu19YAh3IGHBT13O7_q9We9q4J5gLdUSwIO5o5OQp-LTiFWKMoWXeiN-V7pX8urXUDgSTwbCa7osgSmx4brR3006RUZza0PnPP5-SN6-ABNEfoVVnI08tJ4IR0Sa28r4GT7f7GNMeMR6fiF20XJRo12O6QVPM9TKLhzQlhGLPNyh1kLlXrMS1vPdFKv2_IZAmH2wWDrhwMmzFGJw1z-qxB0m6thKI3Rat6pNQ6pKUGfUtDe99Mp1qh2C1FjEYW_KSsOkAgIVMOWYFPbRzs6Uy2AHcxcBv_t28GKwKzWS4MP10R7y3WvdV2KOK2hVM&sai=AMfl-YSF2fgCvlEfYN1T8whFt6OZylFCnxVDx-t_wopggfCgC4oTmgCGTKfRXPnQPnEzIUIeH3NFv2h27EPVxuRQEPjcewqkm_Cjxx72ixa2d2jBsxmJaPoIlcFdUUlPsTBPEkqAWLJrsCjfDVy6EQvYRpxBW0oRmZkEL9xzLisE&sig=Cg0ArKJSzLswARixxT11&pr=13:YxeMowAAAABSVj39_Uir4LVRCbyNK-Zj7bBA6Q&fbs_aeid=[gw_fbsaeid]&urlfix=1&adurl=https://www.hbomax.com/series/house-of-the-dragon?utm_id=cm|27770977|2414963|337969533|176266753&dclid=%edclid!.

B QUESTIONNAIRES

B.1 Feature Testing Interview

Welcome. Thank you for participating in our study. The purpose of this study is to inform the design of an app to help users like you learn more about browsing the internet and online trackers. You are allowed to leave at any time. Today's study has two parts. First, I have a couple of background questions about your experiences with online tracking, and I'll also explain what the app is supposed to do. Second is the main part, where we'll have you visit some web pages and then walk through a section of the app. At the end, we have a few short overall questions. As a reminder, as stated in the consent form, we will record your screen and what you say for later analysis. We will remove any identifying information before we analyze the recordings. I will now start the recording, please make sure your video is off and you are sharing your screen.

- In your own words, could you explain to me what you know/-think about online tracking?
- How do you feel about online tracking?

Today we are testing an app called "Tracking Transparency." It was developed by researchers at the University of Chicago and the University of Maryland. The app is a browser extension that gives you an advertiser's (i.e., trackers) perspective of your online habits—what they can learn **about you**, what might be **sensitive** and **unique**, and what this means in terms of what **ads** are shown to you.

We were hired by the researchers to get feedback on their Tracking Transparency app. There is not one particular design the researchers hope you'll like better than the others; they're most interested in your honest and blunt feedback. As you go through the app, I would like you to think aloud for me as you answer.

Thinking aloud means saying, out loud, whatever comes into your head as you use the site and decide what to do next. As an example, if I thought aloud while trying to remember what I had for dinner earlier this week... *<Do example>*

Now, I'd like you to give it a try: Think aloud while answering the question: How many windows are there in the home where you grew up?

Let's get started. I'm going to have you click on the link which will start our demo. *<Give link>*

Please click this link. This is the starting page. Today, we're going to be working on [feature set {about you, unique you, sensitive you, ads}]. Please click on [feature set]. This software will guide you through a few websites with prompts and then we'll show you the parts of the app we're hoping to get your feedback on. Ok, now please follow the prompts, let me know at any time if you get stuck, and please remember to think aloud.

Per-Visual Questions

1. What do you think of this *(open-ended)*
   a. What do you think this visualization is showing you?

2. Did you learn anything from this?
   a. [If yes] What?
   b. Is anything here new or surprising to you?
3. Is anything here confusing
   a. [If yes] What? And what would you change or add to make this less confusing?
4. Is there anything you'd like to know that isn't covered here, or anything you want to see added to this visualization?
5. How does this make you feel
   a. Does anything about this make you **happy**, or is something you find enjoyable
   b. Does anything about this make you **sad/upset,** or is something you do not find enjoyable
      i. Would you say this is creepy
      ii. [if sad/creepy] If this could be stopped without impacting your experience online, would you stop it?
         1. If it was a little harder, and possibly made things difficult to do online (like logging in more times) would you feel the same way
6. Let's look at another visualization here [loop back to 1]

Overall Questions

1. Do you have any thoughts about all of what you've just seen as a group (*open ended*)
2. What was most surprising (or interesting) to you
3. What do you think you had the strongest reaction to (most happy or most sad or most creeped out)
4. Do you think you would want to use a tool like this with your real browser? Why or why not?
   a. What do you think you would use it for?
5. After looking at a tool like this, do you think you would change anything about your web browser or your browsing habits? Why or Why not?
   a. If yes, what would you change?
6. Did you learn anything about online tracking during this session today? If so, what did you learn?
7. Have your feelings about online tracking changed at all after this session? Why or why not?
   a. If yes, how have they changed?

Thank you for your participation. Is there anything else you would like to share about your experience today? We are very grateful for all your comments today, and will be passing them on to the researchers for their final design. You will be paid through the survey provider and if you have any questions about this research, you may contact our Principal Investigator or the IRB at the contact info on the consent form. Thank you again!

B.2 Field Study, Part I

<validate desktop device> <note about Apple CPUs>

Please note, this study requires a desktop computing device using Google Chrome and will require you to download an extension from the Chrome Web Store. Additionally, the extension you will download is largely incompatible with newer Apple computers relying on Apple Silicone processing chips. If you have an Apple computer (apple icon in the top left-hand corner of your computer screen > "About This Mac" > "Chip" > "Apple [M1, M2, or variant]") with an M1 or M2 chip, then please return this survey. Please return this task if you are unable to do so.

- I do not have an Apple computer with an M1 or M2 chip
- I am willing to download a Google Chrome extension from the Chrome Web Store

Thank you for participating in Part I of our two-part study! In this first part, you will:

- Install a Chrome browser extension
- Answer preliminary questions about your attitudes and opinions on the Internet ecosystem
- In closing, answer a few demographic-type questions

The survey should take approximately 20 minutes to complete, including time to download the extension.

Consenting Instrument

Which of the following browsers do you regularly use? Select all that apply.

- Chrome
- Firefox
- Safari
- Opera
- Internet Explorer/Edge
- Epic
- Brave
- Firefox Focus
- Tor
- Other *<free-text>*

What percentage of your online browsing is on the **device and browser** you are using right now, compared to other devices or other browsers?
<slider 0-100> <less on this device to more on this device>

How often do you make purchases online using a web browser (as opposed to through an app)?

- Never
- Rarely
- Monthly
- Weekly
- Daily
- Multiple times a day
- Don't know

Have you ever heard of or used the following software, browser extensions, websites, or tools? {don't use it and have never heard of it, don't use it, but have heard of it, previously used it, currently use it}

- AdBlock Plus
- AdBlock
- Disconnect
- Facebook

- Firefox Tracking Protection
- Ghostery
- Gmail
- HTTPS Everywhere
- Privacy Badger
- uBlock Origin

Have you seen this icon <adChoices> while browsing online?

- Yes
- No
- Don't know

Have you ever looked at your Google ad settings (partial example shown below)?

- Yes
- No
- Don't know

Please download the Tracking Transparency Chrome Extension before continuing. Downloading and installing this browser extension is required in order to successfully complete this survey.

- Download the Chrome extension here
- Add the extension to Chrome
- Wait for the automated pop-up page!
- Pin the extension (see the automated pop-up page for how to do this)

<message passing, verify install, set condition>

Thank you for installing the Tracking Transparency extension!

If you do not see the next button:

- Make sure you've installed the Tracking Transparency extension (here)
- Make sure you've seen the automated pop-up page!
- If you still do not see the next button, it means there was an error installing the extension, please return the task or contact us at: trackingtransparency@gmail.com.

During the rest of this survey, we use the term "online advertising companies" to refer to companies that show you advertisements online. Note that the companies selecting and displaying advertisements are distinct from the companies whose products are being advertised.

Please select the answer choices that best describes your agreement or disagreement with the statements shown below.

Attitudes Block

I would like to see ads that are relevant to my interests, as opposed to generic ads.

- Strongly agree
- Agree
- Somewhat agree
- Neither agree nor disagree
- Somewhat disagree
- Disagree
- Strongly disagree

I would be comfortable with "online advertising companies" guessing my interests based on which websites I visit. *<same selectors as previous question>*

If it were available, I would like to use a system that shows me what information has been collected about me online. *<same selectors as previous question>*

I feel that "online advertising companies" adequately explain why I received a particular ad. *<same selectors as previous question>*

I feel that I understand how "online advertising companies" determine which advertisements I see. *<same selectors as previous question>*

I would consider it **fair** for advertising companies to track which websites I visit in order to show me ads that are relevant to my interests. *<same selectors as previous question>*

I would consider it **creepy** for advertising companies to track which websites I visit in order to show me ads that are relevant to my interests. *<same selectors as previous question>*

Intents Block

How likely are you to seek out more information about online advertising?

- Extremely likely
- Likely
- Neutral
- Unlikely
- Extremely unlikely
- Don't know

How likely are you to use a browser's private browsing mode? *<same selectors as previous question>*

How likely are you to use browser extensions that block ads and/or online tracking? *<same selectors as previous question>*

The Do Not Track (DNT) setting is a browser setting to indicate to web pages you visit that you do not want to be tracked online. How likely are you to use the DNT setting? *<same selectors as previous question>*

How likely are you to click on ads? *<same selectors as previous question>*

Imagine that online advertising companies provided a page to show you what topics they guessed you are interested in. How likely are you to spend time looking at such a page? *<same selectors as previous question>*

IUIPC

{strongly agree, agree, somewhat agree, neutral, somewhat disagree, disagree, strongly disagree} for each

- Consumer online privacy is really a matter of consumers' right to exercise control and autonomy over decisions about how their information is collected, used, and shared.
- Consumer control of personal information lies at the heart of consumer privacy.
- I believe that online privacy is invaded when control is lost or unwillingly reduced as a result of a marketing transaction.
- Companies seeking information online should disclose the way the data are collected, processed, and used.
- A good consumer online privacy policy should have a clear and conspicuous disclosure.
- It usually bothers me when online companies ask me for personal information.
- When online companies ask me for personal information, I sometimes think twice before providing it.
- I'm concerned that online companies are collecting too much personal information about me.

Demographics

With what gender do you identify?

- Female
- Male
- Non-binary
- Other
- Prefer not to say

What is your age?

- 18-24
- 25-34
- 35-44
- 45-54
- 55-64
- 65 or older
- Prefer not to say

What is the highest degree or level of school you have completed?

- Some high school
- High school
- Some college
- Trade, technical, or vocational training
- Associate's degree
- Bachelor's degree
- Master's degree
- Professional degree
- Doctorate
- Prefer not to say

Which of the following best describes your educational background or job field?

- I have an education in, or work in, the field of computer science, engineering, or IT.
- I do not have an education in, or work in, the field of computer science, engineering, or IT.
- Prefer not to say

What is the time where you currently live right now?

(Optional) Do you have any final thoughts or questions about today's survey?

Thank you for completing Part 1 of our survey! In order to be eligible to complete *Part 2* of this study, you MUST:

- Keep our extension downloaded in your browser until you are contacted (via Prolific) to complete the second survey (in approximately one week)
- View the extension's dashboard page **at least three times** during the week by opening the Tracking Transparency dashboard (**the video below shows how**)
- If you uninstall and re-install the extension, your data will no longer be valid and payment for Part 2 will not be processed
- Please try to refrain from using private browsing mode while using the Tracking Transparency extension throughout the week

In addition to payment for completing Part 1 (this survey, $3.00), you will be compensated $7.00 for successful completion of Part 2. When you hit next, you will be redirected to Prolific. Remember, you must "Open the Tracking Transparency Dashboard" page **a total of three times** over the next week to become eligible for the second survey!

B.3 Field Study, Part II

<validate extension install, applicable visualizations, participant_ID>

Thank you for participating in Part II—**the final part**—of our study! This survey is about Tracking Transparency, the extension you installed about a week ago. The survey will take approximately 30 minutes to complete and, roughly, consists of two sections:

- Your attitudes and opinions on the Internet ecosystem
- Opinions on your data—as visualized by the Tracking Transparency extension you installed about a week ago

Part of this survey requires you to answer questions while looking at the extension's dashboard; you must take this survey using Google Chrome, on the browser where you installed the Tracking Transparency extension.

Attitudes Block Repeat

Intents Block Repeat

IUIPC Block Repeat

<for each visual in each condition (loop)>

You have this visualization in your extension. This is an example (above) of what the visualization looks like. Answer the following questions using your own extension. *{strongly agree, agree, neither agree or disagree, disagree, strongly disagree, don't see visual}*

- I find it creepy that this information is associated with me.
- I find it creepy that data brokers could sell this information to anyone who wants to pay for it.
- The information presented about me in this visualization accurately reflects me as a person.
- This visualization increases my concern about my privacy.
- I want to take privacy-protective actions based on this visualization.
- The information presented about me in this visualization accurately reflects my web browsing.

Please explain why you find this information creepy or not creepy.

<for participants in TT1-Dashboard >

(1) In your own extension, in the "Your Top Trackers" visualization, what is the #2 tracker listed?
(2) In your own extension, in the "Your Top Interests" visualization, what is the #1 interest listed?
(3) In your extension, in the "Recent Interests" and "Recent Sites" visualization, what is the first interest listed?
(4) In your extension, of the "Trackers encountered, Pages visited, and Potential interests," which of these numbers did you find most surprising?

<for participants in TT1-Everything >

(1) In your own extension, after clicking on one of the sections in the chart (marked with a star above), what "interest" is shown in the center of the circle?
(2) In your own extension, in the "Who is tracking you" visualization, which "tracker" is found on the highest percentage of pages from your browsing history? If there are none, say "none."
(3) In your extension, in the "Where were you tracked" visualization, what is the name of one of the sites listed (shown with a star in the image above) in the "Sites without trackers" list? If there are none, say "none."
(4) In your extension, when hovering over one of the dots in the "When were you tracked" visualization (noted with a star in the example image shown above), how many pages were visited at this time?

<for participants in TT2 (dependent on per-participant data)>

(1) In your own extension, which piece of information presented in the "Your demographics" visualization was most surprising to you (e.g., Age, Household Income, Marital Status, etc.)?
(2) In your extension, in the "Your inferred interests" visualization, of the "Most specific interests" listed (top-left box), what interest is listed first?
(3) In your extension, in the "Your interests over time" visualization, what is the highest number of interests ever recorded (e.g., 68 interests is the all-time high for the example image shown above)?
(4) In your extension, in the "When you're engaged" visualization, what was your "Top-Time Interest" (noted in "Slice Notes" in the bottom right-hand corner of the example image shown above)?
(5) In your extension, in the "How you spend your time" visualization, which of the listed interests did you find most surprising?
(6) In your extension, in the "When you go to sleep" visualization, hover your cursor over one of the dots; what "Late night interest" is shown?
(7) In your extension, in the "Search habits" visualization, what is one of the grouped search terms shown (e.g., "Gift" as shown in the example image above).
(8) In your extension, in the "Possible sensitive interests" visualization, list one of the categories of sensitive websites (e.g., from the example image above: politics).
(9) In your extension, in the "Ads you've been served (overview)" visualization, how many ads have you been served?
(10) In your extension, in the "Ad explanations" visualization, what is one of the ad categories shown (e.g., "Business Services" and "Living Room Furniture" are shown as ad categories in the example image provided above)?

How much would you be willing to pay—**per month**—to stop this kind of information (all of what you've seen today) from being associated with you? Please enter a number.

I am willing to take the following actions to stop this information from being associated with me:

- Stop using email (e.g., Gmail, Outlook, Yahoo)
- Install a privacy-focused browser extension which will likely slow down my Internet connection
- Use a privacy-focused browser which may slow down my Internet connection
- Stop using Google as my search engine
- Only use encrypted text messaging services (e.g., iMessage, Signal, Telegram)
- Only browse the Internet from multiple-user devices (e.g., public libraries or shared cell phones)

<show most-creepy visualization (Likert) plus accuracy response (Likert)> Do you think creepiness and accuracy are related? *<yes or no>* How are creepiness and accuracy related? *<alternative>* Why are creepiness and accuracy not related?

<show most-creepy visualization (Likert) plus privacy concern response (Likert)> Do you think creepiness and privacy concern are related? *<yes or no>* How are creepiness and privacy concern related? *<alternative>* Why are creepiness and privacy concern not related?

<show most-creepy visualization (Likert) plus willingness to act response (Likert)> Do you think creepiness and willingness to take

privacy-protective actions are related? *<yes or no>* How are creepiness and willingness to take privacy-protective actions related? *<alternative>* Why are creepiness and willingness to take privacy-protective actions not related?

(Optional) Do you have any final thoughts or questions about today's survey?

Thank you for participating in our study about opinions on the online tracking ecosystem. The purpose of this study is to: (1) show participants, visually, how tracking is occurring online; and (2) collect opinions on this tracking by asking questions about the visualizations. Thank you for your participation!

To uninstall your extension, please click the uninstall button in the popup *(instructional video below).*

C EXTENSION PRIVACY POLICY

The goal of this project is to measure and study how users interact with personalized information regarding online tracking. If any data collected pursuant to this project is sensitive, it will be anonymized. This means that any data collected will not be Personally Identifiable Information (PII). We are committed to protecting the privacy of all users of our extension. We have established this privacy policy to help explain what information we collect through the extension and how this information will be used. In this policy, "the researchers," "our," "we," or similar terms refer to any and all researchers or assistants otherwise involved in this project. This project involves personnel from the University of Maryland and the University of Chicago.

1. Information Gathered

As you browse the web, the extension will gather information. This information will be presented to you in order to provide a "tracker's perspective" of you and your browsing habits.

Information gathered by the extension may include:

1.1 Overview

- Data about the web pages you visit, including:
  - The page's title and URL
  - Date and time information about pages visited
  - The trackers present on the page
  - A guess about what the page is about (inferred topic)
  - Google adsSettings information over time AdsSettings
  - Modified (i.e., readable, or stop-words removed) web page content if that page falls into a particular inferred topic
  - Information about advertisements served, including what the ad is about (i.e., the inferred topic) and where the ad links to (i.e., the final destination click-through)
- Analytics regarding interaction with the extension
- Whether you have other ad or tracker blocker extensions installed

1.2 Sharing Data with Researchers

The extension may also share anonymous data with researchers—not PII—which includes the following: configuration of computers, operating systems, browsers, browsers' plugins, browsing patterns, adblockers and other privacy software. Although it is theoretically possible for this data to form a 'fingerprint' that could be used to track individuals, the researchers will not use the data provided for that purpose.

For clarity, here is an example of what anonymous data may look like. Please note, this example shows a particular case that only occurs when: (1) you are logged in to Google adsSettings; (2) Google adsSettings information is updated; and (3) you were visiting web pages within a three-minute window from when adsSettings was updated. If all of these conditions are met, then the extension would send the following information:

- date: 1656702538754
- account: '8d9f19fe73ba0...d172c8'
- inferences:
  - {type: 'demographic', value: '35-44 years old'}
  - {type: 'demographic', value: 'Male'}
  - {type: 'interest - company', value: 'USAA'}
- difference from previous adsSettings data:
  - {type: 'interest', value: '/Beauty & Fitness/Fitness'}
  - {type: 'interest', value: '/Home & Garden/Home Appliances'}
  - {type: 'interest', value: '/Home & Garden/Kitchen & Dining'}
- pages visited
  - "21 Best Yoga Pants For Women, According To Reviews In 2022 [Fitness, womenshealthmag.com]"
  - "Colorblock Studio Legging | Light Oregano – Vuori Clothing [Fashion & Style, googleadservices.com]"

Notably, this entry occurred on Friday, July 1, 2022 at around 3PM. Account information (i.e., 'account') is anonymized, but the inferences are not. The inferences, however, are guesses, made by Google, taken from the Google adsSettings page (https://adssettings.google.com/authenticated). Again, this level of detail only occurs if Google adsSettings is updated while you are browsing the web.

PII will exist *only* on your local copy of the extension (i.e., on the local device). If any of the data listed above is considered PII, then it will be anonymized prior to collection by us.

2. Purposes in Data Collection

Web history data: In order to help you visualize your web browsing, the extension keeps a local database with the pages that you visit while the extension is installed and enabled. While page titles and URLs are stored on the local copy of the extension (i.e., your computer), this information is never sent to the researchers. Instead, anonymized metrics will be sent in order to identify aggregate trends in web browsing and tracker activity or inferred topics.

Tracker data: This extension gathers information about the trackers that you may have interacted with online. This information is stored locally, and is also used to help you visualize what happens

when you browse online. Anonymized information about the trackers will be sent to the researchers to gain insights about online tracking, without connection to you specifically.

Inferred topics: When you browse web pages, our extension will make inferences about the topics of visited web pages and store this information locally, in order to improve the visualizations shown in the extension. Anonymized metrics about the inferred topics will be sent to the researchers to determine trends in web browsing and potential inferences, without connection to you specifically.

Google adsSettings data: As you browse the web, our extension will periodically check the Google AdsSettings web page for new information. We collect this information to improve the visualizations you see in the extension, information that is stored locally. Anonymized information about this data (e.g., number of interests or number of demographics) may be shared with the researchers, but will not include PII.

Advertising data: The extension captures information about advertisements you've been served while browsing the web. This includes the inferred interest of the advertisement, which is gathered by fetching the URL of the advertisement (i.e., the final destination of the click-through link) and guessing the topic of the resulting web page. This information is collected in order to improve visualizations found in the extension. Anonymized information about this data (e.g., number of ads or inferred ad topics or provided ad explanations) may be shared with the researchers, but will not include PII.

Usage data: We collect usage data for the dashboard visualization page in our extension. This includes data about which components were clicked, but not any identifying data about your web browsing habits. We collect this information in order to determine which parts of the dashboard are more frequently used. Usage data will not be connected to you specifically.

Other installed extensions: We access a list of your installed extensions in order to determine if you have another ad or tracker blocker installed. We do not record the specific names of any extensions you have installed, only whether there is such an extension currently enabled. This is so that we can determine whether such extensions change the behavior of the extension.

3. Updating or Removing Your Information

To protect your privacy, we use various techniques to anonymize the data, and have agreed in this policy to refrain from any attempts at re-identification of the data. Because of our use of anonymization, we will be unable to know which entry in our data set is yours. Additionally, we have no way to allow you to access, update, or remove any specific data. If you have any questions about this, please contact us at the links below.

4. Sharing of Your Data

As part of this project, we may share datasets derived from this project with research partners. Before sharing, we will evaluate whether further sanitization or aggregation of data is necessary to reduce the likelihood that inferences about identifiable individuals' activities might be made from the published dataset. Because anonymization is a complex problem, we cannot promise that our techniques will be perfect. If we find that a dataset may contain information that is sensitive or vulnerable to re-identification, we will not publish it, and if we share such data with research partners, we will place them under a contractual obligation to keep the dataset confidential and to refrain from attempts to re-identify.

Furthermore, we may publicly release and publish anonymized information from datasets to further general scientific knowledge. The datasets we may share or publish will not intentionally contain PII. As part of the surveys for this project, you will be asked whether you are willing to allow anonymized data from your responses to be publicly released for scientific purposes. This decision will not affect your participation or compensation in any way.

5. Data Storage and Retention

We will retain the dataset for as long as the data remains useful for research topics related to online tracking, privacy, and personalized web visualizations.

6. Security

We employ industry-standard security measures to protect the loss, misuse, and alteration of the information under our control, including appropriate technical and organizational measures to ensure a level of security appropriate to the risk, such as the pseudonymization, the encryption of personal data, data backup systems, and engaging security professionals to evaluate our system's effectiveness. Although we make good faith efforts to store information collected by us in a secure operating environment, we cannot guarantee complete security.

7. Contact

If you have any questions about our privacy and data protection practices, you can reach our Principal Investigators at:

[contact information appeared here]

8. Changes Made

This privacy policy may change periodically. However, any revised privacy policy will be consistent with the purposes of this research project. If we make any substantive changes to our policies, we will post notice of changes on this page.

- Updated October 28, 2018 to clarify affiliations of researchers.
- Updated November 5, 2018 to update institutions involved.
- Updated January 19, 2022 regarding data collection processes.
- Updated July 6, 2022 to add examples of data collection.
- Updated August 2, 2023 to reflect accurate data collection.

D ADDITIONAL FIGURES

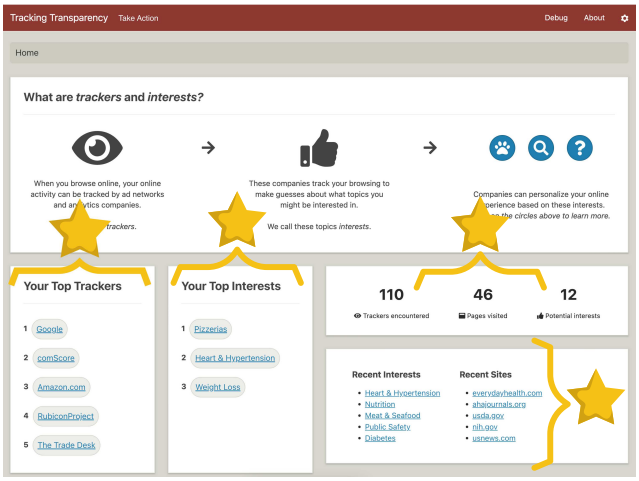

Figure 9: TT1-Dashboard. Stars show locations where the participant was prompted to verify they were responding to the correct visualization (e.g., "[i]n your own extension, in the 'Your Top Trackers' visualization, what is the #2 tracker listed?"). Participants in this condition answered one set of Creep Factor questions per informational box (i.e., Interests, Trackers, Sites, and Activity).

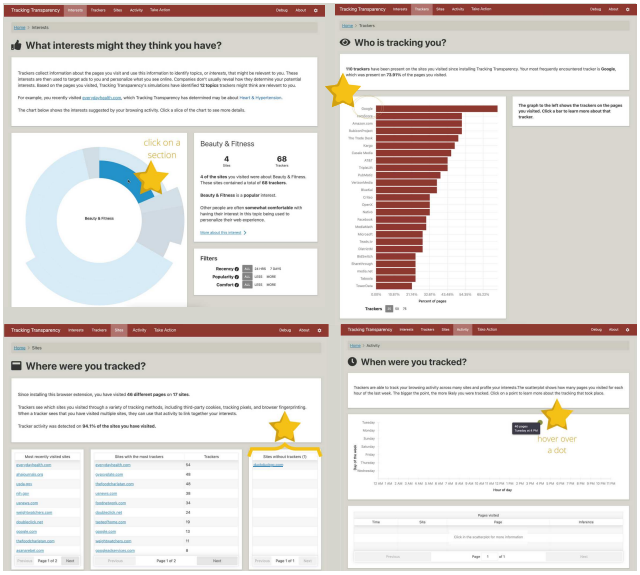

Figure 10: TT1-Everything example. Stars as in Figure 9. Participants in this condition answered one set of Creep Factor questions per tab (i.e., Interests, Trackers, Sites, and Activity).

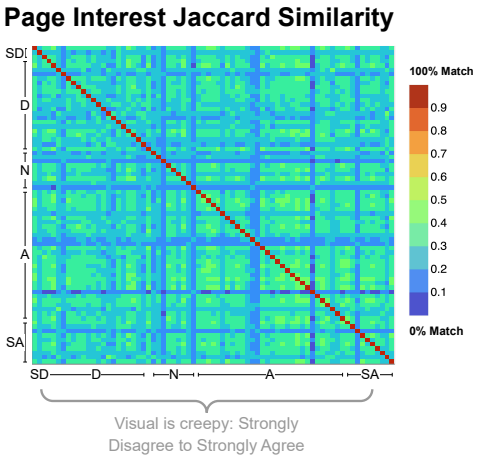

Figure 11: Jaccard similarity for pairwise comparisons between participants' inferred web page interests, grouped by Likert responses to the *General Creepiness* question for the *Time Spent per Interest* visualization. Particular interests do not seem to correlate with creepiness.

E INTEREST INFERENCE ENGINE DETAIL

Table 6: Model performance per type. For a full breakdown of performance on all categories included in the final model, see Table 7.

	Reduced Text Accuracy (%)		
	Train	Test	No. Param.
BoW + SLP	99.8	48.9	15,699,191
TF-IDF + SLP	33.0	24.1	15,280,536
Word2Vec + 1 LSTM + 1FC	37.7	18.4	2,772,859
GloVe + 1 LSTM + 1FC	59.9	29.0	2,772,859
GloVe + 3 Bidirect. LSTM + 2FC	50.0	33.3	24,935,727

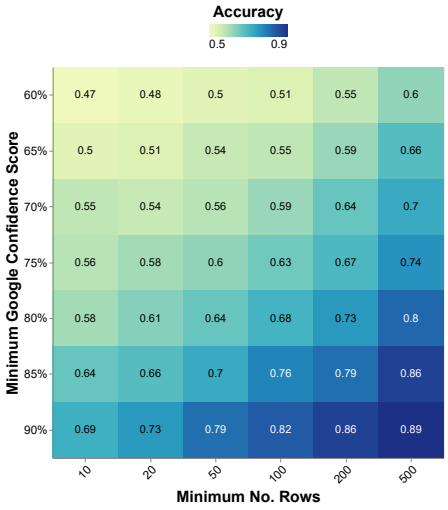

Figure 12: Heatmap of experimental results using different parameters, varying the minimum value of Google confidence score and the minimum number of rows (i.e., examples) per category.

Table 7: Test performance metrics per interest category, listing only categories found in the final model (i.e., some categories found in [43] may be missing from this list).

	Testing Performance (per label)			
	precision	recall	f1-score	support
Adult	0.94	0.96	0.95	106
Arts & Entertainment	0.50	0.12	0.20	8
|----Celebrities & Entertainment News	1.00	0.67	0.80	3
|----Comics & Animation	1.00	1.00	1.00	3
|----|----Anime & Manga	0.83	1.00	0.91	5
|----|----Cartoons	0.00	0.00	0.00	2
|----|----Comics	1.00	0.50	0.67	8
|----|----Film & TV Industry	0.00	0.00	0.00	4
|----|----Recording Industry	0.83	0.83	0.83	6
|----Events & Listings	0.78	0.88	0.82	8
|----|----Bars, Clubs & Nightlife	0.80	0.50	0.62	8
|----|----Concerts & Music Festivals	0.00	0.00	0.00	1
|----|----Expos & Conventions	0.58	1.00	0.74	7
|----|----Film Festivals	1.00	0.67	0.80	3
|----Fun & Trivia	1.00	1.00	1.00	1
|----Humor	0.61	0.92	0.73	12
|----|----Funny Pictures & Videos	0.83	0.83	0.83	6
|----Movies	0.65	0.83	0.73	41
|----Music & Audio	0.56	0.58	0.57	24
|----|----Classical Music	0.79	0.92	0.85	12
|----|----Country Music	0.25	1.00	0.40	1
|----|----Dance & Electronic Music	1.00	0.50	0.67	14
|----|----Jazz & Blues	1.00	0.67	0.80	6
|----|----Music Education & Instruction	0.90	0.60	0.72	15
|----|----Music Equipment & Technology	0.86	0.92	0.89	13
|----|----Music Reference	0.00	0.00	0.00	1
|----|----Radio	0.91	0.91	0.91	11
|----|----Religious Music	1.00	0.33	0.50	3
|----|----Rock Music	0.57	0.80	0.67	5
|----|----Urban & Hip-Hop	0.83	0.71	0.77	14
|----|----World Music	1.00	0.50	0.67	2
|----Online Media	0.60	0.57	0.59	21
|----|----Online Image Galleries	0.50	1.00	0.67	1
|----Performing Arts	0.57	0.33	0.42	12
|----|----Acting & Theater	0.64	0.70	0.67	10
|----|----Circus	1.00	0.20	0.33	5
|----|----Dance	0.75	0.60	0.67	5
|----|----Magic	1.00	0.50	0.67	2
|----TV & Video	0.00	0.00	0.00	2
|----|----Online Video	0.20	0.33	0.25	3
|----|----TV Shows & Programs	0.78	0.82	0.80	17
|----Visual Art & Design	0.40	0.40	0.40	5
|----|----Architecture	0.92	0.79	0.85	14
|----|----Art Museums & Galleries	0.00	0.00	0.00	1
|----|----Design	0.56	0.45	0.50	11
|----|----Painting	0.71	0.62	0.67	8
|----|----Photographic & Digital Arts	0.59	0.59	0.59	17
Autos & Vehicles	0.44	0.70	0.54	10
|----Bicycles & Accessories	0.57	0.72	0.63	18
|----Boats & Watercraft	0.75	0.60	0.67	5
|----Campers & RVs	0.82	0.82	0.82	11
|----Classic Vehicles	1.00	0.67	0.80	3
|----|----Cargo Trucks & Trailers	0.00	0.00	0.00	3
|----Motor Vehicles (By Type)	0.00	0.00	0.00	1
|----|----Hybrid & Alternative Vehicles	1.00	0.90	0.95	10
|----|----Motorcycles	0.75	0.60	0.67	10
|----|----Off-Road Vehicles	0.00	0.00	0.00	3
|----|----Trucks & SUVs	0.86	0.86	0.86	7
|----Vehicle Codes & Driving Laws	0.67	0.86	0.75	7
|----|----Vehicle Licensing & Registration	0.74	0.93	0.82	15
|----Vehicle Parts & Services	0.00	0.00	0.00	1
|----|----Vehicle Parts & Accessories	0.57	0.44	0.50	9
|----|----Vehicle Repair & Maintenance	0.60	0.90	0.72	10
|----Vehicle Shopping	0.63	0.81	0.71	21
|----|----Used Vehicles	1.00	0.29	0.44	7
|----Vehicle Shows	0.00	0.00	0.00	3
Beauty & Fitness	0.33	0.29	0.31	7
|----Beauty Pageants	1.00	0.67	0.80	3
|----Body Art	1.00	0.71	0.83	7
|----Cosmetic Procedures	0.60	0.43	0.50	7
|----|----Cosmetic Surgery	0.30	0.30	0.30	10
|----Cosmetology & Beauty Professionals	0.62	0.71	0.67	7
|----Face & Body Care	0.67	0.55	0.60	11
|----|----Hygiene & Toiletries	0.00	0.00	0.00	2
|----|----Make-Up & Cosmetics	0.78	0.86	0.82	21
|----|----Perfumes & Fragrances	0.75	0.75	0.75	4
|----|----Skin & Nail Care	0.80	0.75	0.77	16
|----|----Unwanted Body & Facial Hair Removal	1.00	0.33	0.50	3
|----Fashion & Style	0.82	0.90	0.86	10
|----|----Fashion Designers & Collections	1.00	0.75	0.86	4
|----Fitness	0.60	0.75	0.67	20
|----Hair Care	0.76	0.84	0.80	19
|----|----Hair Loss	0.75	0.60	0.67	5
|----Spas & Beauty Services	0.80	1.00	0.89	4
|----|----Massage Therapy	1.00	0.60	0.75	5
|----Weight Loss	0.38	0.50	0.43	6
Books & Literature	0.64	0.73	0.68	22
|----Children's Literature	1.00	0.73	0.84	11
|----E-Books	1.00	1.00	1.00	4
|----Fan Fiction	0.80	1.00	0.89	4

Continued on next column

Continued from previous column

	precision	recall	f1-score	support
|----Literary Classics	0.82	0.75	0.78	12
|----Poetry	0.80	0.80	0.80	5
|----Writers Resources	0.95	0.95	0.95	19
Business & Industrial	0.35	0.50	0.41	58
|----|----Public Relations	0.50	0.25	0.33	4
|----|----Space Technology	0.86	0.86	0.86	14
|----Agriculture & Forestry	0.71	0.50	0.59	10
|----|----Agricultural Equipment	1.00	1.00	1.00	3
|----|----Forestry	0.50	0.50	0.50	2
|----|----Livestock	1.00	0.50	0.67	2
|----Business Education	0.75	0.64	0.69	14
|----Business Finance	0.50	0.33	0.40	3
|----|----Venture Capital	0.67	1.00	0.80	2
|----Business Operations	0.73	0.76	0.74	21
|----|----Business Plans & Presentations	0.00	0.00	0.00	1
|----|----Management	0.77	0.62	0.69	16
|----Business Services	0.72	0.79	0.75	107
|----|----Corporate Events	0.50	0.50	0.50	2
|----|----E-Commerce Services	0.74	0.74	0.74	34
|----|----Fire & Security Services	0.00	0.00	0.00	1
|----|----Office Supplies	0.71	0.83	0.77	12
|----|----Writing & Editing Services	0.75	0.50	0.60	6
|----Chemicals Industry	0.72	0.76	0.74	17
|----|----Plastics & Polymers	1.00	0.58	0.74	12
|----Construction & Maintenance	0.53	0.75	0.62	12
|----|----Building Materials & Supplies	0.53	0.56	0.55	16
|----Energy & Utilities	0.80	0.73	0.76	11
|----|----Electricity	0.00	0.00	0.00	2
|----|----Oil & Gas	0.88	0.75	0.81	20
|----|----Renewable & Alternative Energy	0.67	1.00	0.80	10
|----Hospitality Industry	0.50	0.50	0.50	4
|----|----Event Planning	0.00	0.00	0.00	1
|----Industrial Materials & Equipment	1.00	1.00	1.00	3
|----|----Heavy Machinery	0.80	0.67	0.73	6
|----Manufacturing	0.00	0.00	0.00	1
|----Metals & Mining	0.75	0.43	0.55	7
|----Pharmaceuticals & Biotech	0.83	1.00	0.91	5
|----Printing & Publishing	0.50	0.40	0.44	5
|----|----Retail Equipment & Technology	1.00	0.50	0.67	2
|----|----MLM & Business Opportunities	0.58	0.78	0.67	9
|----Textiles & Nonwovens	1.00	0.83	0.91	6
|----Transportation & Logistics	0.65	0.82	0.72	38
|----|----Freight & Trucking	0.77	0.59	0.67	17
|----|----Mail & Package Delivery	0.83	0.56	0.67	9
|----|----Maritime Transport	1.00	0.50	0.67	2
|----|----Moving & Relocation	0.75	0.86	0.80	7
|----|----Packaging	0.00	0.00	0.00	1
|----|----Parking	0.82	0.60	0.69	15
|----|----Rail Transport	1.00	0.20	0.33	5
|----|----Urban Transport	0.00	0.00	0.00	2
Computers & Electronics	0.20	0.08	0.11	13
|----CAD & CAM	0.33	0.33	0.33	6
|----Computer Hardware	0.50	0.33	0.40	6
|----|----Computer Components	1.00	0.78	0.88	9
|----|----Computer Drives & Storage	0.89	0.89	0.89	18
|----|----Computer Peripherals	1.00	0.50	0.67	6
|----|----Desktop Computers	0.00	0.00	0.00	1
|----|----Laptops & Notebooks	0.91	0.83	0.87	12
|----Computer Security	0.79	0.89	0.83	54
|----|----Hacking & Cracking	0.00	0.00	0.00	1
|----Consumer Electronics	0.56	0.64	0.60	14
|----|----Audio Equipment	0.73	0.80	0.76	10
|----|----Camera & Photo Equipment	0.68	0.89	0.77	19
|----|----Drones & RC Aircraft	1.00	0.33	0.50	3
|----|----GPS & Navigation	1.00	0.33	0.50	3
|----|----Game Systems & Consoles	1.00	0.70	0.82	10
|----|----TV & Video Equipment	0.83	0.77	0.80	13
|----Electronics & Electrical	1.00	0.67	0.80	3
|----|----Electronic Components	0.50	1.00	0.67	2
|----|----Power Supplies	1.00	0.33	0.50	3
|----Enterprise Technology	0.33	0.17	0.22	6
|----|----Data Management	0.83	0.62	0.71	8
|----Networking	0.77	0.77	0.77	13
|----|----Data Formats & Protocols	0.00	0.00	0.00	2
|----|----Network Monitoring & Management	0.88	0.64	0.74	11
|----|----VPN & Remote Access	0.71	1.00	0.83	10
|----Programming	0.66	0.65	0.65	86
|----|----Java (Programming Language)	0.50	0.33	0.40	3
|----Software	0.29	0.36	0.32	11
|----|----Business & Productivity Software	0.89	0.73	0.80	22
|----|----Device Drivers	0.89	0.89	0.89	9
|----|----Internet Software	0.50	0.36	0.42	11
|----|----Multimedia Software	0.71	0.81	0.76	27
|----|----Operating Systems	0.80	0.36	0.50	11
Finance	0.33	0.20	0.25	5
|----Accounting & Auditing	0.33	0.33	0.33	3
|----|----Tax Preparation & Planning	0.88	0.88	0.88	8
|----Banking	0.67	0.92	0.77	13
|----Credit & Lending	0.80	0.57	0.67	7
|----|----Credit Cards	1.00	0.57	0.73	7
|----|----Credit Reporting & Monitoring	1.00	1.00	1.00	4
|----|----Loans	0.90	0.96	0.93	46
|----Financial Planning & Management	0.88	0.70	0.78	10
|----|----Retirement & Pension	0.82	1.00	0.90	14
|----Grants, Scholarships & Financial Aid	0.83	0.96	0.89	25
|----|----Study Grants & Scholarships	1.00	0.25	0.40	4
|----Insurance	0.82	0.96	0.88	24

Continued on next column

Continued from previous column

	precision	recall	f1-score	support
|----|----Health Insurance	0.80	0.80	0.80	5
|----Investing	0.78	0.83	0.81	30
|----|----Commodities & Futures Trading	1.00	0.67	0.80	3
|----|----Currencies & Foreign Exchange	0.88	0.85	0.86	33
Food & Drink	0.57	0.44	0.50	9
|----Beverages	0.80	1.00	0.89	4
|----|----Alcoholic Beverages	0.97	0.86	0.91	36
|----|----Coffee & Tea	0.71	0.91	0.80	11
|----|----Juice	0.00	0.00	0.00	3
|----Cooking & Recipes	0.60	0.88	0.71	24
|----|----BBQ & Grilling	0.00	0.00	0.00	1
|----|----Desserts	0.62	0.50	0.56	10
|----|----Soups & Stews	0.00	0.00	0.00	1
|----Food	0.20	0.33	0.25	3
|----Food & Grocery Retailers	0.80	0.80	0.80	5
|----|----Baked Goods	0.00	0.00	0.00	2
|----|----Breakfast Foods	1.00	1.00	1.00	1
|----|----Candy & Sweets	0.83	0.50	0.62	10
|----|----Grains & Pasta	0.00	0.00	0.00	1
|----|----Meat & Seafood	1.00	0.38	0.55	8
|----|----Snack Foods	1.00	0.67	0.80	6
|----Restaurants	0.73	0.73	0.73	11
|----|----Pizzerias	1.00	1.00	1.00	8
Games	0.25	0.67	0.36	3
|----Arcade & Coin-Op Games	0.80	1.00	0.89	8
|----Board Games	1.00	0.33	0.50	3
|----|----Chess & Abstract Strategy Games	1.00	0.78	0.88	9
|----|----Miniatures & Wargaming	1.00	0.77	0.87	13
|----Card Games	1.00	0.80	0.89	5
|----|----Collectible Card Games	1.00	0.88	0.93	8
|----|----Poker & Casino Games	0.88	0.94	0.91	16
|----Computer & Video Games	0.54	0.79	0.64	19
|----|----Casual Games	0.00	0.00	0.00	1
|----|----Driving & Racing Games	1.00	0.80	0.89	5
|----|----Fighting Games	0.00	0.00	0.00	4
|----|----Music & Dance Games	1.00	1.00	1.00	1
|----|----Sandbox Games	0.00	0.00	0.00	1
|----|----Shooter Games	1.00	0.67	0.80	9
|----|----Sports Games	0.00	0.00	0.00	1
|----|----Strategy Games	1.00	0.71	0.83	7
|----|----Video Game Emulation	1.00	1.00	1.00	1
|----|----Drawing & Coloring	0.00	0.00	0.00	1
|----Gambling	0.82	0.90	0.86	10
|----|----Lottery	1.00	0.86	0.92	7
|----|----Massively Multiplayer Games	0.00	0.00	0.00	5
|----Puzzles & Brainteasers	0.75	1.00	0.86	12
|----Roleplaying Games	0.46	0.69	0.55	16
|----Table Games	1.00	0.20	0.33	5
|----|----Billiards	0.93	0.93	0.93	15
|----Word Games	1.00	1.00	1.00	1
Health	0.31	0.38	0.34	13
|----Aging & Geriatrics	0.71	0.62	0.67	8
|----Health Conditions	0.48	0.50	0.49	26
|----|----AIDS & HIV	0.92	1.00	0.96	12
|----|----Allergies	0.77	1.00	0.87	10
|----|----Arthritis	1.00	1.00	1.00	6
|----|----Cancer	0.69	1.00	0.81	11
|----|----Diabetes	0.76	0.76	0.76	17
|----|----Ear Nose & Throat	1.00	0.77	0.87	13
|----|----Eating Disorders	1.00	0.80	0.89	10
|----|----Endocrine Conditions	0.62	0.71	0.67	7
|----|----Genetic Disorders	0.80	0.57	0.67	7
|----|----Heart & Hypertension	0.85	0.88	0.87	26
|----|----Infectious Diseases	0.80	0.50	0.62	8
|----|----Neurological Conditions	0.50	0.25	0.33	8
|----|----Obesity	0.90	0.82	0.86	11
|----|----Pain Management	0.80	0.67	0.73	12
|----|----Respiratory Conditions	0.91	0.77	0.83	13
|----|----Skin Conditions	0.83	0.71	0.77	14
|----|----Sleep Disorders	0.75	1.00	0.86	12
|----Health Education & Medical Training	0.62	0.57	0.59	23
|----Health Foundations & Medical Research	0.71	0.71	0.71	7
|----Medical Devices & Equipment	0.50	0.20	0.29	5
|----Medical Facilities & Services	1.00	0.29	0.44	7
|----|----Hospitals & Treatment Centers	0.64	0.82	0.72	11
|----|----Medical Procedures	0.66	0.59	0.62	49
|----|----Physical Therapy	0.86	0.67	0.75	9
|----Men's Health	1.00	0.67	0.80	3
|----Mental Health	0.50	0.93	0.65	14
|----|----Anxiety & Stress	0.82	0.82	0.82	11
|----|----Depression	0.60	0.50	0.55	6
|----Nursing	0.76	0.73	0.74	22
|----|----Assisted Living & Long Term Care	0.80	0.67	0.73	6
|----Nutrition	0.90	0.69	0.78	13
|----|----Special & Restricted Diets	1.00	0.50	0.67	2
|----|----Vitamins & Supplements	0.58	0.64	0.61	11
|----Oral & Dental Care	0.88	0.71	0.79	21
|----Pharmacy	1.00	0.64	0.78	11
|----|----Drugs & Medications	0.50	0.25	0.33	4
|----Public Health	0.29	0.33	0.31	6
|----|----Occupational Health & Safety	0.50	0.25	0.33	8
|----Reproductive Health	0.33	0.29	0.31	7
|----Substance Abuse	0.00	0.00	0.00	2
|----|----Drug & Alcohol Treatment	0.89	0.89	0.89	9
|----|----Smoking & Smoking Cessation	0.86	0.86	0.86	14
|----|----Steroids & Performance-Enhancing Drugs	1.00	0.80	0.89	5
|----Vision Care	0.86	1.00	0.92	6

Continued on next column

Continued from previous column

	precision	recall	f1-score	support
|----Women's Health	0.84	0.78	0.81	27
Hobbies & Leisure	0.54	0.60	0.57	70
|----Clubs & Organizations	0.00	0.00	0.00	1
|----|----Youth Organizations & Resources	1.00	0.75	0.86	4
|----Crafts	0.60	0.60	0.60	5
|----|----Fiber & Textile Arts	0.88	0.88	0.88	8
|----Merit Prizes & Contests	0.33	0.40	0.36	5
|----Outdoors	0.50	0.33	0.40	6
|----|----Fishing	1.00	1.00	1.00	8
|----|----Hiking & Camping	0.88	0.83	0.86	18
|----Paintball	1.00	0.60	0.75	5
|----Radio Control & Modeling	0.83	0.62	0.71	8
|----|----Model Trains & Railroads	0.88	1.00	0.93	7
|----Special Occasions	0.27	0.38	0.32	8
|----|----Holidays & Seasonal Events	0.20	0.33	0.25	6
|----|----Weddings	0.20	0.10	0.13	10
|----Water Activities	0.83	0.42	0.56	12
|----|----Boating	0.90	0.82	0.86	11
|----|----Surf & Swim	0.77	0.94	0.85	18
Home & Garden	0.43	0.38	0.40	8
|----Bed & Bath	0.75	0.75	0.75	8
|----|----Bathroom	0.80	0.80	0.80	10
|----|----Cleaning Services	0.59	1.00	0.74	16
|----Gardening & Landscaping	0.76	0.87	0.81	15
|----HVAC & Climate Control	0.66	0.68	0.67	31
|----|----Fireplaces & Stoves	0.00	0.00	0.00	1
|----Home & Interior Decor	0.57	0.57	0.57	7
|----Home Appliances	0.56	1.00	0.71	5
|----Home Furnishings	0.55	0.71	0.62	24
|----|----Curtains & Window Treatments	0.86	0.75	0.80	8
|----|----Lamps & Lighting	0.73	0.85	0.79	13
|----|----Living Room Furniture	0.60	0.75	0.67	4
|----|----Rugs & Carpets	0.72	0.72	0.72	18
|----Home Improvement	0.41	0.57	0.48	21
|----|----Construction & Power Tools	1.00	0.82	0.90	11
|----|----Doors & Windows	0.75	0.60	0.67	10
|----|----Flooring	0.83	1.00	0.91	5
|----|----House Painting & Finishing	0.75	1.00	0.86	3
|----|----Plumbing	0.86	0.86	0.86	7
|----Home Safety & Security	0.00	0.00	0.00	0
|----Home Storage & Shelving	0.75	0.86	0.80	7
|----Home Swimming Pools, Saunas & Spas	0.96	1.00	0.98	27
|----Kitchen & Dining	0.50	0.33	0.40	3
|----|----Cookware & Diningware	1.00	0.50	0.67	2
|----|----Major Kitchen Appliances	0.50	0.25	0.33	4
|----|----Small Kitchen Appliances	0.75	0.60	0.67	5
|----Laundry	0.00	0.00	0.00	1
|----|----Washers & Dryers	0.00	0.00	0.00	1
|----Pest Control	1.00	0.75	0.86	8
|----Yard & Patio	0.88	0.67	0.76	21
|----|----Lawn Mowers	1.00	1.00	1.00	1
Internet & Telecom	0.00	0.00	0.00	1
|----|----Radio Equipment	1.00	0.50	0.67	4
|----Email & Messaging	0.80	0.80	0.80	5
|----|----Voice & Video Chat	0.60	0.75	0.67	4
|----|----Mobile & Wireless Accessories	0.67	0.40	0.50	5
|----|----Mobile Apps & Add-Ons	0.42	0.53	0.47	15
|----|----Mobile Phones	0.57	0.73	0.64	11
|----Service Providers	0.87	0.93	0.90	14
|----|----Cable & Satellite Providers	1.00	1.00	1.00	4
|----Web Services	0.65	0.70	0.67	60
|----|----Affiliate Programs	0.75	0.75	0.75	4
|----|----Web Design & Development	0.66	0.70	0.68	30
Jobs & Education	0.00	0.00	0.00	1
|----Education	0.43	0.63	0.51	51
|----|----Colleges & Universities	0.70	0.81	0.75	57
|----|----Distance Learning	0.75	1.00	0.86	3
|----|----Homeschooling	0.67	1.00	0.80	4
|----|----Primary & Secondary Schooling (K-12)	0.70	0.79	0.75	24
|----|----Standardized & Admissions Tests	0.55	0.67	0.60	18
|----|----Teaching & Classroom Resources	1.00	0.86	0.92	7
|----|----Training & Certification	0.33	0.33	0.33	3
|----|----Vocational & Continuing Education	0.00	0.00	0.00	4
|----Jobs	0.00	0.00	0.00	2
|----|----Career Resources & Planning	0.67	1.00	0.80	2
|----|----Job Listings	0.77	0.94	0.85	18
|----|----Resumes & Portfolios	0.88	1.00	0.93	14
Law & Government	0.00	0.00	0.00	1
|----Government	0.50	0.25	0.33	8
|----|----Courts & Judiciary	0.88	1.00	0.93	7
|----|----Visa & Immigration	0.88	0.88	0.88	8
|----Legal	0.69	0.69	0.69	26
|----|----Bankruptcy	1.00	0.83	0.91	6
|----|----Legal Education	1.00	0.82	0.90	11
|----|----Legal Services	1.00	1.00	1.00	2
|----Military	0.74	0.85	0.79	20
|----Public Safety	0.58	0.76	0.66	41
|----|----Crime & Justice	0.80	0.44	0.57	9
|----|----Emergency Services	0.67	0.57	0.62	7
|----|----Law Enforcement	0.71	0.56	0.63	9
|----|----Security Products & Services	0.89	0.76	0.82	21
|----Social Services	1.00	0.62	0.77	8
News	0.50	0.42	0.45	12
|----|----Company News	0.00	0.00	0.00	1
|----Politics	0.85	0.93	0.89	44
|----Sports News	0.00	0.00	0.00	0
|----Weather	1.00	0.91	0.95	11

Continued on next column

Continued from previous column

	precision	recall	f1-score	support
Online Communities	0.80	0.67	0.73	6
|———-Blogging Resources & Services	0.65	0.83	0.73	18
|———-Dating & Personals	0.94	1.00	0.97	48
|———-|———-Matrimonial Services	1.00	0.33	0.50	3
|———-File Sharing & Hosting	1.00	1.00	1.00	1
|———-|———-Clip Art & Animated GIFs	1.00	0.80	0.89	5
|———-|———-Skins, Themes & Wallpapers	1.00	0.55	0.71	11
|———-Social Networks	0.38	1.00	0.55	3
|———-Virtual Worlds	0.00	0.00	0.00	2
People & Society	0.56	0.42	0.48	12
|———-Family & Relationships	0.00	0.00	0.00	2
|———-|———-Family	0.57	0.62	0.59	13
|———-|———-Marriage	0.65	0.78	0.71	36
|———-|———-Troubled Relationships	0.00	0.00	0.00	1
|———-|———-Children's Interests	1.00	0.43	0.60	7
|———-Religion & Belief	0.78	0.86	0.82	114
|———-Seniors & Retirement	1.00	0.40	0.57	5
|———-Social Issues & Advocacy	0.38	0.43	0.40	21
|———-|———-Charity & Philanthropy	0.67	0.89	0.76	27
|———-|———-Discrimination & Identity Relations	0.00	0.00	0.00	1
|———-|———-Green Living & Environmental Issues	0.00	0.00	0.00	3
|———-|———-Poverty & Hunger	1.00	0.33	0.50	3
|———-|———-Work & Labor Issues	0.90	0.82	0.86	11
|———-Social Sciences	0.59	0.59	0.59	17
|———-|———-Political Science	0.00	0.00	0.00	1
|———-|———-Psychology	0.54	0.88	0.67	8
|———-Subcultures & Niche Interests	0.00	0.00	0.00	2
Pets & Animals	0.68	0.64	0.66	53
|———-|———-Pet Food & Supplies	0.89	0.53	0.67	15
|———-|———-Veterinarians	0.86	0.96	0.91	70
|———-Pets	0.00	0.00	0.00	2
|———-|———-Birds	1.00	1.00	1.00	1
|———-|———-Cats	1.00	0.57	0.73	7
|———-|———-Dogs	0.57	0.44	0.50	9
|———-|———-Exotic Pets	0.00	0.00	0.00	1
|———-|———-Fish & Aquaria	0.83	1.00	0.91	10
|———-|———-Horses	0.75	0.75	0.75	4
|———-|———-Rabbits & Rodents	0.33	0.50	0.40	2
|———-|———-Reptiles & Amphibians	0.75	0.60	0.67	5
|———-Wildlife	0.54	0.64	0.58	11
Real Estate	0.68	0.68	0.68	19
|———-Real Estate Listings	0.33	0.40	0.36	10
|———-|———-Bank-Owned & Foreclosed Properties	1.00	0.50	0.67	8
|———-|———-Commercial Properties	1.00	1.00	1.00	2
|———-|———-Lots & Land	0.50	0.33	0.40	3
|———-|———-Residential Rentals	1.00	1.00	1.00	14
|———-|———-Residential Sales	0.69	0.75	0.72	12
|———-|———-Timeshares & Vacation Properties	1.00	1.00	1.00	5
|———-Real Estate Services	0.73	0.57	0.64	14
Reference	0.00	0.00	0.00	5
|———-|———-Business & Personal Listings	0.67	0.67	0.67	3
|———-|———-Biographies & Quotations	0.00	0.00	0.00	2
|———-|———-Calculators & Reference Tools	1.00	0.50	0.67	2
|———-|———-Dictionaries & Encyclopedias	0.80	0.80	0.80	10
|———-|———-Forms Guides & Templates	1.00	1.00	1.00	1
|———-|———-Public Records	0.43	1.00	0.60	6
|———-|———-Time & Calendars	0.75	0.50	0.60	6
|———-|———-Maps	0.75	0.75	0.75	12
|———-|———-History	0.83	0.83	0.83	6
|———-|———-Myth & Folklore	1.00	0.25	0.40	4
|———-|———-Philosophy	0.92	0.92	0.92	12
|———-Language Resources	1.00	0.88	0.93	8
|———-|———-Foreign Language Resources	0.80	0.57	0.67	7
|———-Libraries & Museums	0.70	1.00	0.82	23
|———-|———-Museums	0.48	0.73	0.58	15
Science	1.00	0.17	0.29	6
|———-Astronomy	1.00	0.67	0.80	9
|———-Biological Sciences	0.67	0.80	0.73	5
|———-|———-Neuroscience	0.50	1.00	0.67	1
|———-Chemistry	0.87	1.00	0.93	13
|———-Computer Science	0.40	0.49	0.44	39
|———-Earth Sciences	0.72	0.72	0.72	18
|———-|———-Atmospheric Science	0.75	0.75	0.75	4
|———-|———-Geology	1.00	0.75	0.86	12
|———-Ecology & Environment	0.00	0.00	0.00	2
|———-|———-Climate Change & Global Warming	0.00	0.00	0.00	3
|———-Engineering & Technology	0.00	0.00	0.00	4
|———-|———-Robotics	0.88	1.00	0.93	7
|———-Mathematics	0.80	0.80	0.80	5
|———-|———-Statistics	1.00	0.71	0.83	7
|———-Physics	0.86	0.75	0.80	16
Sensitive Subjects	0.68	0.67	0.67	63
Shopping	0.10	0.20	0.13	10
|———-Antiques & Collectibles	0.73	0.89	0.80	9
|———-Apparel	0.50	0.67	0.57	24
|———-|———-Athletic Apparel	0.00	0.00	0.00	2
|———-|———-Casual Apparel	0.00	0.00	0.00	3
|———-|———-Children's Clothing	1.00	0.86	0.92	7
|———-|———-Clothing Accessories	0.91	0.89	0.90	36
|———-|———-Costumes	1.00	0.80	0.89	5
|———-|———-Eyewear	0.81	1.00	0.89	17
|———-|———-Footwear	0.84	0.87	0.86	31
|———-|———-Formal Wear	0.00	0.00	0.00	1
|———-|———-Headwear	1.00	0.25	0.40	4
|———-|———-Men's Clothing	1.00	0.20	0.33	5
|———-|———-Swimwear	0.88	0.78	0.82	9
|———-|———-Undergarments	1.00	0.60	0.75	10

Continued on next column

Continued from previous column

	precision	recall	f1-score	support
|———-|———-Women's Clothing	0.75	0.43	0.55	7
|———-Auctions	0.57	0.67	0.62	6
|———-Classifieds	0.00	0.00	0.00	1
|———-Consumer Resources	0.00	0.00	0.00	1
|———-|———-Consumer Advocacy & Protection	1.00	1.00	1.00	1
|———-|———-Coupons & Discount Offers	0.75	0.50	0.60	6
|———-Gifts & Special Event Items	1.00	0.67	0.80	3
|———-|———-Cards & Greetings	0.67	1.00	0.80	2
|———-|———-Flowers	0.88	0.92	0.90	24
|———-|———-Gifts	0.67	0.50	0.57	4
|———-Photo & Video Services	1.00	0.50	0.67	6
|———-Tobacco Products	0.50	0.14	0.22	7
|———-Toys	1.00	0.50	0.67	4
|———-|———-Building Toys	1.00	1.00	1.00	1
Sports	0.53	0.59	0.56	29
|———-Animal Sports	0.50	0.67	0.57	3
|———-College Sports	0.67	0.90	0.77	31
|———-Combat Sports	0.00	0.00	0.00	5
|———-|———-Boxing	0.88	0.70	0.78	10
|———-|———-Martial Arts	0.91	0.91	0.91	32
|———-|———-Wrestling	0.75	0.75	0.75	8
|———-Extreme Sports	1.00	0.67	0.80	6
|———-Fantasy Sports	0.88	0.88	0.88	8
|———-Individual Sports	0.40	0.33	0.36	6
|———-|———-Cycling	0.83	0.71	0.77	7
|———-|———-Golf	0.90	0.75	0.82	12
|———-|———-Gymnastics	1.00	0.57	0.73	7
|———-|———-Racquet Sports	0.92	0.92	0.92	13
|———-|———-Skate Sports	1.00	0.50	0.67	4
|———-|———-Track & Field	0.67	0.67	0.67	3
|———-|———-Olympics	0.50	0.80	0.62	5
|———-Motor Sports	1.00	0.50	0.67	6
|———-Sporting Goods	0.43	0.33	0.38	18
|———-|———-Sports Memorabilia	1.00	0.83	0.91	6
|———-|———-Winter Sports Equipment	1.00	0.17	0.29	6
|———-Sports Coaching & Training	0.00	0.00	0.00	1
|———-Team Sports	0.00	0.00	0.00	5
|———-|———-American Football	0.86	0.80	0.83	15
|———-|———-Baseball	0.60	0.43	0.50	7
|———-|———-Basketball	0.75	0.43	0.55	14
|———-|———-Cheerleading	1.00	0.88	0.93	8
|———-|———-Cricket	1.00	0.80	0.89	5
|———-|———-Hockey	0.86	0.95	0.90	19
|———-|———-Rugby	0.67	0.67	0.67	3
|———-|———-Soccer	0.80	1.00	0.89	4
|———-|———-Volleyball	1.00	0.86	0.92	7
|———-Winter Sports	0.00	0.00	0.00	3
|———-|———-Ice Skating	1.00	0.83	0.91	6
|———-|———-Skiing & Snowboarding	0.57	1.00	0.73	8
Travel	0.53	0.74	0.62	35
|———-Air Travel	0.84	0.89	0.86	18
|———-|———-Airport Parking & Transportation	0.00	0.00	0.00	1
|———-Bus & Rail	0.71	0.71	0.71	7
|———-Car Rental & Taxi Services	0.89	0.80	0.84	10
|———-Cruises & Charters	1.00	0.82	0.90	17
|———-Hotels & Accommodations	0.57	0.84	0.68	19
|———-|———-Vacation Rentals & Short-Term Stays	1.00	0.33	0.50	9
|———-Specialty Travel	0.00	0.00	0.00	2
|———-Tourist Destinations	0.29	0.17	0.21	12
|———-|———-Beaches & Islands	0.79	0.73	0.76	15
|———-|———-Mountain & Ski Resorts	1.00	0.80	0.89	10
|———-|———-Regional Parks & Gardens	0.50	1.00	0.67	2
|———-|———-Theme Parks	0.75	0.50	0.60	6
|———-|———-Zoos-Aquariums-Preserves	0.00	0.00	0.00	2
accuracy	0.71	0.71	0.71	0
macro avg	0.66	0.60	0.61	5893
weighted avg	0.73	0.71	0.71	5893

F PROTOTYPE VISUALIZATIONS

The following figures relate to the mock-ups we used in feature-testing interviews (Section 3.3). These mock-ups are prototypes and necessarily changed during development. Additionally, not all prototypes are included below (originally a set of 23 possible visualizations). Given time constraints, the impacts perceived by participants, and the data requirements of some visualizations, only a select set of visualizations were picked for development. It is also noteworthy that some visualizations were not known to be possible before development. For example, we did not know about the trends needed to make *Google Interests Dynamic* until we had developed *Google Interests*. When applicable, participant quotations are provided to add context for why we picked a certain mock-up for development.

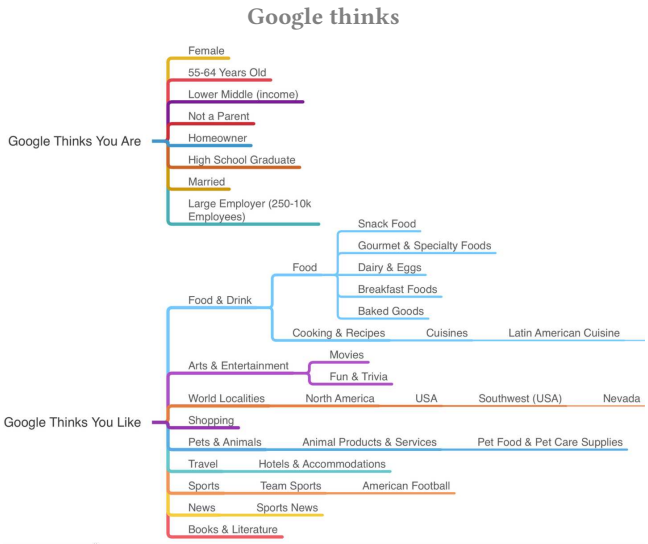

Figure 13: Although this visualization was presented to participants as a unit, in development we split *Google Demographics* (Google thinks you are) and *Google Interests* (Google thinks you like) into separate visualizations. Reactions to demographic visualizations are captured by P-01: "It feels a bit creepy at this point, like they seem to like make deductions about you based on what you look at."

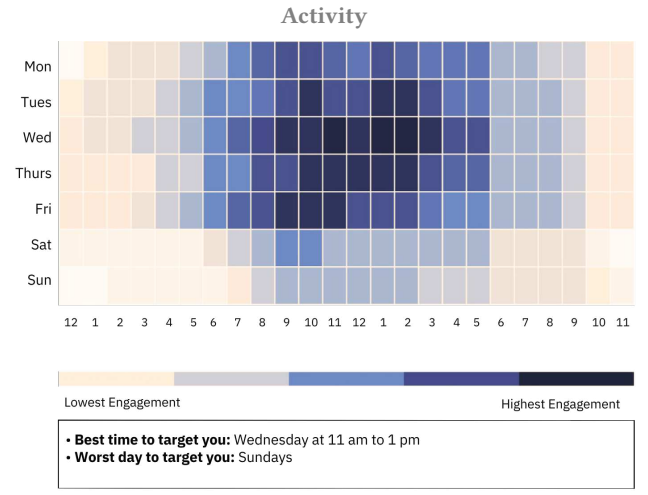

Figure 14: Prototype for the *When You Are Engaged* visualization. Participants reacted strongly to the inferences being drawn from the chart, as P-02 notes: "[T]his is slightly creepy because it says best time to target you and worst time to target you, it knows when I'm, you know, using the Internet a lot, it feels a bit scary [like] 'she's ripe for the plucking at this time of the day."'

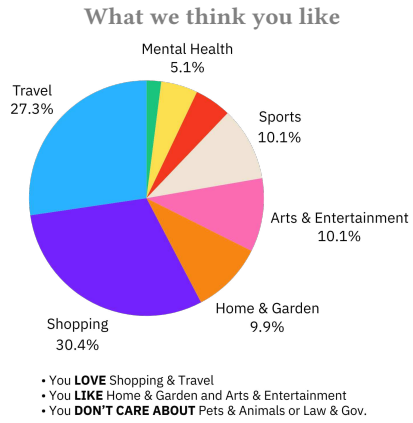

Figure 15: Prototype for the *Time Spent per Interest* visualization.

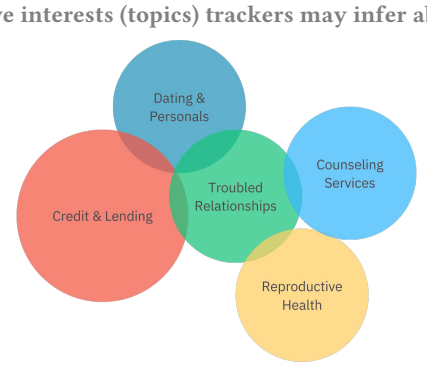

Figure 16: Bubblechart used in combination with Figure 17 to create *Possible Sensitive Interests*.

What pages you visited on plannedparenthood.com	Sensitive topics trackers may think you are interested in	Who tracked you on this page
"Abortion Clinics Near You"	Reproductive Health	Google, Amazon
"Emergency Contraception"	Reproductive Health, Dating & Personals, Troubled Relationships	Amazon
"For teens"	Dating & Personals, Troubled Relationships	Google, Amazon

Figure 17: Wordcloud and tracker information paired together with Figure 16 to create the *Possible Sensitive Interests* visualization. Participants found the visualization invasive and wanted to be removed from it. When asked "is there anything you would want added to this visualization" P-19 said "take it all away."

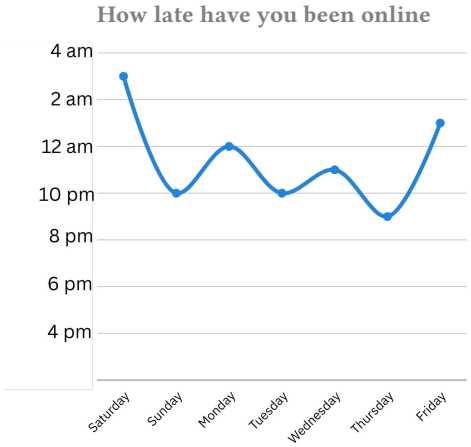

Figure 18: Prototype for the *Late-Night Engagement* visualization. Participants, like P-05, noted how this visualization's precision may be unnerving: "getting into like that, that level of data, like, you know, having advertisers know what time you're going to sleep is just a little unnerving."

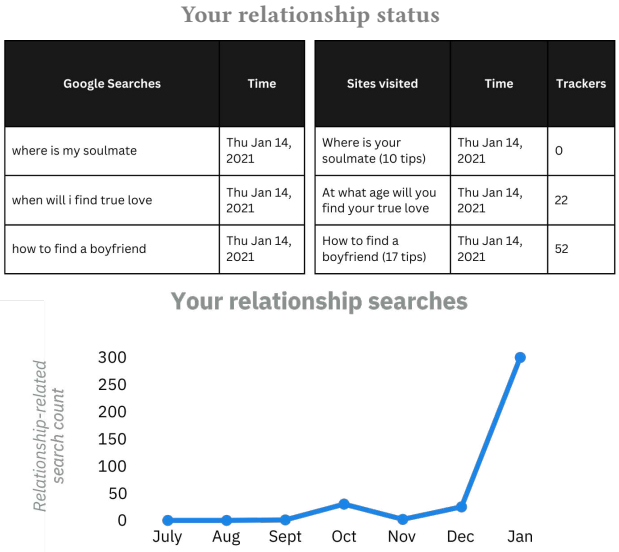

Figure 19: Relationship status and relationship searches combined to produce *Search Habits* visualization. Participants here, like P-01, noted how personal the *Search Habits* visualization felt: "it steps a little bit closer to a little too personal. . . . I'm not sure what a tracker would want to do with that information, and it does make me a little uncomfortable."

Figure 20: Combined with Figure 21 to create the ad-vault-style [51] *Ad Explanations* visualization. The exclusive focus on health was broadened because this type of data was hit-or-miss in participant data. Health ads in particular seemed to capture feelings of creepiness, as P-14 states: "The health ones make me a little more concerned because . . . let's say, . . . as you to type in something and then they might think—whoever is tracking you—might think that you have it so this is definitely more uncomfortable."

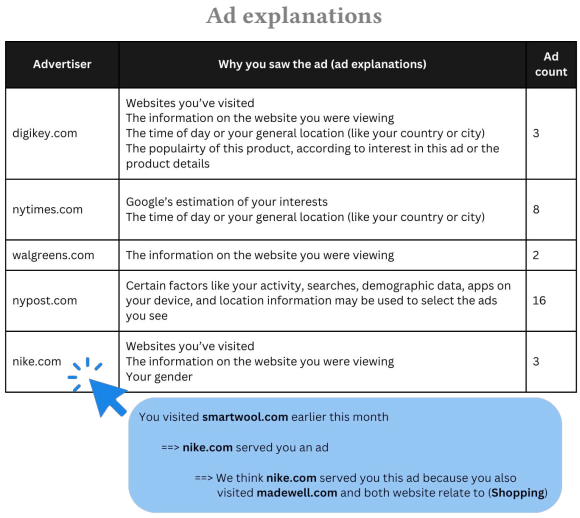

Figure 21: Combined with Figure 20 to create the *Ad Explanations* visualization.

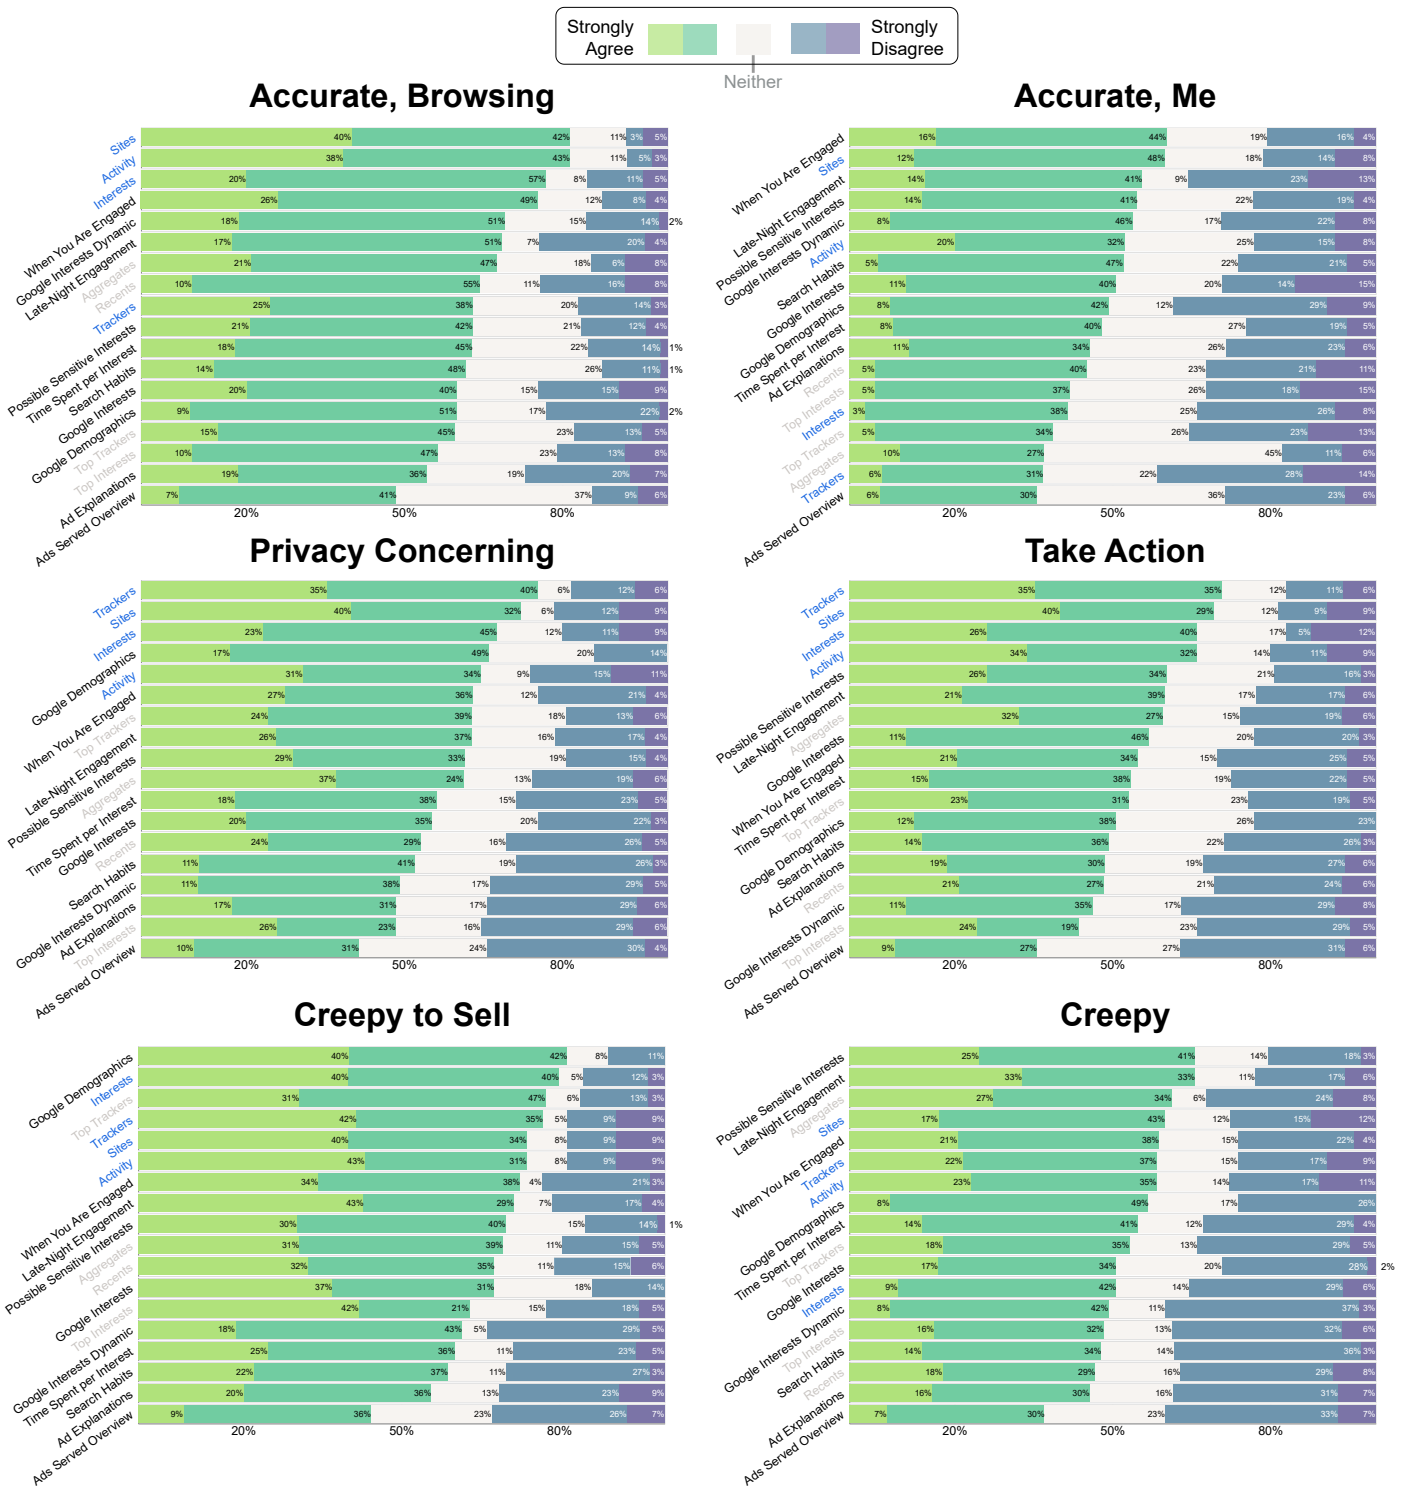

Figure 22: Ordering visualizations by the most participants who agreed or strongly agreed to each question. For example, more participants found the *Sites* visualization (TT1-Everything) to accurately reflect their browsing habits (82%) as compared to the *Ads Served Overview* visualization (TT2) where only 48% of participants agreed or strongly agreed that it accurately reflected their browsing. Visualizations are noted per condition as TT1-Dashboard, TT1-Everything, and TT2.