

William Melicher, Mahmood Sharif, Joshua Tan, Lujo Bauer, Mihai Christodorescu, and Pedro Giovanni Leon

(Do Not) Track Me Sometimes: Users' Contextual Preferences for Web Tracking

Abstract: Online trackers compile profiles on users for targeting ads, customizing websites, and selling users' information. In this paper, we report on the first detailed study of the perceived benefits and risks of tracking—and the reasons behind them—conducted in the context of users' own browsing histories. Prior work has studied this in the abstract; in contrast, we collected browsing histories from and interviewed 35 people about the perceived benefits and risks of online tracking in the context of their own browsing behavior. We find that many users want more control over tracking and think that controlled tracking has benefits, but are unwilling to put in the effort to control tracking or distrust current tools. We confirm previous findings that users' general attitudes about tracking are often at odds with their comfort in specific situations. We also identify specific situational factors that contribute to users' preferences about online tracking and explore how and why. Finally, we examine a sample of popular tools for controlling tracking and show that they only partially address the situational factors driving users' preferences. We suggest opportunities to improve such tools, and explore the use of a classifier to automatically determine whether a user would be comfortable with tracking on a particular page visit; our results suggest this is a promising direction for future work.

Keywords: online tracking, human factors

DOI 10.1515/popets-2016-0009

Received 2015-08-31; revised 2015-11-19; accepted 2015-12-02.

1 Introduction

Online tracking has become a widespread practice. It allows for an increasingly personalized user experience and more effective ads, which support free online ser-

vices. However, users also have serious privacy concerns about online tracking. Current policy initiatives to limit online tracking have difficulty establishing guidelines that all stakeholders agree on. Technical solutions have usability problems that interfere with their adoption and efficacy [26].

Studies that investigate online tracking and behavioral advertising have found that users' preferences are complex and have suggested that they can be highly contextual and difficult to capture [24, 41]. Quantitative work has also suggested that online privacy decisions are contextual—users' preferences in hypothetical scenarios have been shown to vary considerably based on details of framing and context [24, 30, 37]. Overall, although the research community has achieved a general understanding of users' attitudes towards tracking (and a more specific understanding of preferences for behavioral advertising), it is far from clear under what specific circumstances users would consent to be tracked and how specific situational factors (e.g., properties of websites or trackers) lead to their decisions. At the same time, a more precise understanding of these factors is necessary if we are to develop effective solutions—technical or otherwise—to control tracking in a non-binary way.

To improve our understanding of how specific situational factors inform user preferences for online tracking, we conducted 35 in-person, semi-structured interviews. We examined participants' feelings and concerns about tracking in actual web-browsing situations that they experienced outside of the study setting. To that end, we first extracted participants' web-browsing history using a web-browser plugin, and structured each interview around a series of web-browsing tracking situations (i.e., visits to specific pages or pairs of pages) that the participant had experienced during the two weeks preceding the interview.

Following this methodology, we believe we have achieved a more precise and thorough understanding of users' comfort with tracking as an interplay between situation-specific factors, the harms and benefits these factors cause users to perceive, and users' general attitudes about tracking. By basing our study on concrete experiences, we are able to confirm and reinforce previous results derived from more hypothetical methods, as

William Melicher, Mahmood Sharif, Joshua Tan,
Lujo Bauer, Pedro Giovanni Leon: Carnegie Mellon Uni-
versity; {billy, mahmoods, jstan, lbauer, pedrogl}@cmu.edu
Mihai Christodorescu: Qualcomm; mi-
hai@qti.qualcomm.com

well as uncover specific new findings. For example, some users link their tracking preferences to how frequently they visit a site; and participants were less comfortable with the more invisible outcomes (price discrimination, revenue to web sites, etc.) than with more noticeable outcomes (ads, customization, etc.).

In general, we found participants to have reasoned, complex, and nuanced preferences, with their comfort with tracking in specific situations depending on a number of situation-specific factors. Our participants' general attitudes about tracking were often at odds with their tracking preferences in specific tracking situations; for example, individuals who generally preferred to avoid tracking were comfortable with being tracked in some situations (and vice versa). This interplay between general attitudes and behaviors in specific situations is well documented in prior work [8, 13, 28, 29]. Whereas previous work focused on the properties of trackers and the information they collect, we found that participants' comfort with tracking was usually based on the properties of the sites they visited, which led them to focus on specific potential harms or benefits. For example, factors such as participants' trust in or frequency of visits to a particular site played a significant role in determining comfort with tracking. The topic of sites also often played a role, sometimes as a proxy for the type of information participants believed trackers could learn about them.

In light of this understanding of how situational factors affect user comfort, we revisit several popular web-browser plugins and mechanisms for limiting tracking, and examine the extent to which they are (in principle) able to effectively implement user preferences. We find that while existing tools can often help address some of the outcomes of tracking that participants seek to achieve or avoid (e.g., avoiding unwanted ads), they are too coarse-grained and generally unable to take into account the situational factors on which users base their preferences. We examine where the existing tools fall short, often failing to support users' situation-specific needs and preferences; we offer a set of design suggestions for approaches that would better fit users' needs. Leveraging our understanding of how specific factors influence comfort with tracking, we explore whether it is possible to automatically determine, on a per-page-visit basis, whether a user is comfortable with tracking. Our initial results suggest this to be a viable direction: we train a classifier that can with minimal false positives identify about 50% of the situations in which a user is comfortable being tracked. We believe this is a promising direction for the development of web browser addons

that will automatically or semi-automatically allow or disallow tracking to match individual users' preferences.

2 Background and Related Work

Online tracking is used for a variety of purposes. It allows customized search results, tailored recommendations on websites, better understanding of audience characteristics through analytics, and personalized advertisements. Given these many uses, it is not surprising that online companies make extensive use of this practice. In 2011, third-party trackers were found on 79% of the Alexa top 500 websites [34].

In this section, we first discuss users' concerns about online tracking and the benefits they might gain from it. Then we present previous work that studies which factors affect users' preferences for online tracking. Finally, we discuss current approaches to provide tools to enforce users' preferences.

Concerns about online tracking By using online tracking, advertisers can target ads based on sensitive information [19, 22], discriminate against users [19, 39], or even manipulate users' purchasing intentions [16].

Due to these practices, surveys of internet users have found high levels of concern about online tracking. Turow et al. found that 87% of telephone survey respondents would not allow advertisers to track them online if given a choice [40]. Wills et al. built a website that users could visit to learn what trackers were present on the pages they had visited in the past, finding that users' expressed privacy concerns [43]. A more recent Pew telephone survey found that 68% of respondents did not like targeted ads due to concerns about online tracking [32]. However, qualitative research has found that users are not completely against targeted ads, but they are concerned about the lack of transparency and control over the tracking that enables it [41]. Apart from transparency and control concerns, users have also expressed concerns about the type of targeted ads that they might see, which can lead to embarrassment [9].

Prior work has discovered that there is a difference between the preferences and concerns that users' express in the context of research studies and their actual privacy behaviors [8, 13, 28, 29]. Our research supports this finding. In addition, the difference between hypothetical surveys and real behavior motivates our choice to ground our study in users' real behavior.

Benefits of tracking While online tracking raises privacy concerns, it also has benefits, both to users and advertising companies. The revenue that online tracking provides to companies enables some websites to provide free content [45]. Online advertising (which often employs tracking) was responsible for over \$49 billion in revenue for US practitioners in 2014 [15]. Although the advertising industry argues that behavioral advertising is a significant contributor to online advertising revenue [12], others have questioned its effectiveness [20, 25, 38].

In addition, previous studies have shown that users often find the outcomes of online tracking to be beneficial. For example, researchers conducting semi-structured interviews found that while 41 of 48 participants expressed privacy concerns, 31 participants still perceived targeted ads as useful [41]. Others found that participants preferred Google's personalized search results over "vanilla" results in 38% of the cases, whereas the opposite happened in only 23% of cases [31].

The desire to reap these benefits of advertising while minimizing annoyance from ads has prompted the creation of an "Acceptable Ads" program by industry players [6]. However, this program is focused on removing intrusive ads, not protecting privacy.

Factors that affect preferences Previous work has attempted to identify factors that affect users' comfort regarding online tracking. The majority of study participants were comfortable with tracking for the purpose of personalizing search results, as long as the stored information is not sensitive [31]. In another study where tracking companies' practices were shown to users, users relied mostly on tracking companies' sharing and retention practices to decide what types of information they would disclose for the purpose of receiving targeted ads [27]. In a follow-up work, participants expressed willingness to share more data if given prior control mechanisms to select the type of shared information, restrict first and third parties from collecting their info, and customize the topics of targeted ads [17]. In the same work, researchers found that frequency of website visits, having an IT background, general privacy attitudes (measured via the Westin Index), and the intention of exploring online ads are correlated with participants' willingness to share. Different browsing scenarios (e.g., planning a vacation, looking for a job, etc.) affect participants' comfort with tracking [9, 41]; furthermore, users' awareness that tracking spans visits to multiple sites is positively correlated with their concerns about tracking [33].

Our work differs from previous work in that we gain a more comprehensive view of the factors that affect users' comfort with online tracking due to not focusing on particular aspect of tracking via semi-structured interviews. Our study is also the first to explore participants' tracking preferences based on their real browsing histories, which we believe allows us to capture users' privacy preferences more faithfully.

Tools to control online tracking Many tools to control online tracking have been developed. Popular browser plugins such as Adblock Plus, Ghostery, and Blur [1–3] allow users to selectively block third-party trackers. Another plugin, ShareMeNot, focuses on preventing social widgets from tracking users [35]. Modern web browsers include privacy settings to manage cookies, incognito browsing modes, and a setting to signal a global preference to not be tracked (i.e., Do Not Track). The W3C recently published tentative standards for Do Not Track (DNT) [5], providing concrete guidance on how companies should respond to DNT requests and eliminating a common reason for industry non-compliance. The ad industry allows users to install specific cookies to opt out of targeted ads from ad companies affiliated with a self-regulatory organization [7]. However, the plugins and tools largely offer users binary alternatives to enable or disable online tracking or targeted ads on a per-company basis. Furthermore, opt-out options sometimes only prevent the showing of targeted ads, rather than preventing tracking. As previous research has shown, users are not satisfied with these features [9, 17].

Several systems (e.g., Adnostic [36], Privad [23], RePriv [21], and CoP [14]) have been proposed to protect users' privacy by treating the collection of users' data and the usage of the data as two separate tasks. While promising, these systems face several obstacles hindering their adoption. Primarily, they require changes to the advertising ecosystem that advertising companies have little incentive to perform.

As browser cookies are a popular method to track users, effective management of cookies has the potential to limit online tracking. However, users struggle to use the available tools for managing cookie-based tracking due to usability issues. Researchers have tested the usability of nine popular tools to limit tracking and found that participants faced problems understanding and using them; many users believed they had configured the tools to limit tracking when they had failed to do so [26]. There is a need for user-friendly tools that align with users' skills and mental models.

In this paper, we examine users' preferences for online tracking to inform the design of tools that prevent the drawbacks of tracking while simultaneously allowing the benefits. To that end, we aim to uncover the situational factors that a tool should provide controls for in order to best implement users' tracking preferences.

3 Methodology

We conducted 35 in-person, semi-structured interviews with internet users about online tracking. We elicited participants' online-tracking preferences in the context of their own browsing history. Participants were required to use either the Chrome or Firefox browsers in order to install a web browser plugin that would allow them to send their sanitized browsing history to researchers. The interviews lasted approximately 60 minutes and participants were paid \$15 each. The study protocol was approved by the Carnegie Mellon University Institutional Review Board.

3.1 Participant Demographics

We recruited 17 women and 18 men from the Pittsburgh area via Craigslist, flyers, and university forums to participate in “a study about online privacy behaviors.” Participants' ages ranged from 18 to 58 years, with an average of 27 and standard deviation of 8.2. The participants had a variety of occupations including students, nurses, artists, and business-related jobs. Twenty-five held a bachelors or higher degree. All of them identified as either very or somewhat technically savvy, and 12 had an IT-related background. Twenty-one participants reported that they use ad-blocking software. Our participants skew towards being more technically savvy and are from a specific demographic group, Firefox and Chrome users.

3.2 Interview Procedure

We began each interview by examining participants' awareness of and general opinions about online tracking. Participants were asked to read a short description of online tracking to establish a baseline understanding (see Appendix A.1). Then, we collected participants' general perceptions of benefits and risks of online tracking and their knowledge of and experience with tools to control tracking. Finally, we elicited attitudes towards



Fig. 1. Example of web history print outs that are shown to participants during interviews.

online tracking for a set of page visits from participants' own browsing history. The interview script is shown in Appendix A.

With the exception of general attitudes towards tracking, all results are based on web pages participants actually visited. For these visits, participants were asked what they perceived as the positives and negatives of tracking, what they thought would happen as a result of tracking, and their comfort with tracking. Responses were based on participants' understanding of what did or could have happened.

Participants were shown information about specific pages they had visited, including the time each was visited, the web page's title, the website's URL, and a screen shot of its homepage. Figure 1 shows an example of what participants saw. We describe how web pages were selected for this purpose in Section 3.3. If a participant did not remember visiting a specific web page, we showed the participant a different web page that met the same selection criteria. On average, participants were interviewed eleven days after they submitted their browsing history. We first showed participants up to four individual web pages and elicited their first-party tracking preferences. We then showed participants up to seven pairs of web pages and elicited third-party preferences. In each situation, participants were prompted first to think about the specific benefits and risks of being tracked, and then to think about the information that could be learned from tracking participants, and how it could be used. In order to reduce negativity bias [11], we asked participants positive questions before negative ones. We did this because we observed negativity bias during pilot interviews, in which participants had difficulty discussing the benefits of tracking after discussing the negative aspects.

3.3 Website Selection

We screened participants to ensure their preferred browser was one of the two major browsers—Firefox or Chrome—for which we developed plugins to extract browsing history. Participants submitted their browsing history for the last two weeks before signing up for the interview. This length of time was chosen to capture a representative fraction of participants' online behaviors while allowing them to remember the situational factors of their web page visits. Participants had the opportunity to remove any web pages that they did not wish to share via a software interface that we provided. We collected browsing history from 48 participants, though we only present data for the 35 eligible participants that were able to attend the interview.

Prior to each interview we analyzed the received history and selected the web pages to show during the interview. Web pages were chosen to cover a wide range of situations about which participants might have different feelings. Participants were invited to interviews only if the websites in their history covered at least seven of the eleven situations of interest. Since we expect that most internet users would encounter most of these situations within a two-week period, this criteria was primarily intended to detect cheating. Only one prospective participant did not meet this requirement and they were not compensated.

To allow us to sample web pages from different categories for use during the interviews, we first automatically categorized all web pages in participants' histories using information from DMOZ, a crowd-sourced database of websites tagged by topic [4]. The types of websites we selected to ask participants about are described in Appendix A.2. This method of selecting websites was not meant to perfectly represent all websites in a particular category or topic, but to sample a wide range of situations about which we could ask participants. Since each participant saw slightly different situations, some participants had the opportunity to comment on different issues. Our analysis is meant to show the range of opinions participants have about tracking based on their own web history.

The majority of participants (80%) removed at least one web page from their history, though the median percentage of pages removed (2%) was low. We did not collect information about the websites participants sanitized beyond counting them.

3.4 Data Analysis

Interviews were recorded and transcribed. We performed an exploratory and qualitative analysis of participants' responses. Transcripts were unitized and segmented by the tracking situations that participants were asked about. Two coders independently analyzed a subset of interviews and developed their own codes. They then discussed the individually derived codes and agreed on a final set of codes. Using the agreed-upon codes, they individually coded each of the interview questions, and met to resolve disagreements. Disagreements were resolved by the two coders discussing each instance of disagreement and reaching agreement by clarifying coding definitions. Before coding resolution, the coders agreed on 91% of the coding instances. After coding, we examined the relationship between the perceived outcomes and the decision-making factors. We also quantified the frequencies of different decision-making factors.

3.5 Limitations

Because participants were asked to submit their browsing history, our sample of participants was likely less privacy sensitive than the general population. To reduce this bias and mitigate privacy risks, we allowed participants to filter their browsing history before submitting it to us. As participants could prune their browsing history, the websites that we asked participants about might be skewed toward less privacy-sensitive ones; however, we did not discover evidence of systematic pruning.

Our selection of websites was based on information in the open directory project [4], which includes over four million sites but may have nevertheless resulted in a skewed sample. While we asked about only a handful of websites for each participant, they were carefully selected to explore preferences about a wide range of different situations.

Our data was collected via after-the-fact interviewing. Interviews conducted in this way are not without bias—participants may alter their opinions in the interview to rationalize their past behavior. However, this bias is likely to oppose the bias in previous work, which uses hypothetical survey methodology [24, 30, 37]. In addition, we believe that this bias is countered by the fact that we grounded interviews in participants' real browsing histories to more realistically capture preferences that matter to them.

The number of participants and websites we investigated was purposely limited to enable us to explore

nuanced and detailed preferences about specific situations. We would not have been able to explore preferences to the detail we desired with a larger sample. We only recruited participants who could attend an in-person interview in our geographic location. Our sample is not representative of the population as a whole, and as such, some of our findings may not generalize beyond our sample. Our sample of participants skewed towards being more tech-savvy, only used the Firefox and Chrome web browsers and used browser plugins more than the general population.

Regardless of the small sample size and self-selection, participants expressed a wide range of opinions enabling new findings and helping better inform the design of online privacy tools. We do not make claims that our participants are representative of the population, but only that our analysis yields insight about promising directions for future privacy plug-in design.

4 Results

Participants expressed a variety of different opinions about online tracking, which we now present. First, we discuss our participants' general attitudes and conceptions, in Section 4.1. In Section 4.2, we discuss the benefits and harms of tracking perceived by participants. Finally, in Section 4.3, we examine how situational factors affected participants' perceptions of these harms and benefits.

4.1 General Attitudes and Conceptions

Here we describe the wide range of general attitudes that participants had toward online tracking. Although participants varied in the amount of information they believed could be learned about them by trackers, participants who believed that trackers would learn either nothing or everything about them were a minority.

Attitudes toward tracking Participants' general attitudes toward tracking often guided their comfort with specific situations; however, general attitude was not always the deciding factor. For example, P9 generally disliked tracking, but felt comfortable with tracking in each situation we presented. Our participants had a variety of general opinions about online privacy. Based on participants' responses to introductory questions, we identified four categories of opinions toward online tracking: generally negative (“dislike”); generally neutral (“OK”);

both positive and negative (“mixed”); and positive, provided certain conditions held (“conditional”). Fourteen participants saw tracking as conditionally positive. For example, P1 mentioned that he is okay with tracking as long as it has a “limit,” and does not collect personal information. Eight participants saw tracking as generally neutral, nine saw it as generally negative, and the remaining four had mixed feelings. A summary of participants' comfort levels in different tracking situations, as well as their general opinions toward tracking, is shown in Figure 2.

Twelve participants felt resigned to tracking. Although their feelings about online tracking varied, these participants believed that online tracking was unavoidable and that it was futile to even attempt to control it. These participants still had nuanced preferences for tracking in particular situations, but did not see a reason to put effort into limiting it. Additionally, six participants distrusted the effectiveness of available tools and three felt that using them was too much work.




Fig. 2. Participants' general opinion about tracking related to their comfort with tracking in our tracking situations. Each participant's general opinion, shown in parentheses, is one of: dislike, would like if some conditions were met, mixed feelings about tracking, and “OK” with tracking. * indicates participant was generally resigned to tracking.

Conceptions about revealed information In addition to general attitudes about tracking, participants had a range of different perceptions of what tracking companies could learn about them. At one end of the spectrum, P12 thought that trackers could only access a limited subset of the information on websites being tracked. When asked what she thought her Google searches might reveal about her, she explained “I don’t give my credit card details or ... put my age here or anything. It’s just a normal search ... so I don’t think they know anything.” At the other end, P18 held more extreme views, believing that trackers could access all information from tracked websites. This included the contents of her bank statements and private communications via email and social networking sites. P18 was uncomfortable with a third party tracking his Facebook visits because “it’s a kind of private medium to talk ... I might have some private data that I extend with them. So I don’t want at all to be tracked on Facebook.” Despite these varied perceptions of what level of access trackers have, participants rarely considered the information that could be inferred from their visits to the web page, or the website in general, over time.

It is worth noting that all participants were provided a description of tracking at the beginning of the interview, which stated “credit card numbers, Social Security numbers, passwords, and other sensitive personal information is normally out of reach of online trackers.” Seven participants mentioned that they did not trust this description and believed sensitive information could be easily obtained by trackers. This view often coincided with the belief that trackers would use malicious means to gather data. Additionally, after reading our description, seven participants expressed being comforted to learn that sensitive personal information was unlikely to be tracked, as this was something they were previously uncertain of.

Misconceptions about online tracking For each tracking situation, we coded whether participants demonstrated a misconception about online tracking. Although we provided participants with a description of online tracking at the beginning of the interview, 14 participants demonstrated a misconception about online tracking at a subsequent point in the interview. These misconceptions related to online tracking mechanisms: ten participants conflated online tracking with malware (e.g., hidden scripts) and four participants believed it directly involved local browser website history.

Participants’ misconceptions about online tracking influenced their comfort with tracking to different de-

grees. For each tracking situation in which participants demonstrated a misconception, we coded whether the misconception seemed to be the primary factor influencing their comfort with tracking in that situation. Our results suggest that misconceptions had little overall impact on participants’ elicited preferences. Only three participants demonstrated a misconception that appeared to significantly affect their comfort with tracking in a specific tracking situation. In these cases, participants believed the tracker could use malicious scripts to obtain sensitive information and were therefore less comfortable being tracked. The eleven remaining participants with misconceptions about tracking did not appear to be strongly influenced by those misconceptions; instead, their comfort seemed to be determined by factors such as whether personal information was involved or whether the participant trusted the tracker.

4.2 Users’ Perceptions of Outcomes

Participants discussed a variety of online-tracking outcomes: some that were overtly visible—including targeted ads and customized web sites—and some that were not—such as more revenue for companies. Participants also discussed different harms and benefits that are possible from each one of the outcomes (e.g., targeted ads could be useful or embarrassing). In this section, we separately discuss the noticeable (Section 4.2.1) and the less overt ones (Section 4.2.2), their benefits and harms as conceptualized by the participants, and how these contributed to participants feeling comfortable or uncomfortable with tracking in specific situations. Table 1 summarizes participants’ perceptions of different perceived outcomes of tracking; Figure 3 shows the relationship between these perceived outcomes and participants’ comfort with tracking.

4.2.1 User Noticeable Outcomes

One type of outcome that participants experienced as a result of tracking were things that they would be able to directly notice. Three of these outcomes—targeted ads, customization of websites, and legal harms—received a mix of positive and negative reactions from participants. At the same time, online tracking also triggered a “weird feeling” for some participants. Here we present a detailed analysis.

	Outcome	(#) Positive	(#) Negative
Noticeable	Targeted ads	(26) Relevant, useful; preferred to typical ads	(21) Distracting, annoying; Others might notice
	Feeling “stalked”		(24)
	Customized websites	(11) Better search, recommendation experience	(5) “tracking bubble”
	Legal action		(8) Legal consequences, e.g., when file sharing
Invisible	Company revenue	(11) Access to free services	(8) Feel used by companies
	Price discrimination	(11) Lower prices, sales, coupons	(5) Potential for higher prices
	Data linked to identity		(12) Feels privacy invasive

Table 1. Number of participants that mentioned possible outcomes of tracking. Participants saw trade-offs of different outcomes depending on the situation.




Fig. 3. Percentage of tracking situations in which participants see a particular outcome and their comfort for tracking in that situation. Situations in which participants perceive invisible harms are also those in which they are most uncomfortable.

Targeted ads Among the most common of noticeable outcomes were targeted ads. The majority of participants (74%) saw targeted advertisements as beneficial or at least preferable to non-targeted ads in at least one tracking situation. In these cases, participants expressed that advertisements targeted to their interests could help them save time or money. P20, for example, explained that targeted ads could help remind her of products that she wanted to buy. Sometimes, participants felt that these ads might be helpful for other people, but not for themselves. As an example, P21 said: “it just depends on the kind of person that you are. Some people might find that beneficial.”

More than half of participants (60%) found targeted ads harmful in at least one tracking situation. In these cases, participants cited the repetitive nature of targeted ads as both annoying and “pushy.” Moreover, targeted ads were sometimes distracting to participants if they were targeted to their interests on topically different sites. Additionally, several participants said that sometimes they feel embarrassed if ads targeted to them are observed by others. For example, P19, a fan of anime, was concerned that she might be shown, and other people might notice, anime ads when she is not browsing anime-related websites.

Customized websites About a third of participants (31%) perceived potential customization of web pages as beneficial in at least one scenario. In these cases participants felt that customization might save them time if they are searching for something or allow targeted content. P29, for instance, appreciated that such customization could be used to target Google’s search results. Participants similarly liked that social media feeds and shopping sites are customized to their interests. P14, for example, thinks that it is “the whole point of Facebook and social media websites.”

However, a handful of participants (14%) saw customization as being harmful in at least one situation. In these cases, which mainly involved search and shopping websites, participants did not want a “tracking bubble” phenomenon to dictate what content they saw. These participants were uncomfortable with a company making decisions based on past browsing behavior about what content to show them. Participants felt that inferences based on browsing history might be incorrect or inaccurate. P7 wanted customization to help her, but not bias what she was shown. Common to these cases was participants’ wish to just “see the website” as it is, not manipulated content.

Legal action Several participants (23%) were concerned that tracked information can lead to real-world harms in the form of legal repercussions or persecution based on their beliefs. P33 said she would be concerned if she were buying incriminating items. Although participants expressed general concerns of this type, our data set is understandably skewed towards not containing this information for specific situations, as participants were given the opportunity to cleanse sensitive data from their browsing history before submission. One participant, P8, expressed that tracking for the purpose of persecution or legal action was potentially beneficial if done against “the bad guys.”

Negative feeling Over two thirds of participants (69%) expressed occasionally feeling “weird” or “creeped out” by tracking. P23 said she felt like she was being “stalked” by trackers in some cases. In most cases, targeted ads were the catalyst for that feeling. For instance, P28 feels “weird” when receiving emails about volunteering in Tanzania, a place where she volunteered five years ago.

4.2.2 User Invisible Outcomes

In contrast to outcomes that participants could easily notice, participants also mentioned several outcomes that are invisible to internet users. The invisible outcomes that were mentioned include: companies' profit, price discrimination, and data being linked to the identity of the users. While the perceived benefits did influence their comfort level (Figure 3a), the perceived harms had a more significant effect on their discomfort level (Figure 3b). Scenarios in which participants saw invisible outcomes as harmful are also situations in which participants were least comfortable. In this section we present the participants' perceptions of these outcomes.

Price discrimination Participants indicated that prices for online goods might be manipulated based on the information a company had about them. About a third of participants (31%) were hopeful that prices would be lowered in the form of sales or deals, but five participants recognized that prices could also be raised.

Company revenue Participants also expressed the belief that tracking would be beneficial for companies. Whether participants saw that as beneficial or harmful to themselves varied. About a third of participants (31%) liked this outcome, saying that the revenue that companies get from this practice allows them to provide

free or better services. P28 explained “It’s just helping businesses stay in business so I can use their services for free.” However, about one fourth of participants (23%) felt used by this practice in at least one instance, saying that tracking purely for the purpose of profit was harmful or unpleasant in that the company benefits but the user does not. P15 exclaimed, “This is just typical profiting off my information.”

Inference Participants were generally unsure what information about them could be inferred. They were understandably wary of this uncertainty, and were concerned that it was difficult to make decisions about tracking without knowing what could happen. P13 says: “Maybe they get something out of it. But just to be sure I would rather have no tracking.” Others were more generally open to being tracked in situations where inference could have occurred, provided it did not harm them. While we did not test the technical knowledge of our participants, many participants revealed an incomplete understanding about what could be inferred from their web history. For example, participants rarely considered that detailed information about them could be inferred from pages that do not directly reveal that information. P3, for example, did not think that her income information would be revealed from her searches on Google. However, targeting users based on their income bracket is an optional feature of the Google AdWords product. In addition, participants tended to reason separately about each pair of websites in the third-party scenarios, rather than consider what an advertiser with knowledge of both pages' visits might infer.

Data linked to identity Twelve participants felt that linking data to their identity or other identifying information would be an invasion of their privacy. P16 explained that she is OK with a third party tracking her on eBay as long as the tracker has no access to the personal information on the page, which can be used to tie the collected data to her identity. Related to this idea, P31 felt that aggregating data from multiple sources was undesirable. When asked to specify how she would like to control tracking, she explained that tracking should be limited to the visited website, as aggregating data from multiple websites may create a detailed image of her linked to her identity, and that would be too “invasive.”

4.2.3 Third-Party and First-Party Differences

Participants perceived a large difference between tracking by a first party and by a third party. We coded par-




Fig. 4. Comfort of participants in different scenarios. Participants are often more uncomfortable with third-party tracking than first-party tracking, particularly for websites of sensitive topics.

Participants' comfort in each scenario to which we exposed them; the results are shown in Figure 4. Participants were less comfortable overall with tracking in all third-party scenarios than any first-party scenario. Interestingly, over 75% of participants in the “sensitive” first-party tracking scenario were comfortable with tracking. However, less than 25% of participants were comfortable with visits to these same websites being tracked by third parties. A plausible explanation for this phenomenon is that the participants felt safe being tracked by the first party with whom they willingly share sensitive information, but were concerned about the potential of a third party getting hold of this information. This is exemplified by P3, who felt comfortable being tracked by her banking website, which she perceives as trustworthy, but was uncomfortable being tracked by a third party on the same website, as she was worried that personal information learned by the third party may be used to cause her harm.

4.2.4 Effect of Benefits and Harms on Comfort

We were also interested to see how the benefits and harms resulting from outcomes relate to the comfort levels of the participants. We examined the comfort level of participants in scenarios where they saw only beneficial outcomes, only harmful outcomes, both beneficial and harmful outcomes, or no outcomes at all. The results are illustrated in Figure 5. Participants were rarely uncomfortable when they saw only benefits, but when they perceived both benefits and harms, participants became less decisive and conditioned their comfort on different factors. Additionally, discomfort increased in cases where there were no perceived benefits. While these results are not very surprising, they imply that




Fig. 5. Comfort of participants in scenarios where they saw benefits, harms, both, or neither.

participants' preferences in our study were guided by the perceived benefits and harms.

4.3 Factors that Affect Users' Preferences

In Sections 4.1 and 4.2 we described participants' general attitudes toward tracking, the outcomes they perceived tracking to have, and how these made them more or less comfortable with tracking. In this section, we examine how users decide, when considering visits to specific web pages or specific pairs of pages, whether they are comfortable with tracking. We find that specific properties of websites, including their topics and the types of information they have about users, as well as specific properties of trackers, influence which outcomes users will perceive, whether they perceive these outcomes as positive or negative, and consequently influence their comfort with being tracked. We generally find that these situational factors frequently lead users to have different preferences about specific instances of tracking than suggested by their general attitudes.

We find that some of the specific factors that influenced participants' preferences about online tracking are related to the information being tracked, such as whether it was considered personal; others involved non-informational properties, such as trust, awareness, and consent. A summary of these situational factors and their frequency is shown in Table 2.

4.3.1 Factors Related to Information Tracked

Personal info Nearly all participants (97%) expressed that they were less comfortable when personal information was tracked. Although participants' responses varied when we asked what they meant by “personal information,” they typically listed some combination of information that could uniquely or nearly uniquely identify them, including email address, location information (e.g., physical address or more detailed data), and their

	Factor	Effect	# 1 st pty	# 3 rd pty	# Overall	Why
Informational properties	Has personal info	—	27	31	34	More invasive, perceived as more risky
	Has social info	+/-	7	14	18	Opportunity for customization, but also privacy invasive, and others might see information about the user
	Has search info	+/-	14	14	18	Opportunities for customization, but also “tracking bubble” effect, and reveals other information, specific searches are sometimes private
	Has correspondence	—	7	16	18	Private info disclosed
	Has financial or health info	—	4	10	11	More opportunity for price discrimination; privacy sensitive
	Has shopping info	+/-	8	7	10	Better customization and prices, shopping recommendations, but chances for mistargeted ads, distracting or unwanted ads
	Has education info	+/-	7	5	11	Innocuous info to some, not to others
	No volunteered info on site	+	6	4	9	No apparent harms
	Sharing with other 1 st parties	—	0	26	26	More opportunity for data aggregation
Non-info. properties	Trust	+	11	8	13	Less chance of bad invisible and visible effects
	Lack of consent	—	6	11	14	Less transparency into invisible effects, no chance to limit bad effects
	Lack of awareness	—	2	8	8	Less transparency into invisible effects
	Sites infrequently visited	+/-	5	6	9	Infrequently visited sites don't reveal lots of information, but can misrepresent user's preferences for ads and customization
	Sites frequently visited	+/-	2	4	5	Frequently visited sites represent interests for customization and ads better, but might learn a lot of information about the user

Table 2. Description of situational factors. “Effect” shows whether this factor has a positive or negative effect on tracking preference. “Overall #” shows how many participants mentioned this factor. “1st pty” and “3rd pty” shows how many participants mentioned this factor as mattering in first-party and third-party tracking situations.

full name (for some people). Participants often felt that trackers knowing their personal information was an invasion of privacy and that it resulted in a greater potential for harm.

Search info Slightly over half of participants (51%) felt that information about their general searching activities, for example on Google or Yahoo, impacts their feelings about tracking. Participants felt that being tracked on search websites was potentially beneficial in 36% of situations where this was a factor. At the same time, they also expressed concern about the level of detail that some searches can provide in 63% of situations where this was a factor. Participants found these situations to be conflicting: on one hand, tracking search information can help save time by identifying popular web sites; on the other hand, it might reveal information about the user. In addition to information about general interests, participants were concerned that some searches might reveal information like the participant's address from map searches, or clues about that participant's social circle. P30, for example, was concerned that Google might know her address based on her history of searching for directions on Google Maps.

Correspondence About half of participants (51%) felt that their private communications and correspondence with others should not be tracked for any reason. This included their emails in online email interfaces, private chat logs, and private messages on other sites. Participants did not feel this way about public communications, for example, a public post on Facebook, which they felt was appropriate to track because there is no expectation of privacy.

No volunteered info Roughly half of participants (51%) were asked about at least one site for which they expressed they were comfortable with being tracked because they felt the site had no information or inaccurate information about them. Many (26%) felt this way in situations in which they did not volunteer any information to the website. Participants generally thought that sites on which they did not specifically enter information revealed nothing about them. While participants felt comfortable in these situations, some expressed that they did not know how tracking those visits would be useful to trackers.

Social info Approximately half of participants (51%) felt that tracking social information, such as the identity

of a user's friends and family, affected their privacy preferences. In 88% of cases where participants mentioned social information, it made them more uncomfortable. P30 was concerned that her information from Facebook might be used to advertise products from Macy's; she found this uncomfortable. Participants were worried that this information might be used to target products to them or to their friends. However, sometimes participants thought there could be a specific customization benefit to this. P18, for example, thought that the website LinkedIn might use third-party information to customize the website by showing him more articles for companies he is interested in.

Financial and health info Slightly less than one third of participants (29%) felt that information about finances or health, for example, information from financial-service sites like banks or online payment sites (e.g., Paypal), was particularly sensitive and had greater potential for misuse. Participants were sometimes unsure why they felt this way, but a few participants believed that they might be the target of price discrimination when determining loans or rates. P26 thought that she would be very uncomfortable if a tracker learned any of her financial information from the TurboTax website that she uses to do her taxes. Even for the purpose of targeting ads, these participants felt that using income information about their purchase history was not appropriate. This situational factor was most often voiced as a concern in third-party tracking situations. As shown in Figure 4, this suggests that third-party tracking is particularly unpleasant when it involves tracking sensitive information.

Shopping info About a third of participants (29%) felt that a tracker learning information about their shopping habits affected their preferences. Participants were for the most part comfortable with tracking on shopping websites; in 73% of the cases where shopping information affected their opinions, it made them more comfortable with tracking. Indeed, participants felt that such tracking would be beneficial to them when finding new products or new sales. However, there are also occasions where this tracking information would make participants feel less comfortable. In these situations, participants felt that there was greater potential for price discrimination or that this data might manipulate their behavior, causing them to spend more money.

Educational info Almost a third of participants (31%) felt that educational information was appropriate to track both by first parties and third parties.

While participants did not see a particular benefit to this tracking, they also felt that such information was "already out there" or not harmful. Sites related to education, such as a school's web page, were often seen as trustworthy.

4.3.2 Factors Not Related to Information Tracked

Participants' comfort with tracking was often informed by factors such as trust in the tracking party or whether they had consented to tracking by that particular party.

Trust About a third of participants (37%) described their trust in the tracking party as affecting how they feel about tracking. This might be a first-party or third-party tracker. Participants used "trust" to indicate that they trusted a tracking party to use information appropriately and safely. Trust was occasionally important even when tracking relatively unimportant data about the participant. Participants were more comfortable with tracking when they trusted the website, and more uncertain when they did not. Participants described trust as a company having a "big name," or even in some cases if they personally knew the owners of the website.

When asked about a website he says he does not visit frequently, P28 expressed concern that he did not know if the site was ethical or would handle his data ethically. For him and other users, trust was linked to the frequency of visits to a site. In counterpoint, P29 expressed distrust of Google, whose website she visits frequently.

Awareness and consent Many participants stated that their decisions would depend on whether they were aware (23%) or consented (40%) to a tracker collecting information about them. While participants also sometimes stated that they were unlikely to read or understand EULAs or "legalese," they also felt that other, more reasonable means for promoting awareness and consent could be used. Participants voiced these concerns primarily in the context of third-party tracking.

Visit frequency A few participants (14%) expressed that the frequency with which they visit a website could be used as a proxy for whether their visits should be tracked. Some participants preferred for visits to infrequently visited websites not to be tracked. This is because they may have gotten there from a search or did not know what the website was about before visiting it. They felt that websites they visit infrequently do not

accurately reflect their preferences for either ads or customization, whereas websites they frequently visit do accurately capture such preferences. In addition, participants thought websites they frequented deserved to get more revenue from advertising. P20, for example, said he would be comfortable with tracking that takes place only on websites that he frequently visits. Interestingly, a few participants (26%) felt that they were more comfortable with tracking on sites they visit infrequently. Participants believed these sites would have less information about them, and that they would have less exposure to harmful effects of tracking on these sites.

5 Controlling Tracking

In this section, we evaluate current tools to control tracking in light of our findings. We also propose new ways to leverage situational factors to ameliorate harms of tracking while allowing benefits, and describe an initial exploration of a classifier that could be used to help automate the implementation of tracking preferences.

5.1 Current Tools to Limit Tracking

Current tools to limit tracking are technically capable of preventing any specific instance of (cookie-based) tracking. However, users rarely install or correctly configure tools, often due to usability problems [26]. Indeed, six of our participants volunteered that they did not trust the effectiveness of such tools and three saw these tools as requiring too much effort to use. Eleven additional participants felt generally resigned to tracking, feeling that it was unavoidable or pointless to attempt to stop. At the same time, all participants felt that there could be benefits to online tracking in certain situations.

We evaluated the following privacy tools and approaches: Adblock, Ghostery, Blur, Lightbeam, browser configuration, and private browsing mode. In Appendix B we describe the criteria, derived from the results reported in Section 4, by which we judge whether a tool is capable of allowing or preventing tracking according to users' preferences. We discuss how these criteria apply to each of the privacy tools in Appendix B. The evaluation of the ability of tools to address situational factors is summarized in Table 3; their ability to control possible outcomes of tracking is shown in Table 5 (Appendix B).







	Adblock	Ghostery	Blur	Lightbeam	Browser Config.	Private Browsing
Factor						
Trust	○	●	○	○	○	○
Lack of awareness	●	●	●	●	○	○
Lack of consent	○	○	○	○	○	○
Sharing with 1 st parties	○	○	○	○	○	○
Visit frequency	○	○	○	○	○	○
Has personal info	○	○	●	○	○	○
Has social info	○	●	○	○	○	○
Has search info	○	○	●	○	○	○
Has shopping info	○	○	○	○	○	○
Has financial, health info	○	○	○	○	○	○
Has correspondence	○	○	○	○	○	○
No volunteered info	○	○	○	○	○	○

Table 3. Summary of tools' ability to support situational factors on which users base decisions about tracking. ● indicates that a plugin or approach can account for differences in the factor; ○ indicates that it cannot.

Almost none of the current tools give easy ways for users to make privacy decisions based on preferences that involve situational factors. Instead, many tools focus on preventing harmful outcomes of tracking. However, this often requires blocking a particular outcome altogether, including the beneficial aspects of such outcomes. Some tools do not allow users to make fine-grained decisions about tracking. Adblock, for example, allows users to block ads, which limits some harmful outcomes of tracking, but does not allow the user to easily control blocking based on situational factors.

Other tools, notably Ghostery and Blur, allow a high degree of configuration, each capable of making automated decisions based on situational factors on behalf of the user. For example, Ghostery is capable of making automated decisions about social information, makes users aware of tracking, and provides users with information about the trustworthiness of third parties. Blur allows users to automate decisions to protect their personal information and search information in addition to making users more aware of tracking. Even so, many of the situational factors exposed in these tools do not align with our findings of what situational factors are important to users.

We primarily consider the degree of configurability that each plugin allows. There is a range of other trade-offs when designing privacy tools other than their degree of configurability. For example, Adblock Plus appeals

to some people primarily because it requires negligible configuration. Conversely, modifying browser settings to limit tracking also appeals to many because it does not require installing an addon. We do not address the usability of these plugins, which certainly also plays a role in meeting users' needs.

5.2 Design Recommendations

Designers of tools to manage tracking should be aware of the types of situational information that define privacy boundaries for users. In our study, we observed that the tracked information, properties of websites being tracked, and properties of tracking parties all defined boundaries for our participants. We discuss the extent to which these boundaries could be enforced by different means: Some boundaries, such as those that depend on properties of information on tracked sites, might be enforceable by a semi-automated user agent or via interaction with users; others, such as requiring more transparency, might require policy intervention. Our goal is to identify opportunities for online privacy tools to better enforce users' preferences.

We make the following recommendations:

- Automate the detection and enforcement of the most common preferences (e.g., those regarding personal information).
- Automated enforcement should take advantage of additional contextual information.
- When an automated agent may not be able to enforce preferences due to diverse or inconsistent preferences, it should judiciously prompt the user for a user-specific default.
- Provide methods for users to understand the effect online activities have on what information might be inferred about them.

Boundaries that could potentially be automatically detected by machine-learning tools include those defined by the type of information on the first-party site, e.g., whether the site contains personal information, financial or health information, and correspondence. For example, all our participants but one were concerned about a website learning personal information about them without them knowingly providing it. To address this concern, tools could detect when users enter their personal information into a web form and prevent tracking by sequestering cookies on that site from all other sites. It may be possible to similarly detect the presence of other information by using machine-learning tools which de-

tect the topic of a page, or by compiling a list of popular pages and the information they contain.

Technical solutions are sometimes not sufficient to mitigate all users' concerns. For example, whether the user has consented to tracking is difficult for plugins or browser-based solutions to detect. Similarly, our participants often did not realize trackers could infer more about users than was explicitly described in the tracked web pages. In such situations, tools might focus on improving user awareness and guiding users to manually decide whether to permit the tracking.

Yet other situations may require regulatory intervention. For example, neither a browser-based plugin nor a user can determine which trackers share information with which first-party websites in the absence of regulation to enforce and codify such relationships.

5.3 Comfort Prediction

In this section, we report on a preliminary examination of the feasibility of developing an automated user-agent that allows or blocks tracking based on the predicted comfort of the user. Specifically, we explored the use of several classifiers (including Asymmetric AdaBoost [42], Support Vector Machines [18], and Generalized Linear Mixed Effects Regression [10]) to predict a user's comfort with tracking of specific page visits based on properties of the web page and that user's demographics and general attitudes toward tracking; for brevity, we report only on the performance of Asymmetric AdaBoost, which performed best in our experiments.

We split our data (285 situations) into training and testing sets, trained a classifier on the training set, and then measured its accuracy at predicting comfort on the test set. Tracking situations in which participants are "maybe" comfortable are treated as if participants were uncomfortable, which is the conservative, safer option.

This prediction task is asymmetric: incorrectly predicting a user is comfortable with tracking when in fact they are not is more harmful than incorrectly predicting discomfort. Thus, we focus on reducing the more harmful error at the expense of the overall error rate. In training AdaBoost, we impose increased weights to "uncomfortable tracking" situations in order to decrease the false-positive rate (FPR), predicting an instance of uncomfortable tracking as comfortable. In consequence, the true-positive rate (TPR), correctly predicting an instance of tracking as comfortable, also decreases.

Our data set contains more features than would be easy to automatically detect (and hence base classifica-

tion on) in practice, for example, whether the user finds customization on websites beneficial. In addition, some features are impossible to measure with certainty, for instance, whether a tracker aggregates information from different sources. Consequently, we categorize features (independently by two researchers) into three overlapping sets: features that can be easily estimated or measured automatically; a superset of the previous features that also includes features that can be measured with a larger effort (e.g., by taking a crowdsourcing approach); and the set of all features. We then measure the performance of our classifier with each of these sets of features.

Figure 6 summarizes the performance of our classifier. To estimate average performance, each curve is computed by calculating the mean of results achieved on 20 random splits of the data into training and testing sets. Our results show that a user who wishes to automatically block *all* uncomfortable tracking (0% FPR) could simultaneously allow only 2-8% of the tracking that she is comfortable with. On the other hand, a user that is ready to permit tracking in a few situations in which she might be uncomfortable can allow 48-60% of the tracking she is comfortable with, while blocking over 90% of undesirable tracking. Interestingly, for FPR lower than 10%, the performance of the classifier that uses only the easily automatable features is just slightly below the performance of the other classifiers. This suggests that easily automatable features should suffice for the implementation of a conservative classifier that blocks most of the undesired tracking while allowing a significant portion of “good” tracking.

We investigate not penalizing the classifier for predictions made in “maybe” situations (i.e., when the participant also wasn’t sure); this makes it possible to allow over 60% of desirable tracking while incorrectly permitting less than 5% of undesirable tracking (see Appendix C). We also examine the classifier’s robustness to feature selection and the importance of specific features (also in Appendix C).

These results are preliminary and should be treated as a proof of concept. Nevertheless, they strongly suggest that automated tools that outperform current tools in enabling users to enforce their tracking preferences are an interesting direction of future research.

6 Conclusion

In this paper, we described the first in-depth investigation of users’ tracking preferences carried out in the

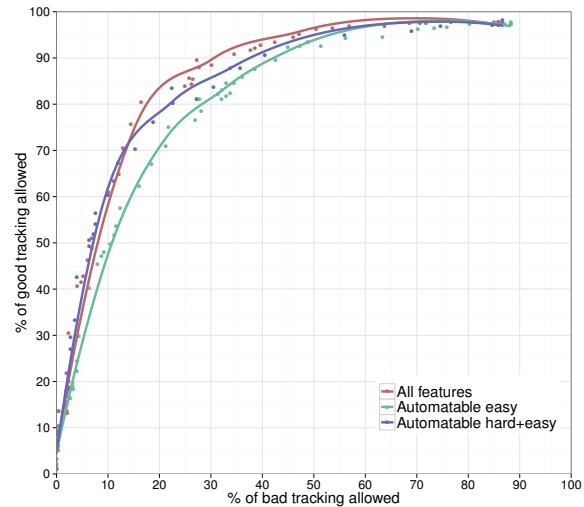


Fig. 6. ROC curves for predicting comfort

context of their own browsing histories, which we believe allowed us to more precisely capture participants’ concerns and comfort with tracking under specific circumstances than had previously been done. Using this methodology we studied the interplay between participants’ general attitudes about tracking, the outcomes of tracking they perceived in specific tracking situations, and the situational factors that guide their comfort or discomfort with specific instances of tracking.

Our examination both confirmed existing and provided specific novel insights; e.g., that users are less comfortable with the invisible outcomes of tracking (price discrimination, revenue for web sites, etc.) than with more noticeable outcomes (ads, customization, etc.), and that users commonly base their tracking preferences on specific properties of first-party websites, such as the topic of the site and frequency of visits.

In light of the perceived outcomes and the situational factors important to users, we examined a selection of current tracking tools, and found that they were rarely able to account for the situational factors important to our participants. We identified design guidelines for future tools, and showed experimentally that machine-learning tools can in many scenarios be an effective aid in implementing users’ preferences.

Acknowledgments

We would like to thank Samantha Gottlieb for her comments on early drafts of this work. This research was partially supported by NSF award CNS-1330596.

References

- [1] Adblock plus. <https://adblockplus.org/>.
- [2] Blur. <http://www.abine.com/>.
- [3] Ghostery. <https://www.ghostery.com/en/>.
- [4] Open directory project. <http://www.dmoz.org>. Accessed: Nov, 2014.
- [5] W3C Do Not Track Standard. <http://www.w3.org/TR/2015/WD-tracking-compliance-20150714/>.
- [6] Acceptable ads. <https://acceptableads.org/>, 2015.
- [7] AAAA, ANA, BBB, DNA, and IAB. *Self-regulatory principles for online behavioral advertising*. Digital Advertising Alliance, July 2009.
- [8] A. Acquisti. Privacy in electronic commerce and the economics of immediate gratification. In *Proc. EC. ACM*, 2004.
- [9] L. Awarwal, N. Shrivastava, S. Jaiswa, and S. Panjwani. Do not embarrass: Re-examining user concerns for online tracking and advertising. In *Proc. SOUPS*, 2013.
- [10] D. Bates, M. Maechler, B. M. Bolker, and S. Walker. lme4: Linear mixed-effects models using eigen and s4, 2014. ArXiv e-print; submitted to *Journal of Statistical Software*.
- [11] R. F. Baumeister, E. Bratslavsky, C. Finkenauer, and K. D. Vohs. Bad is stronger than good. *Review of general psychology*, 5(4):323, 2001.
- [12] H. Beales. The value of behavioral targeting. *Network Advertising Initiative*, 2010.
- [13] B. Berendt, O. Günther, and S. Spiekermann. Privacy in e-commerce: stated preferences vs. actual behavior. *Communications of the ACM*, 48(4):101–106, 2005.
- [14] M. Bilenko, M. Richardson, and J. Y. Tsai. Targeted, not tracked: Client-side solutions for privacy-friendly behavioral advertising. In *HotPETs*, 2011.
- [15] I. A. Bureau. IAB internet advertising revenue report, Apr. 2015.
- [16] R. Calo. Digital market manipulation. *George Washington Law Review*, 2013.
- [17] F. Chanchary and S. Chiasson. User perceptions of sharing, advertising, and tracking. In *Proc. SOUPS*, 2015.
- [18] C. Cortes and V. Vapnik. Support-vector networks. *Machine Learning*, 20(3), 1995.
- [19] A. Datta, M. C. Tschantz, and A. Datta. Automated experiments on ad privacy settings. In *Proc. PETS*, 2015.
- [20] A. Farahat and M. C. Bailey. How effective is targeted advertising? In *Proc. WWW*, 2012.
- [21] M. Fredrikson and B. Livshits. Repriv: Re-imagining content personalization and in-browser privacy. In *IEEE S&P*, 2011.
- [22] S. Greengard. Advertising gets personal. *Communications of the ACM*, 2012.
- [23] S. Guha, A. Reznichenko, K. Tang, H. Haddadi, and P. Francis. Serving ads from localhost for performance, privacy, and profit. In *HotNets*, 2009.
- [24] A. N. Joinson, U.-D. Reips, T. Buchanan, and C. B. P. Schofield. Privacy, trust, and self-disclosure online. *Human-Computer Interaction*, 25(1):1–24, 2010.
- [25] A. Lambrecht and C. Tucker. When does retargeting work? Information specificity in online advertising. *Journal of Marketing Research*, 2013.
- [26] P. Leon, B. Ur, R. Shay, Y. Wang, R. Balebako, and L. Cranor. Why Johnny can't opt out: A usability evaluation of tools to limit online behavioral advertising. In *Proc. CHI*, 2012.
- [27] P. G. Leon, B. Ur, Y. Wang, M. Sleeper, R. Balebako, R. Shay, L. Bauer, M. Christodorescu, and L. F. Cranor. What matters to users?: Factors that affect users' willingness to share information with online advertisers. In *Proc. SOUPS*, 2013.
- [28] M. Malheiros, S. Brostoff, C. Jennett, and M. A. Sasse. Would you sell your mother's data? personal data disclosure in a simulated credit card application. In *WEIS*, 2012.
- [29] M. Malheiros, S. Preibusch, and M. A. Sasse. "Fairly truthful": The impact of perceived effort, fairness, relevance, and sensitivity on personal data disclosure. In *TRUST*, 2013.
- [30] K. Martin. Addressing privacy online—Initial analysis and draft findings. Unpublished research presentation, 2014.
- [31] S. Panjwani, N. Shrivastava, S. Shukla, and S. Jaiswal. Understanding the privacy-personalization dilemma for web search: a user perspective. In *Proc. CHI*, 2013.
- [32] K. Purcell, J. Brenner, and L. Rainie. Search engine use 2012. Technical report, 2012.
- [33] E. J. Rader. Awareness of behavioral tracking and information privacy concern in Facebook and Google. In *Proc. SOUPS*, 2014.
- [34] F. Roesner, T. Kohno, and D. Wetherall. Detecting and defending against third-party tracking on the web. In *Proc. NSDI*, 2012.
- [35] F. Roesner, C. Rovillos, T. Kohno, and D. Wetherall. Balancing privacy and functionality of third-party social widgets. *USENIX Magazine*, 2012.
- [36] V. Toubiana, A. Narayanan, D. Boneh, H. Nissenbaum, and S. Barocas. Adnostic: Privacy preserving targeted advertising. In *Proc. NDSS*, 2010.
- [37] H. Treiblmaier and I. Pollach. Users' perceptions of benefits and costs of personalization. In *Proc. ICIS*, 2007.
- [38] C. Tucker. Social advertising. *SSRN eLibrary*, 2012. <http://ssrn.com/abstract=1975897>.
- [39] J. Turow. *The daily you: How the new advertising industry is defining your identity and your worth*. Yale University Press, 2012.
- [40] J. Turow, J. King, C. J. Hoofnagle, A. Bleakley, and M. Hennessy. Americans reject tailored advertising and three activities that enable it. *SSRN*, 2009.
- [41] B. Ur, P. G. Leon, L. F. Cranor, R. Shay, and Y. Wang. Smart, useful, scary, creepy: Perceptions of online behavioral advertising. In *Proc. SOUPS*, 2012.
- [42] P. A. Viola and M. J. Jones. Fast and robust classification using Asymmetric AdaBoost and a detector cascade. In *NIPS*, 2001.
- [43] C. E. Wills and M. Zeljkovic. A personalized approach to web privacy: Awareness, attitudes and actions. *Information Management & Computer Security*, 2011.
- [44] C. Wilson. If you use a Mac or an Android, e-commerce sites may be charging you more, Nov. 2014. <http://www.washingtonpost.com/posteverything/wp/2014/11/03/if-you-use-a-mac-or-an-android-e-commerce-sites-may-be-charging-you-more/>.
- [45] D. A. A. with Zogby Analytics. Poll: Internet Users Recognize the Importance of Online Advertising and the Value of Self-Regulation. <http://www.aboutads.info/ZogbyDAAOct13PollResults.pdf>.

A Interview Script

I'd like to record this interview, may I start now? Today we will talk about online tracking. Before we start, I want you to know that we will be asking you questions about your personal preferences. We are only interested in your opinions, we are not testing your knowledge.

1. Have you heard of the terms "online tracking?"
 - (a) *If heard of online tracking* What does online tracking mean to you?
 - (b) *If heard of online tracking* What benefits can online tracking offer you?
 - (c) *If heard of online tracking* What harms can online tracking cause to you?
 - (d) *If heard of online tracking* Are you aware of any online tracking technique?
 - (e) *If aware of techniques* Which online tracking techniques are you aware of?
 - (f) *If Not heard of online tracking* What is the first thing that comes up to your mind when you hear "online tracking?"

Article: I am going to let you read some information about online tracking. Then I am going to ask you a few questions to collect your opinion about it. Please note that the questions are not intended to measure your comprehension skills or memory. We are only interested in your opinions about online tracking. *[The text shown to participants is in Section A.1]*

1. What benefits do you believe online tracking can offer to you?
2. What harms do you believe online tracking can cause to you?
3. In general, how do you feel about being tracked online? Why?
4. *If negatively* what would make you more comfortable with online tracking?
5. Can you think of any particular situations in which being tracked would benefit you?
6. Can you think of any particular circumstances in which being tracked could cause harm to you?
7. Are you aware of any ways in which you can limit online tracking?
8. Have you taken any actions to limit online tracking?
9. *If yes* When have you taken this action?
10. Imagine that you could specify how companies can track the websites that you visit. Could you describe how you would specify this?
11. Do you have any additional comments?

Single web page preferences Next we will show you some of the web pages you have visited over the past few days, and we will ask you some questions about them. For each of these web pages that you visited, I'm going to ask you how you feel about tracking only on that website.

[For the following questions, participants will be shown only one web page they visited.]

1. Do you remember visiting this page?
2. *If No* Move on to next question
3. *If Yes* What do you use this website for?
4. Would it be acceptable to you if this website collects and aggregates your actions on this web page?
5. What benefits can you think of from site X logging and analyzing your visit to this web page?
6. What harms can you think of from site X logging and analyzing your visit to this web page?
7. What do you think site X would learn from logging and analyzing this visit?
8. *If answers affirmatively* How do you feel about site X learning that?
9. *If answers affirmatively* With whom do you think site X can share the information it learns about you?
10. *For each piece of information the participant mentions* How do you think site X could use this information about you?
11. In general, How comfortable are you if site X tracks your online activities? Why?

Pairs preferences [Web pages are referred to as pages X and Y below.] Now I will show you two web pages that you visited. Imagine that a tracking company is able to know that you have visited both the first and the second web page. I will be asking you questions about your preferences for tracking by a third party across these two pages.

1. Do you remember visiting these pages?
2. *If yes* Are these pages related?
3. *If yes* How are they related?
4. *If No* Skip to the next question
5. How comfortable are you if the tracking company tracks your online activities on these two web pages? Why?
6. What benefits can you think of from the tracking company knowing that you visited both of these two web pages?
7. *Possible follow up* What benefits can you think of from the tracking company logging and analyzing your visits to both of these two web pages over time?

8. What harms can you think of from the tracking company knowing that you visited both of these two web pages?
9. *Possible follow up* What harms can you think of from the tracking company logging and analyzing your visits to both of these two web pages over time?
10. What do you think this tracking company would learn from your two web page visits?
11. *Possible follow up* How do you feel about the third party learning that (ask this for each data type mentioned before)?
12. How do you think this tracking company could use this information about you (ask for each data type mentioned above)?
13. What do you think would be appropriate for the company to learn?
14. What do you think would be inappropriate for the company to learn?
15. How would you feel if the tracking company shared the information it learns from you with another third-party?
16. How comfortable would you be if the tracking company shares this information with site X? Why?
17. How comfortable would you be if the tracking company shares this information with site Y? Why?
18. In general, how comfortable are you if the tracking company tracks your online activities on these two web pages? Why?
19. Do you have any other comments that you'd like to share with us?

A.1 Information About Tracking

What is online tracking? Online tracking companies (a.k.a. online trackers) partner with websites to be able to track the activities of visitors on those websites. For example, a tracking company can know the web pages a user visits on a website, what links a user clicks on, how much time they spend on certain pages, what search terms they type, etc. Generally speaking, there are two types of online tracking. Tracking that is contained within a given website and tracking that expands across different websites. Online tracking is often imperceptible to users because online trackers operate “behind scenes” and there is normally no clear indication that trackers are present on websites.

What can be learned about users and how is it used? In general, online trackers are interested in learning users' interests, preferences, demographics, on-

line habits, interactions with website features, purchasing behaviors, and more. However, credit card numbers, Social Security numbers, passwords, and other sensitive personal information is normally out of reach of online trackers. Information that online trackers collect and aggregate may be used for different purposes. The ultimate use depends on who buys or otherwise can have access to information collected from users and on the intentions of the recipient.

Common uses include, but are not limited to:

- Analytics (e.g., understand how users interact with the website, analysis of website traffic, types of visitors, etc.)
- Personalized (a.k.a. targeted or tailored) website advertising
- Website customization (e.g., modifying the website design for a specific user or for the general audience of a website)
- Marketing (e.g., contacting a user to sell something they previously showed interest in)

A.2 Website Selection

Sensitive topics were defined as those which deal with financial services, medicine, health, file sharing, insurance or employment. This was chosen based on prior research [30], and based on a pilot study which found these types of sites to be particularly sensitive. Insensitive topics were any topics not fitting these criteria. Search websites were: google.com, bing.com, and yahoo.com. Shopping websites included any category from the open directory that includes shopping. Social networking sites were: facebook.com, twitter.com, linkedin.com, or websites whose topic in the open directory include *Online Communities*. Two websites were considered close in time if they were visited less than 30 minutes apart. Two websites were considered far away in time if they were visited more than 2 days apart. Single websites were selected such that no website was repeated in the first-party tracking section of the interview. Pairs of websites were selected such that no pair contained the same two websites. In addition, websites that satisfied one situation, would not be chosen for satisfying a different situation.

1 st party	1	Insensitive
	2	Shopping
	3	Search engines
	4	Sensitive
3 rd party	5	Two insensitive sites
	6	Shopping site and insensitive site
	7	Search engine and insensitive site
	8	Social networking site and insensitive site
	9	Two sensitive sites
	10	Two sites visited close together in time
	11	Two sites visited far away in time

Table 4. Criteria to select websites for the interviews.

B Plugin Evaluation

Evaluation Criteria Using data from our interviews, we focus on the ability of tools to help users make decisions based on situational factors and address the perceived outcomes of tracking.

Tools were marked as accounting for a situational factor if the tool has a control to toggle tracking based on that factor. Even though some tools in theory allow the user to perfectly control the tracking of, e.g., their search information if they use private browsing mode on every search, this would take a large amount of effort for users and they would likely be prone to neglecting to do this. When a tool can account for a situational factor only using such methods, we do not consider it to successfully account for that factor. We examine whether tools are able to account for the following situational factors based on our interviews.

- **Informational properties**—Can a tool be configured to control tracking based on the type of information on the site (e.g., some tools can turn off tracking on search engines)?
- **Search**—Can a tool block tracking specifically for search activities?
- **Has Correspondence**—Can a tool block tracking of the contents of users' email, instant messages, etc.
- **Lack of awareness**—Does a tool help make a user aware of tracking?
- **Trust**—Does a tool help the user decide whether a tracker can be trusted?

We also examine whether tools are able to prevent harms due to specific outcomes.

- **Targeted ads**—Tools were counted as addressing this outcome if they blocked any of the harmful aspects of advertising that participants listed (i.e.,

other people might see, they are annoying if repetitive).

- **Price discrimination**—This outcome was addressed by no tool, since the information used to change pricing could be revealed by any page visit (e.g., by the users' device) [44].
- **Customization**—Tools satisfied this requirement if they were able to limit the harms of customization independently of preventing other (potentially beneficial) outcomes.
- **Negative feeling**—We believed all tools might make the user feel as if they had more control over tracking.
- **Revenue**—Tools were counted if they were capable of disabling advertisements altogether; we do not consider tools that might indirectly affect company revenue by decreasing the value of a particular visitor's advertising profile to advertisers.
- **Persecution**—No tools seemed capable of limiting persecution by the government.
- **Linked**—We count any tool that is capable of limiting the ability of third parties to link information to the user as satisfying this requirement.

Tools and Evaluation We chose the following tools to balance both popular tools and those offering unique features to limit tracking. While this is not a complete list of such tools, we believe it captures a range of different features and approaches to specifying preferences for limiting tracking.

AdBlock is a plugin that blocks ads based on preconfigured lists that identify advertisements on web pages. AdBlock was not counted as supporting any situational factor because of its inability to allow users to configure tracking beyond a per-website basis. However, it does limit the harms people associate with targeted advertisements by blocking all ads.

Ghostery is a plugin that shows users which trackers are on a website and provides information about these trackers to help users decide whether to trust them. Users can configure the Ghostery plugin to block all or specific third parties or allow (but not block) all tracking on a particular first-party website. In addition, users are able to make decisions about particular kinds of trackers (e.g., advertising, analytics, beacons, privacy, or widgets).

Blur is a plugin that blocks third-party tracking, similar to AdBlock. It uses an indicator to alert users when third-party tracking occurs. It also allows users to limit tracking during web searches. Blur allows users to au-







Outcome	AdBlock 	Ghostery 	Blur 	Lightbeam 	Browser Config. 	Private Browsing 
Ads	●	●	●	●	●	●
Price	○	○	○	○	○	○
Customization	○	●	○	○	○	●
Feel "stalked"	●	●	●	●	●	●
Revenue	●	●	●	●	○	○
Persecution	○	○	○	○	○	○
Linked	○	●	●	○	○	●

Table 5. The extent to which different plugins are capable of ameliorating the outcomes that concern users. ● indicates that the plugin is fully capable of ameliorating the harms of this outcome and allowing benefits if configured to the users' wishes. ○ indicates that the plugin is unable to do this. ◐ indicates that the plugin is able to ameliorate the harms but not also allowing benefits.

tomatically “mask” some types of personal information, such as their email address or phone number.

Lightbeam is a plugin for the Firefox browser that shows how first-party websites are linked to third-party trackers. It seeks to provide transparency into how websites are tracked, and focuses on showing these links to users rather than controlling the tracking that does happen. Users are limited to blocking tracking globally, on specific first parties, or by specific third parties.

Browser configuration Most web browsers give users the option of blocking all third-party cookies via preference configuration. Browser configuration is capable of limiting the potential for data aggregation, but is not situationally configurable except through manual intervention, which it is not designed for.

Private browsing mode can allow users a high degree of situational avoidance of cookie-based tracking. However, many users might forget to use private browsing mode, or later wish that they had used it.

C Revisiting Comfort Prediction

A Different Perspective, Maybe? In Section 5.3 we showed a proof of concept for classifiers that predict the comfort of our participants in different tracking situations. There, we conservatively treated situations in which participants were “maybe” comfortable as situa-




Fig. 7. ROC curves for predicting comfort.

tions in which they were not (i.e., if the classifier predicted that users were comfortable, we counted this as incorrect). If we focus on the arguably more important web page visits for which users had definitive opinions, excluding “maybes” from testing, the classifier performs better, as shown in Figure 7.

Important Features To examine which of the features are most important for predicting user comfort, we ran a forward feature selection process with 20 random splits into validation and training sets. The top eight features are shown below, along with the average accuracy for predicting comfort using that and all higher-ranked features. The classifier performs very well even in the absence of the top eight features, achieving 78.57% accuracy with the remaining automateable features.

1. Participant's identity; one of 35 possibilities (67.14%).
2. Website type(s); one of 11 options as shown in Table 4 (71.96%).
3. “Has financial info”: whether the participant is concerned about her financial info being tracked (73.75%).
4. Whether the frequency of visit to the website(s) matters to the participant (74.64%).
5. “Has search info”¹ (75.54%).
6. “Has correspondence info”¹ (76.79%).
7. “Has shopping info”¹.
“Has social info”¹ (77.14%).²

¹ Analogous to item 3.

² Both factors needed to be added for classification accuracy to improve.