

Daniel Smullen*, Yaxing Yao, Yuanyuan Feng, Norman Sadeh*, Arthur Edelstein, and Rebecca Weiss

Managing Potentially Intrusive Practices in the Browser: A User-Centered Perspective

Abstract: Browser users encounter a broad array of potentially intrusive practices: from behavioral profiling, to crypto-mining, fingerprinting, and more. We study people’s perception, awareness, understanding, and preferences to opt out of those practices. We conducted a mixed-methods study that included qualitative (n=186) and quantitative (n=888) surveys covering 8 neutrally presented practices, equally highlighting both their benefits and risks. Consistent with prior research focusing on specific practices and mitigation techniques, we observe that most people are unaware of how to effectively identify or control the practices we surveyed. However, our user-centered approach reveals diverse views about the perceived risks and benefits, and that the majority of our participants wished to both restrict and be explicitly notified about the surveyed practices. Though prior research shows that meaningful controls are rarely available, we found that many participants mistakenly assume opt-out settings are common but just too difficult to find. However, even if they were hypothetically available on every website, our findings suggest that settings which allow practices by default are more burdensome to users than alternatives which are contextualized to website categories instead. Our results argue for settings which can distinguish among website categories where certain practices are seen as permissible, proactively notify users about their presence, and otherwise deny intrusive practices by default. Standardizing these settings in the browser rather than being left to individual websites would have the advantage of providing a uniform interface to support notification, control, and could help mitigate dark patterns. We also discuss the regulatory implications of the findings.

Keywords: Browsers, privacy, security, usability, settings, preferences, understanding, mental models, interaction design

DOI 10.2478/popets-2021-0082

Received 2021-02-28; revised 2021-06-15; accepted 2021-06-16.

***Corresponding Author: Daniel Smullen:** Carnegie Mellon University, E-mail: dsmullen@cs.cmu.edu

1 Introduction

The modern web browsing experience is trickier than ever before. As techniques such as machine learning, fingerprinting, profiling, and other forms of automated reasoning become increasingly pervasive, users may experience them nearly constantly during everyday Internet browsing. The application of these techniques can often provide users with improved, safer, and more relevant web experiences. However, they potentially become invasive when there is a mismatch between users’ expectations for privacy and security, and reality.

In this work, we use “potentially intrusive practices” (PIPs) to generically refer to common third-party tracking methods, as well as other types of scripts that run in the browser to collect data, monitor activity, redirect users’ attention, or operate in the background to gather something of value. We refer specifically to 8 categories of practices that fit this definition: identity/sign-in services, targeted advertising, behavioral profiling, reporting and analytics, fingerprinting, nag screens, session replay, and crypto-mining. Each of these PIPs can raise concerns associated with different dimensions of privacy captured by Solove’s taxonomy [64], which we detail in § 3.1. However, whether any of these practices are viewed as intrusive or not is determined by the personal perspective of the individual – this user-centric aspect is the subject of interest in our work. These 8 PIPs can offer clear benefits to websites and perceived benefits to users. Generally, websites increasingly employ profiling, reporting and analytics, and session replay to improve their products and services, increase business intelligence, and capitalize upon data broker-

Yaxing Yao: University of Maryland, Baltimore County, E-mail: yaxingyao@umbc.edu

Yuanyuan Feng: Carnegie Mellon University, E-mail: yuanyua2@cs.cmu.edu

***Corresponding Author: Norman Sadeh:** Carnegie Mellon University, E-mail: sadeh@cs.cmu.edu

Arthur Edelstein: Mozilla, E-mail: aedelstein@mozilla.com

Rebecca Weiss: Mozilla, E-mail: rweiss@mozilla.com

age [34]. Many websites use nag screens, crypto-mining, or targeted advertising to highlight new features, generate revenue from monetization, or make ads more relevant, respectively [5, 41]. Sign-in services and fingerprinting are used ostensibly for user convenience and security [67, 71]. However, PIPs are increasingly ubiquitous [61], and lack transparency – many users experience annoyance, frustration, fear, or feelings of being spied upon when they find out that they had been subjected to them without their consent [6, 19, 81].

The predominant approach to managing PIPs online is through notice and control [21]. Using a patchwork of settings (e.g., tracking protection [78]), built-in tools (e.g., private browsing mode), and third-party extensions (e.g., ad blockers), users have a variety of controls to align their browsing experiences with their privacy expectations. However, these controls often fail to achieve their goals due to users’ unwillingness to make the effort to use them, various usability issues, or misconceptions about how and when to use them [43].

Our work focuses on the notice and control made available by the browser itself rather than the ever-increasing array of third-party add-ons and tools. Often, add-ons require technical expertise and may not be intended for the average user [47] who may not understand the diverse implications of their use [40]. Outside of this tool-centered perspective, few settings are available in browsers or on websites for users to manage PIPs. Moreover, restricting PIPs using mechanisms that are not explicitly supported by websites can be fragile. Websites are constantly updated, and breakage can occur when their contents are manipulated. As a result, rather than risking breakage and losing users, many browsers’ default settings are limited and there is little that can be done to control PIPs [25, 51]. Of the few controls which are supported explicitly on websites, many involve redirecting users through complex third-party opt-out procedures [8].

Thus, our work answers the following research questions, each focused on a different aspect of design:

- RQ1 What are the signals that users rely on to determine whether they are being subjected to PIPs or not during browsing?
- RQ2 What affordances do users associate with allowing or restricting PIPs?
- RQ3 Are there PIPs that users want to control (e.g. allow/opt-in, or deny/opt-out), and subject to what factors?
- RQ4 What are users’ preferences to be notified about PIPs on different types of websites?

RQ5 How well can the existing settings capture users’ preferences, how often would they ideally need to be adjusted from the default, and what are the trade-offs associated with alternative settings?

In short, managing online PIPs effectively is a significant problem and would benefit from user-centered research; PIPs are complex, pervasive, and the extent to which users feel they have adequate notice and control over them is unclear. Moreover, prior research has shown that users’ privacy expectations are not currently fulfilled [8, 35, 36, 53, 79]. By modeling users’ expectations, understanding, and preferences, we propose ways to improve the settings offered by browsers and shed light on some of the potential implications of alternative designs for settings to manage PIPs.

Main Contributions

This work makes the following three main contributions:

1. This work contributes to an understudied niche and focuses on the user’s perspective, rather than on specific tracking technologies and tracking mitigation techniques. It provides new insights into the understanding, preferences, and expectations of users toward PIPs beyond the tool-centric approach seen in the prior art. Our user-centered approach exposes users’ various misunderstandings and misconceptions about practices on different websites (e.g., believing there is no PIPs present if ads are not present), which call for further attention from researchers and practitioners.
2. Our results reveal ways to address participants’ unfulfilled desire to be notified about and opt out of PIPs in different contexts, highlighting that their preferences can extend across categories of websites.
3. Our findings present opportunities to revisit the settings that browsers make available, characterize their accuracy and user burden trade-offs, and highlight new research challenges for these settings to be better aligned with users’ expectations. We emphasize the unique role browsers should play in centrally managing and enforcing preferences in a neutral manner, given the different incentives browsers have in contrast to website operators.

2 Related Work

In this section, we distinguish our work from prior research exploring understanding and attitudes towards potentially intrusive practices (PIPs), restriction mechanisms, opt-out and notification preferences.

2.1 User Modeling

Prior work in user modeling sought to elicit understanding and attitudes towards specific types of tracking, narrow definitions of PIPs, and in specific contexts. Individual understanding and expectations have been shown to vary [38, 58, 77], and these can impact behavior [13, 19, 66]. Attitudes towards PIPs on the web may also vary significantly across demographics and user populations. Studies have shown that some users encounter difficulty and usability issues managing tracking related to targeted ads [2, 6, 35, 43] and social media [22, 57]. Many other studies have investigated mental models and preferences toward tracking management and online tracking in general [10, 39, 81]. Some have focused on narrow and specific contextual threats, such as browser history [50], opt-out and data deletion on websites [8, 35, 36], privacy expectations on the Internet of Things [53, 79], and in mobile app contexts [44, 45]. However, broader definitions of PIPs in different browsing contexts still require further studies, such as those that we explore in our work. In our study, we focus on opt-out and notification preferences, as well as individual perspectives and expectations with the goal of improving browser affordances.

2.2 Browsers, Settings, and Tools

It is unclear whether current browser-based settings and user interfaces are seen as sufficient by users, particularly when considering PIPs holistically. Chrome, the most popular browser [46], currently offers only a small set of limited settings for managing privacy online, though a number of security-related settings for password management, in particular, are offered [33]. Browsers such as the Tor Browser [42, 56], Firefox [29], the Brave browser [37], and services such as the search engine DuckDuckGo [74] offer a variety of settings and provide alternative channels which offer extra security and privacy. More recently, some browsers have integrated new features to limit some PIPs. “Intelligent Tracking Protection” in Safari [78] blocks some “track-

ers” (which they define as third-party cookies and fingerprinting in some contexts) by default but does not offer additional settings to users. The default settings in Firefox block crypto-mining and provide information about PIPs [17].

Prior work has also studied specific features which are built-in to browsers, such as private browsing modes [1, 51], finding that users have unrealistic expectations about their effectiveness at limiting online tracking and providing security. Online privacy and security risks are rapidly evolving and pervasive, making them difficult for individual techniques to address [25, 61]. However, there have been standards-based solutions proposed by both researchers and practitioners to limit certain practices online. For example, Do-Not-Track (DNT), is an HTTP header field that signifies the user wants to opt-out of being tracked when enabled [16]. The Platform for Privacy Preferences (P3P) is a machine-readable disclosure so that the web browsers can automatically retrieve the policy and make users aware of PIPs [20]. Both DNT and P3P are not only technical solutions but also voluntary web standards recommended by the World Wide Web Consortium (W3C). Unfortunately, there is a historical reluctance for these standards to be accepted [12, 62], which renders them largely ineffective [21].

Another type of mechanism to limit PIPs is third-party browser extensions, a number of which have been developed into commercial products. For example, Ghostery and Disconnect.me are very popular tracking protections [23, 32], though not without flaws [14, 54]. When users visit websites, they render a list of “trackers” that are presented on privacy dashboards and allow the users to selectively block or unblock each one. However, these tools are largely intended for more technical users [14]. The literature has also uncovered a variety of misunderstandings about how users believe third-party browser extensions can protect them [30, 47, 48, 76]. A similar tool, Privacy Badger, adopts a slightly different approach and makes blocking decisions based on behavior across different websites [28]. These tools are rapidly evolving and their impact is still not fully understood, particularly because user perceptions of PIPs are not well understood. In this work, we limit our scope to improving what is offered by browsers in terms of built-in settings and affordances, acknowledging that there is a separate effort to improve browsers by way of third-party extensions. Rather than proposing new tools, our study is intended to address whether users understand their existing ability to control PIPs, and whether the available settings are cognitively associated

with enabling or disabling them (and in what contexts). We also study how they interpret the signals that they observe.

2.3 Notification Preferences

Privacy and security notifications, including those which incorporate nudging [3], are also a relevant area. It has been shown that notifications can be tailored in many contexts, such as authentication [55], and online social media [4]. Some forms of notification show promise to be tailored to users' specific decision-making and personality traits [73]. Other work has studied the design space of notifications in detail, such as "just in time" notices [60]. Some have studied ways to warn users about PIPs, particularly in the context of social media [72]. Prior work has also explored profiling to tailor notifications and controls [24, 75, 79, 80]. We distinguish our work from the prior art by sampling notification preferences about PIPs using a broad taxonomy, and among a variety of contexts. Our goal is to establish whether there is a connection between users' desire to be notified and their overall perception and understanding of PIPs.

3 Methodology

Our study employs a mixed-methods approach, incorporating both qualitative ($n = 186$) and quantitative surveys ($n = 888$) which were administered to separate groups of participants. This way, we were able to gather a rich set of qualitative perspectives and a large quantitative dataset of preferences from participants. Our surveys were contextualized to 8 different website categories, which we detail in § 3.1 below, and each survey presented one PIP to a participant. We describe and justify the 8 PIPs we chose to study in § 3.1.

Qualitative surveys underwent grounded analysis [70] to collect and categorize general themes – the code book can be found in the appendix. Quantitative surveys measured preferences to opt out of and be notified about PIPs. The corpus of opt-out preferences collected in our quantitative survey was also used to test alternative models of settings for managing PIPs in the browser. Both surveys are part of an IRB-approved protocol and incorporated attention checks to ensure data quality. Our full survey instruments are in the appendix.

3.1 Contextual Categories

Our study was designed to answer questions about a variety of browsing contexts. To do this, we identified 8 major categories of websites, as well as individual popular and less popular websites within each category. We used 8 categories from Alexa [7] which we believed were broadly representative, and selected the 1st (popular) and 500th (esoteric) within each category. The categories we selected were: News and Information, Entertainment and Games, Shopping, Travel, Finance, Adult, Health and Well-being, and Social Media and Blogging.

To capture holistic categories of practices, we created a novel taxonomy elicited from experts at Mozilla. In total, we cover 8 potentially intrusive practices (PIPs): identity/sign-in services (e.g. "sign in with Google"), targeted advertising, behavioral profiling (including associated predictions and data collection about users), reporting and analytics (focusing on technical data collection), fingerprinting, nag screens (which forcefully redirect the user), session replay, and crypto-mining. Each practice in the taxonomy is commonly encountered while browsing, is seen by experts to have some potential privacy and security problems based on Solove's taxonomy [64], and met our overarching definition of PIPs. Crypto-mining and nag screens may involve overtly invasive acts, redirecting computing resources and attention respectively. All 8 PIPs may involve some form of surveillance, and collected data may be involved in aggregation. Data collected through these PIPs also have the potential for insecurity or harm related to dissemination. In particular, data collected during behavioral profiling, reporting and analytics, and session replay may be subject to secondary uses. Fingerprinting may be used for identification (or de-anonymization).

To maximize construct validity, we developed internal technical PIPs definitions and non-technical PIPs descriptions for surveys that were consistent and simple. We used abstract categories of practices instead of specific privacy and security threats, to avoid biases against potentially beneficial aspects of practices. We chose to create descriptions that were suitable for laypersons to easily understand so that we were not limited by how well the average user could understand the technical specifics. Using a top-down brainstorming exercise, we listed candidates for categories of practices, wrote technical descriptions, and summarized the associated risks and benefits neutrally. Our taxonomy intentionally included categories of PIPs that are not necessarily mutually exclusive, such as behavioral profiling and fin-

gerprinting, which may often be closely associated with targeted advertising from a technical standpoint. We included these categories despite their potential overlap in order to tease out whether they were perceived differently by our participants. Each PIP is presented to separate participants, has separate descriptions, risks, and benefits, and is analyzed separately.

Neutral non-technical descriptions of PIPs intended for participants were iteratively refined. In the language used throughout the surveys, we always referred to PIPs as “web technologies” and avoided priming language, such as “intrusion”, “threat” or “attack”. Our descriptions were first piloted with two focus groups of non-technical employees at Mozilla. After each focus group, the text was modified based on the feedback. Clarifying details were added (e.g., fingerprinting is not referring to biometrics, giving specific examples of sign-in services) and priming language was eliminated wherever possible. Then, experts from our research team and external experts at Mozilla judged whether the corresponding PIPs opt-out scenarios were realistic and non-speculative. One PIP, paywalling, was initially considered (due to the potential for invasion [64]) but was later removed as it was controversial whether it would be possible to realistically opt out without circumventing websites’ legitimate business functions. The full list of PIPs descriptions, risks, and benefits can be found in the appendix.

3.2 Qualitative Survey

Our first survey focused on eliciting perceptions of PIPs that participants believed they had encountered, their attempts to control them, and associated experiences.

Recruitment was performed on Amazon Mechanical Turk, implemented and hosted using Qualtrics, with a combined consent form and pre-screening survey. 186 participants were recruited and compensated \$6 for the 20-minute average duration. Pre-screening required participants to affirm that they were over 18 years of age, resided in the US, regularly browsed the Internet, understood the consent form, and wished to participate voluntarily in research. Participants were each randomly assigned to a single PIP only. We report on our recruitment and dropout statistics in § 4.

The pre-survey was a free-listing exercise about the website categories described in § 3.1 in random order. Participants were instructed to look at their browsing history to find examples of websites that they would routinely visit if they did not immediately come to mind.

This exercise focused on popular websites to evoke examples that were representative of their categories and properly contextualize their responses. All participants were required to provide two examples from at least four out of the eight website categories.

The main survey was a qualitative survey with free-text responses. Participants were asked to describe the personal risks and benefits of their assigned “web technology”, and how they believed it might benefit companies who employ it. We then asked participants if and where they believed that they had encountered this “web technology” before. We also asked questions about how to protect themselves from the potential risks; whether they had attempted to opt out, how they approached this, and whether they had succeeded.

The post-survey asked participants about basic demographics; age range, gender, education, employment status, and city size. In addition, we administered the SA-6 questionnaire, a standard measure of security and privacy awareness [26]. Up until this point we had avoided using value-laden terms, such as “privacy” and “security”, but an exception was made in our post-survey because such terms are required as part of SA-6. SA-6 was used as a proxy for measures of technical aptitude in our analysis, and participants with higher values were considered more tech-savvy.

Analysis began with removing responses where participants did not pass attention checks. Next, Glaser’s grounded analysis was chosen to mitigate interpretation bias to systematically search for common themes [31, 70]. First-cycle coding identified general themes and trends. Several follow-up coding iterations were performed until saturation. Annotators were all usability, privacy, and security experts. Analysis occurred in unison, proscribing measures of inter-rater reliability [49]. The qualitative results were used to design a follow-up, large-scale, quantitative survey, as described below.

3.3 Quantitative Survey

The quantitative survey aimed to elicit the opt-out and notification preferences of browser users towards PIPs. Recruitment was performed using the same method and criteria as the initial qualitative survey, permitting only individuals who had not already participated in the previous survey. Participants were compensated \$3 for the 10-minute average duration.

Each participant was randomly assigned one PIP only. The second survey began with the neutral PIP description with associated risks, and benefits. This

ensured that all participants would have at least the same level of basic knowledge about their assigned PIP. Throughout the survey, we provided a link for participants to review the description, risks, and benefits. Next, each participant was presented with an example of a popular and unpopular website in each contextual category in random order. This was intended to help contextualize their responses to the category of websites. For each individual website within the category, participants were required to read about the category, when the website was established, the country it was based in, and a detailed screenshot of the website itself. Adult websites were censored to remove explicit content.

Next, participants were required to read hypothetical scenarios describing a novel mechanism for opting out of the PIPs, and respond to questions about their preferences to use the mechanism. Scenarios were each contextualized separately to individual websites, then to whole website categories, and then to all websites broadly. We used bold fonts to emphasize important details in questions, and made it clear that participants could reverse their choice to opt out if they desired. An example of a scenario contextualized to targeted ads on Amazon (the first-ranked Shopping website) follows:

Imagine that you are given a new setting in your browser that allows you to block Targeted Advertising on specific websites you choose (“opting out” of Targeted Advertising on these websites). When enabled, ads that use Targeted Advertising are blocked on websites you opt-out of. Ordinary ads which do not use Targeted Advertising are not affected by this setting. By default, you are still shown Targeted Advertising on websites you are not opted out of. Assume that you will be able to reverse this setting on any website, at any time. Consider Amazon, which is a Shopping website. How likely would you be to use the setting described above, to **opt-out of Targeted Advertising for Amazon?**

After individual websites, the participant was then asked questions about entire categories of websites, then all websites. PIPs descriptions can be seen in Table 6, risks and benefits in Table 5, and the opt-out scenarios for each PIP can be seen in Table 7 in the appendix along with the full survey texts.

Our qualitative survey showed that participants expressed difficulty in identifying the presence or absence of PIPs (see § 4.1). Therefore, at the end of our quantitative survey, we asked participants whether they would prefer to be notified about the presence (and/or absence) of the “web technology” as they browsed the various categories of websites. These questions specifically did not allude to any implementation details of the notifications – we were concerned with capturing

participants’ general perspectives, rather than testing a particular notification style. For each website category, participants could choose between one of “Notify me every time I visit”, “Notify me only once per week”, “Notify me only once per month”, “Notify me only the first time I visit”, or “Never notify me”, corresponding to ordinal levels of notification desire.

Post-surveys evaluated the participants’ SA-6 score and collected more detailed demographics; age, gender, marital status, education, employment, whether they worked and/or were educated in a STEM field, city size, when they last looked at and modified their browser privacy and security settings, their browser preference, and prior experience filling out privacy-related surveys online. We chose to examine these demographic factors as they had been previously shown to correlate with some privacy preferences, particularly opt-out choices, in prior work [63]. A final question was asked about whether the participant believed they belong to a category or group of individuals who are especially at risk, due to surveillance or some form of systematic oppression.

3.3.1 Regression Analysis

In order to answer RQ3, we needed to determine which demographic factors (e.g., SA-6 score, age, gender, etc.) and/or vignette factors (e.g., website category, individual websites, popularity) impacted participants’ expressed likelihood to opt out of PIPs. We performed regression analysis on our quantitative survey results to determine this. We used regression models as a way to systematically test which factors may have influenced participants’ likelihood to opt out; those which show statistically significant association with changes in opt-out likelihood across all PIPs would be suitable candidates for further testing in alternative settings, in part to answer RQ5. Likert-scale opt-out preferences were collapsed into binary categories (opt-out or opt-in) which served as the outcome variable for binomial generalized linear mixed-effects regression models. One regression was fitted for each PIP, so that they could be analyzed separately. Models were fit by maximum likelihood (Laplace Approximation) [11]. Demographic and vignette factors were modeled as fixed effects, and survey participants were each given a randomized unique identifier modeled as a random effect to account for individual variance between subjects. The 1st level of each fixed effect for ordinal factors was chosen as the intercept. Intercepts for categorical factors are described

in Table 9, along with the odds ratios (Z-test) and p-values.

Each of our regression models underwent model selection. One by one, each factor was added to the model and the resulting candidate model was tested against a null model (with only random effects) using likelihood ratio tests. If the likelihood ratio test showed with $p < 0.05$ that the model including the added factor was statistically significant versus the null model, the added factor was included. In cases where the added factor was not significant ($p > 0.05$), the factor was excluded. We added factors in the same order as is reflected in the header row of Table 1 until all factors were tested. As a final sanity test, we also tested for multicollinearity by reintroducing all factors (except gender and SA-6, which had already been universally excluded) into each regression and calculated the variance inflation factor. We did not find any evidence of moderate or high levels of correlation ($VIF > 5$) between any factors which had been previously included based on our likelihood ratio tests, and factors which showed high levels of correlation ($VIF > 10$) had already been excluded. We also explored whether interactions were present among factors included in the model, but no significant interactions were found.

With a theoretical maximum of 15 demographic and 3 vignette factors, *a priori* power analysis was performed using G*Power [27]. We required at least 997 participants to achieve significance at $\alpha = 0.05$ with a medium effect size (0.4). We report on the observed power in the results.

3.4 Testing Alternative Settings Models

The corpus of preferences collected in our quantitative survey was used to perform a series of simulation experiments which test alternative models of settings for managing PIPs in the browser. These experiments characterize accuracy (i.e., how many instances in which participants' preferences coincided with what is offered by the settings), as well as user burden (i.e., the number of actions participants would need to take to adjust individual settings in order to make what is offered coincide with their preferences). The parameters of the experiment were bounded by the 16 websites collected, spanning the 8 categories of websites, across all 8 PIPs. As such, we are simulating the effect of implementing the hypothetical settings which were introduced in our surveys in a highly constrained setting under conservative assumptions.

The experiments tested 6 different models: (1) No Toggle (closest to the current default in most browsers such as Chrome [33]) where all PIPs are allowed with no additional settings offered, then (2) No Toggle where all PIPs are denied with no additional settings. Next, (3) Default Allow and (4) Default Deny Category-level Toggles, where users can change their category-level preferences when they do not prefer the default to increase accuracy at the cost of additional burden but only based on website categories. Finally, (5) Default Allow and (6) Default Deny Individual Website Toggles, where users can change individual website preferences when they do not prefer the default to increase accuracy at the cost of additional burden across all individual websites. The experiments all assume that changing to different defaults requires one action to change each setting. One decision or changed setting amounts to one unit of user burden, accrued each time the user-preferred setting doesn't match the current default. Changing individual website/category settings requires one decision per individual website/category. Instances where users do not have consensus among categories (e.g., a user has equal numbers of allow and deny preferences within a single website category) do not result in a changed setting. Finally, we assume that we are only changing settings for one PIP at a time – a limitation imposed by our corpus being comprised of data for only one PIP per participant.

4 Results

We organize the results by survey and research questions. In § 4.1, we describe the qualitative survey results, addressing RQ1 and RQ2. In § 4.2, we present findings from the quantitative survey, addressing RQ3 and RQ4. Detailed demographics can be seen in Table 8 in Appendix. Finally, we report the results of our experiments characterizing accuracy and user burden trade-offs between alternative settings models (RQ5) in § 4.3.

4.1 Qualitative Analysis

The aim of our qualitative analysis was to categorize and organize themes in responses. We used these categories to isolate examples of signals that users rely on to determine the presence or absence of PIPs while browsing (RQ1), and the affordances that users associate with controlling PIPs (RQ2). We received 186 responses.

4.1.1 RQ1 - Unreliable Signals

We identified different signals that participants rely on to determine if there are being subjected to PIPs, including the presence of ads, changes in functionality (breakage), recognizing explicit notifications, and recognizing implicit notifications. Participants often told us about heuristics that they had developed to determine whether they are being subjected to (or protected from) PIPs.

Ten separate participants recognized that the presence of advertisements likely implied the presence of targeted ads and behavioral profiling, some participants referred to a lack of advertisements as a signal of not being subject to advertising-related PIPs. Another 21 instances among 19 participants showed that the presence of ads was also used as the signal for PIPs not explicitly related to ads, such as fingerprinting. One participant took note of ads that conflicted with their interests:

“I know that I’ve succeeded [in opting out of fingerprinting] when I see ads for things that I would never eat such as meat or burgers. That’s a very simple example but it tells me some of these sites have no idea what my preferences are because I would never eat animal products of any sort. (Pt. 6f1d71b7) ”

Many participants expressed confidence that they had successfully opted out because they did not see any ads:

“Well with the ad blocker or script blocker program I know [I successfully opted out] because I don’t receive any ads at all. And with the script blocker, I’m pretty sure the website isn’t receiving any information from me based off my limited knowledge of the program. Same goes for private browsing I guess. (Pt. 6d4a0d2e) ”

7 participants recognized connections between breakage and the effectiveness of their opt-out approach. We see evidence of participants using both ad-related signals and breakage in the following example:

“I turn [targeted advertising] off or avoid it on websites that I feel are sketchy. [...] The advertisements that I was shown were different and parts of the website stopped functioning properly. (Pt. 1a64f804) ”

These participants saw breakage as a signal of an effective approach. However, breakage can occur when PIPs are present. It is not definitive evidence of opting out.

99 participants speculated about the technical specifics of PIPs. We observed confusion about PIPs descriptions in focus groups and modified them accordingly. Our survey text specifically pointed out that fingerprinting did not refer to physical fingerprints or biometrics. However, 8 participants were especially con-

fused by fingerprinting, suggesting that the data collected could include authentication tokens, security keys, biometrics, or encourage identity theft.

44 participants purported personal benefits of PIPs, expressing approval and did not mention any risks. 6 participants mentioned that with nag screens, “attention can be drawn to important things”. Though the rest found them annoying, one remarked that their interest “often outweighs or sufficiently overshadows any nag screens. (Pt. 789bf237)” We observe that many of the signals participants rely on may be unreliable.

4.1.2 RQ2 - Incorrect or Missing Affordances

Regarding the affordances users associate with enabling or disabling PIPs, we noted browser settings, third-party tools, extensions, and settings on websites as the most prominent examples which were mentioned in responses. 54 participants in total saw security tools as the most appropriate way to avoid PIPs as well as unrelated threats. This phenomenon was seen among technically sophisticated participants in particular. However, a significant portion of participants (15 instances in 13 responses) asserted that using their browser “safely” alone ensured their safety. These participants were unable to articulate what their approach entailed in terms of specific actions, interfaces, or behaviors. In contrast, more technically sophisticated participants would often recommend specific products. One participant detailed their usage of virtual machines to avoid data collection associated with sign-in services and malware:

“[...] I browse using a VM (virtual machine, a cloned and contained version of my browser) when casually surfing the net or shopping. I use Shadow Defender. I can pick up all the trojans and malware I like, then with a click of a button, that “machine” is destroyed and my real computer is back in play. It’s great. (Pt. 4a171bc9) ”

While using VMs or specific anti-malware security tools is useful and can potentially help mitigate malware risks, they are not effective ways to opt out of PIPs. VMs can be effective at mitigating some forms of fingerprinting, but in general, they do little to mitigate data collection associated with other PIPs. Most notably, the approach is ineffective when users sign in.

There were 12 participants who believed that to opt out of PIPs related to targeted advertising (but not specifically targeted ads), such as behavioral profiling and fingerprinting, they must block ads. These participants mentioned ad-blockers and anti-virus tools as the best way to ensure total protection from the risks associ-

ated with advertising-related PIPs. “Malvertising” (ads with embedded links to malware) [65] was also mentioned as a specific concern by 12 participants. Of these, 5 confused tools such as anti-virus with ad-blocking, perhaps due to the potential to protect from malware.

Some participants suggested that PIPs were intended to improve security. 8 participants mentioned security benefits with fingerprinting, session replay, and identity/sign-in services. We noticed that these participants seemed to emphasize security over privacy concerns, and were primarily concerned about their online accounts being hacked. One participant mentioned that they saw fingerprinting as a way to provide “greater protections against fraud (Pt. 4dc574f7)”. 3 participants perceived session replay as a beneficial feature, confusing it with history or session cookies.

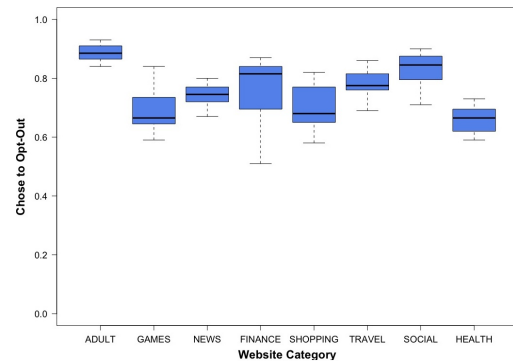
62 participants in total had assumptions about having control over a setting that does not exist. Of these, 53 participants’ responses alluded to settings on websites as ways they had previously used to opt out. Settings to control PIPs are provided on websites in myriad forms, such as cookie-blocking banners and privacy dashboards, though many websites (including the examples we surveyed) do not provide any meaningful settings. 20 participants believed they had control over their surveyed PIPs when in reality there were no settings built into their browser or on the websites they mentioned. Their accounts revealed that they did not take any action to opt out because they were confused about where to find the settings – they had assumed that the settings existed. These participants expressed that it was too difficult to configure these settings because they were unable to find them. The 33 other participants assumed that opt-out settings were available but never attempted to use them. One participant clearly knew where to look for privacy and security settings in their browser, but misinterpreted what was offered:

“I assume the [fingerprinting] features would have their own browser setting location, just like other browser features, where you can disable and enable them. For example, I use Chrome most often, so I would expect to find these features listed under Settings > Advanced > Privacy and Security. (Pt. 2e78b57d)”

Of these 33 participants, 16 assumed their browsers would offer nonexistent settings and insisted that they should be available without specifying where.

In summary, our qualitative results offer insights into users’ perceptions of PIPs. Many of the perceived affordances mentioned by participants are inadequate or non-existent, showing a relatively high level of miscon-

Fig. 1. Aggregate opt-out preferences, per website category.



ception about PIPs. We saw evidence of differences in participants’ technical sophistication, confusion about the risks and benefits of PIPs, and reliance on unreliable signals. These findings informed the development of our quantitative survey.

4.2 Quantitative Analysis

This section details the results from the quantitative data in our second survey. We answer two research questions regarding participants’ preferences to opt out of PIPs (RQ3) and their preferences to be notified (RQ4). Recall that in both surveys, participants were randomly assigned to one of the 8 PIPs we studied (listed in Table 2), so each participant responded to one assigned practice. 888 responses were collected in total, and we observed 92% power in post-hoc power analyses. On average, 111 responses were collected for each practice.

4.2.1 RQ3 - Users Want To Opt Out

Our quantitative data showed that participants want control over PIPs, either by opting in or out, subject to the various contextual factors we studied.

Table 2 shows participants’ opt-out preferences across all surveyed practices. In general, participants were opposed to most practices, wanting to opt out with little variance. Overall, participants preferred to opt out in 81% of instances on average. The outlook is somewhat different with preferences to opt out on a per website category basis as in Figure 1. While the majority of participants prefer to opt out in all website categories, there is some variability in preferences. For example, the preferences we collected about PIPs on finance websites is somewhat skewed towards preferring to allow. This is likely due to fingerprinting and sign-in

services being seen as beneficial here. This echoes our qualitative results that some participants saw security benefits associated with these practices.


4.2.2 RQ3 - Website Category Influences Preferences

In § 3.3.1, we describe our regression analysis used to determine if there was an association between the likelihood to prefer to opt out and the factors we surveyed. As shown in Table 1, many different factors seem to play a statistically significant role in opt-out likelihood for specific PIPs. Age range, education level, STEM education, city size, marital status, employment status, STEM employment, how recently the participant looked at and changed their settings, their browser of choice, whether they had recently participated in online surveys about privacy, and whether they self-identified as being at risk of privacy breach were all shown to be associated with changes in opt-out preferences for at least one PIP. The detailed odds ratios and p-values from our regression models (Z-test versus the intercept) for each PIP can be found in Table 9. We found that only website category was associated with changes in likelihood to opt out in all PIPs, and while other PIPs had factors which may be associated with opt-out likelihood, these findings were inconsistent between PIPs. For this reason, website category was the only factor that we found was appropriate to use in alternative settings models (see § 4.3). We also did not see evidence of interactions among factors.

4.2.3 RQ4 - Users Want To Be Notified About PIPs

RQ4 questions which practices users prefer to be notified about, how often, and in what contexts. We observed clear preferences to be notified about the presence and absence of PIPs on most websites. We expected that participants would prefer not to be notified in most circumstances, reflecting most contemporary browsers’ designs which do not notify without direct user engagement and notifications are very subtle if present at all. However, we see a clear trend towards the preference for notifications, summarized in Figure 2. We saw a similar trend in the responses broken down by website category, with the exception of Adult websites which were preferred to be notified about significantly more frequently. We view the desire to be notified more frequently as evidence of concern. Moreover, our qualitative results show that users are facing difficulty identifying where PIPs

Fig. 2. Aggregate PIPs notification preferences.



are present, and these results seem to provide further evidence.

Note that we deliberately chose not to study the implementation of a specific notification mechanism. Instead, we asked questions about notifications in the abstract. Some users may prefer some types of notifications over others or may find some so annoying that they would prefer to forgo them completely – our results show that there is simply an expectation that may not currently be adequately met.

4.3 Characterizing Alternative Settings

In this section, we show the results of our experiments exploring accuracy and user burden associated with different models of PIPs opt-out settings in browsers. Generally, more accurate settings are more desirable, as they fulfill the expressed preferences of users. Less burdensome settings are also more desirable, as they limit the distraction and annoyance associated with configuring these settings. Accuracy is calculated based on the percentage of individual websites correctly aligned between the settings and expressed preferences of each individual user collected in our surveys. User burden is calculated based on the number of instances where settings must be realigned, within the constraints of the model. Both accuracy and user burden are subject to the constraints of the models we tested. Recall that in § 3.4 we described the alternative settings models. Note that any time the settings are changed from the default, accuracy can increase but user burden is incurred. No Toggle has no settings and only the default applies. With Category Toggles, the default is applied but the settings can be changed per website category at the cost of additional burden. Website Toggles can be adjusted for each indi-

Table 1. Results of our model selection. Included factors are labeled ●. Excluded factors are labeled ○. Website category was the only factor found to consistently

	Age Range	Education Level	STEM Education	City Size	Marital Status	Employment Status	STEM Employment	Looked Settings	Changed Settings	Browser Used	Recent Surveys	At Risk	Gender	SA-6	Website Category
Behavioral Profiling	○	○	●	○	○	○	○	●	○	○	●	○	○	○	●
Reporting and Analytics	○	○	●	○	○	●	●	○	○	●	○	●	○	○	●
Session Replay	○	○	●	○	○	○	○	○	○	○	●	●	○	○	●
Targeted Ads	●	○	●	○	●	○	○	●	●	●	○	○	○	○	●
Crypto-Mining	○	●	○	●	○	○	●	○	○	○	○	●	○	○	●
Identity/Sign-in	●	○	○	○	○	○	○	○	●	○	○	○	○	○	●
Fingerprinting	○	●	○	○	○	○	○	○	○	○	○	○	○	○	●
“Nag” Screens	○	○	○	○	○	○	○	●	○	○	○	○	○	○	●

Table 2. Mean opt-in (Allow) and opt-out (Deny) preferences.

	Prefer Allow	Prefer Deny
Behavioral Profiling	20%	80%
Reporting and Analytics	20%	80%
Session Replay	18%	82%
Targeted Ads	19%	81%
Crypto-Mining	14%	86%
Identity and Sign-In Services	18%	82%
Fingerprinting	23%	77%
“Nag” Screens	17%	83%

vidual website. While the language used in our surveys referred to “opting out” of a practice, for clarity we refer to these settings as “allow” (i.e., to opt in) or “deny” (i.e., to opt out).

4.3.1 RQ5 - Deny By Default Is Less Burdensome

Table 3 describes the accuracy of the different settings models. Here we see that the settings which can offer the greatest accuracy are Category Toggles and Website Toggles, while No Toggle are far less accurate. Our results reveal the trade-off between accuracy and user burden inherent in the number of settings that are offered. This relationship is evident when comparing Table 4 with Table 3. However, deny by default settings based on website categories or individual websites require fewer changes for users to achieve their preferred settings compared to allow by default – between 75% and 65% fewer actions required on average. This is consistent with our findings that users broadly prefer to opt out (Table 2), and website categories are a significant factor in opt-out likelihood (Table 1).

In our simulation, the upper bound of user burden is limited by the number of surveyed website categories (8 total) and websites (16 total, 2 per category). Therefore, more choices are possible when settings are offered based on individual websites rather than categories. We found that both Category and Website Toggles spanned

the entire range of possible choices for individual users, requiring zero changes in the best case and 16 in the worst case (Table 4). Website Toggles offer the best accuracy but are more burdensome even when the setting is allowed. In contrast, Category Toggles can provide very high accuracy with minimal user burden. The optimal trade-off will depend on the specific user, but we speculate that the middle ground would be appropriate for most, suggesting that Category Toggles would be best.

5 Discussion

A significant body of research has shown that users’ privacy expectations are not currently fulfilled by websites or apps [48, 53, 58]. One goal of this work is to determine the extent to which users are aware of this, and another is to investigate ways to address the gap between privacy expectations and reality. Our findings are consistent with prior studies on specific privacy and security threats, particularly those studying tracking, profiling, and certain fingerprinting methods [25], and targeted ads [43]. Our novel contribution informs redesigned browser settings to accurately limit PIPs where preferable and better accommodate users’ expectations.

Similar to prior work, our participants demonstrated a strong desire to be notified about and opt out of PIPs on most websites. Existing solutions, however, do not meet people’s needs for web privacy. Naive solutions—such as attempting to block all PIPs completely—can be too clumsy to be practical as websites offer different sets of controls based on their business goals and privacy policies. Broader efforts—such as new data privacy laws—have emerged to aid in the standardization and uniform requirement for certain privacy options online, restricting specific practices without informed consent in some jurisdictions [9, 69]. The extent to which these efforts have been and will be effective is still unclear [68]. Many other approaches employed by

Table 3. Accuracy of the various alternative setting models. No Toggle (Allow by default) reflects Chrome settings.

	No Toggle	No Toggle	Category Toggles	Category Toggles	Website Toggles	Website Toggles
<i>Default Setting</i>	Allow	Deny	Allow	Deny	Allow	Deny
Profiling	25.8%	74.2%	92.3%	92.3%	100.0%	100.0%
Reporting	27.9%	72.1%	91.9%	91.9%	100.0%	100.0%
Session Replay	24.5%	75.5%	92.7%	92.7%	100.0%	100.0%
Targeted Ads	24.6%	75.4%	90.0%	90.0%	100.0%	100.0%
Crypto-Mining	19.6%	80.4%	95.7%	95.7%	100.0%	100.0%
Identity Services	25.6%	74.4%	90.7%	90.7%	100.0%	100.0%
Fingerprinting	33.6%	66.4%	89.9%	89.9%	100.0%	100.0%
“Nag” Screens	24.5%	75.5%	92.1%	92.1%	100.0%	100.0%
Mean	25.8%	74.2%	91.9%	91.9%	100.0%	100.0%

Table 4. User burden is the average number of settings changed per user, per PIP. No Toggle (Allow by default) is considered the current setting, with zero burden. The maximum possible burden is 16 (given 2 websites in 8 categories).

	No Toggle	No Toggle	Category Toggles	Category Toggles	Website Toggles	Website Toggles
<i>Default Setting</i>	Allow	Deny	Allow	Deny	Allow	Deny
Profiling	0.00	1.00	5.32	1.45	11.87	4.13
Reporting	0.00	1.00	5.12	1.58	11.54	4.46
Session Replay	0.00	1.00	5.45	1.37	12.08	3.92
Targeted Ads	0.00	1.00	5.23	1.17	12.06	3.94
Crypto-Mining	0.00	1.00	6.09	1.23	12.86	3.14
Identity Services	0.00	1.00	5.21	1.30	11.91	4.09
Fingerprinting	0.00	1.00	4.50	1.88	10.63	5.37
“Nag” Screens	0.00	1.00	5.40	1.33	12.08	3.92
Mean	0.00	1.00	5.29	1.41	11.88	4.12

websites result in dark patterns or do not provide any meaningful options [43, 59]. It is also well known that users can be manipulated into choosing settings that do not match their privacy preferences [4] and few websites adopt privacy standards that are considered best practices. [12, 62]. Our work provides insight into novel, better approaches to improve web privacy experiences.

We proceed as follows: first, we remark on our findings about users’ perspectives and expectations (§ 5.1). Next, we discuss the findings of our quantitative studies (§ 5.2). Finally, we discuss policy implications (§ 5.3), limitations (§ 5.4), and future work (§ 5.5).

5.1 Perspectives on Notices and Controls

Our findings confirm that most people find the PIPs we studied to be intrusive and expect to have the ability to restrict them. We found that people want to opt out in most instances (81% of our participants) independent of context. However, many participants assumed that they had the ability to opt out even in cases where such controls are not available. For example, most websites do not offer controls to opt out from data collection associated with behavioral profiling, or reporting and

analytics, yet as described in § 4.1.2 our participants repeatedly asserted that they have some way of controlling these settings but could not precisely articulate what they were. This suggests there is a baseline expectation or assumption which is not meshing with reality.

We also see evidence that users are not looking in the right places to find the settings and signals they need. We recognize that some browsers are incorporating new dashboards and menus which enable users to see whether various practices are present using subtle visual cues and icons. Recently, Firefox and Brave Browser have incorporated icons near the URL bar which can be clicked and expanded to see more details about the presence of practices similar to those seen in our work [52]. Our study found that participants desire to be notified about PIPs in most contexts, but testing the variety of possible notification designs and alternatives is beyond the scope of our work.

As in prior studies [1, 81], our participants suffered from misconceptions about the efficacy of various protections against PIPs, as well as difficulty identifying their presence or absence. Particularly, our qualitative data shows that users are likely to be misled by the signals they observe, placing greater importance on these

signals than other more objective signals. Our results elaborate the circumstances that users may assume that they are being protected simply by blocking ads, meanwhile, unmitigated risks abound. When ads are absent, our participants believed that they were broadly restricting PIPs. Though ad-blocking may be effective at restricting nag screens and targeted ads, it is risky for users to assume that is also the case with fingerprinting, sign-in services, cryptomining, reporting and analytics, or behavioral profiling. However, replacing ad-related signals may be a potential opportunity for a targeted intervention [4, 73]. Effective notifications could redirect users away from misleading signals or explain why some signals are not reliable. In cases that users have already opted out of advertising, browsers could replace ads with notices, or inform users about other unrestricted PIPs.

5.2 Context-Sensitive Settings

Our results show that context (especially website categories) impacts the likelihood to opt out, as seen in Table 1. Preferences to be notified about PIPs appear to vary depending on the context and the practice itself (seen in Figure 2). Conversely, our experiments with alternative settings show that the settings that best balance accuracy with user burden are based on website categories. Though the majority want to opt out of PIPs generally (refer to Table 2), simply denying all PIPs by default while offering no further settings can only achieve accuracy ranging from 66% to at most 80% (refer to Table 3). In other words, one-size-fits-all No Toggle preferences that deny PIPs by default will satisfy many users, but not all. Users would be better accommodated by offering more specific settings which are based on allowing or denying categories of PIPs individually. Alternative settings which are also able to be changed based on website categories can offer 90% to 96% accuracy. Even though this alternative is more complex, settings based on a small number of categories (such as those in our study) are less burdensome while retaining high accuracy. Most importantly, settings should be “deny by default” to achieve high accuracy and low user burden.

There are natural dimensions (such as PIPs and website categories) across which people’s preferences seem to be consistent; if this wasn’t the case, settings provided in the browser would not be useful as people’s preferences would be too diverse to capture with any degree of accuracy without being maximally burdensome. This common perception may also apply to

the categorization of websites; some websites may belong to multiple categories according to some users and not others. We show that the sensitivity individuals associate with categories of websites may also vary. Adult and finance websites may be considered more sensitive.

In summary, we contribute to the existing work by highlighting some concrete ways to improve the settings offered by browsers to help users manage PIPs online. Where we find that users’ preferences are consistent across a small number of dimensions, we show that it is possible to reduce user burden and maintain accuracy by enabling people to specify their preferences accordingly. Determining the most optimal categorization scheme for websites is an important avenue of future work. However, what authority should determine the categorization is potentially a matter of public policy, which we discuss in the section that follows.

5.3 Policy Implications

The current fragmented approach to PIPs management across browsers and websites exposes users to disparate ways of addressing the same preference management and consent issues across different websites. Some browsers and websites have baffling arrays of highly granular settings, while others lack any meaningful settings at all. Our findings lend support for the need for some standardization to give users the ability to control PIPs without imposing unrealistic burden on them. Such standardization would ideally include support for a minimum set of allow/deny settings that users could configure in their browsers, with individual websites being required to honor these settings, as conveyed by browsers to visited websites (e.g., denying specific practices with which the user does not feel comfortable). This would not be dissimilar to the Do Not Track standard developed by W3C [16], except that this standardization would differentiate between different PIP practices. In the simplest case, each PIP could have a single allow/deny setting in the browser, set to deny by default, and with users able to toggle each of these settings to align them with their preferences. Our results indicate that such configuration would significantly reduce the burden on users, saving them the effort of looking for and using equivalent settings on individual websites – to the extent that such settings even exist. This would also empower users to block PIP practices with which they are not comfortable.

A slightly more sophisticated approach, which would further increase accuracy, would involve support-

ing browser settings that do not just differentiate between PIPs but also differentiate between different categories of websites. One would have to ensure that websites do not attempt to defeat such a mechanism by attempting to mask their true category. Browsers would likely be able to defeat most of these attacks by relying on Alexa categories [7], or possibly relying on natural language processing techniques to develop linguistic fingerprints of website categories (e.g., techniques relying on topic modeling) [15, 18]. One would also have to ensure that websites do not attempt to force users to agree to PIPs with which they are not comfortable by simply breaking when users select settings intended to block them. Forcing sites to honor PIP settings specified by users in their browsers without unnecessary breakage would most likely require regulatory support.

User awareness of and control over PIPs would also be enhanced if websites were required to send to the browser standardized notifications about the PIPs running on them. Browsers would then be able to communicate the presence of such practices in a standardized fashion to users, who could in turn adjust their browser-based PIP settings, as desired. While earlier attempts to standardize some practices within this space such as DNT have failed [50], we believe that studies such as the one presented in this paper provide scientific evidence for the need for such standardization. Today, users simply lack practical mechanisms to know what PIPs might be present on websites they visit, and lack practical mechanisms to restrict these practices. The authors believe that the recent development of more stringent privacy regulations such as GDPR or CCPA/CPRA provide a context where the introduction of standards such as the ones outlined above is arguably more within reach than it was when standards such as P3P or DNT were proposed [16, 20].

5.4 Limitations

Like any study, our methodology is subject to some limitations and threats to validity, which we summarize in this section. Our qualitative data collection and analysis method has limitations to generalizability. The resulting corpus of quantitative preferences used in our experiments are similarly limited, and our results may not scale linearly to the number of website categories or individual websites a user typically encounters compared to our sample. Our vignette study incorporates several factors, including website categories. While our work relies on an external source of truth for these cate-

gories, specifically Alexa categories [7], there is room for open interpretation about the most optimal categorization scheme. Individual browser vendors may formulate their own schemes as well. This does not fundamentally alter the concept or impact of controls which are contextualized to website categories.

We crowd-sourced our surveys using Amazon Mechanical Turk, whose users are not completely representative of the broader population, and whose expressed preferences do not completely coincide with actual decisions made by users in-situ. While not fully representative, our respondents' demographics did not differ drastically from general US census categories (Table 8) and were not likely to have been more technically savvy than average based on SA-6, education, or employment.

We have tried to make our methodology as realistic as possible by contextualizing participants' responses using priming exercises. Wherever possible, our surveys excluded language that would unintentionally prime participants or introduce bias. We also piloted our studies several times at each stage, validating them with both focus groups and experts. We rejected abusive and low-effort responses by analyzing the timing of responses and employing attention check questions.

5.5 Future Work

Our work has highlighted several potential areas for future work, which we summarize in this section. The first area is addressing further technical challenges of implementing effective notifications and controls which users would prefer. One problem with offering additional settings is balancing them against the potential to be overly burdensome. We argue that the existing patchwork of settings across websites, browsers, and extensions is already burdensome. There is potential to improve upon this by studying specific types of controls and notifications, which allow more introspection into the state of the browser. What is yet unknown is whether there are more or less important contextual factors which influence users' preference to be notified, or to opt out of PIPs. There may also be different ways to frame and organize these factors which should be explored. What is the best way to notify users about PIPs? What contextual factors other than those we highlight should be considered in opt-out and notification preferences? Assuming browsers implement the changes we propose, what is the most effective approach to encourage users to engage when configuring their settings? Is

it possible to assist users in configuring large numbers of settings, without being overly burdensome?

Future work should also address the potential breakage that may occur if more aggressive protections are enabled. This is still an area of uncertainty for browser vendors and one that could probably benefit from stronger regulatory requirements, given the tendency of many websites to break for no good reason when users attempt to block some PIPs (e.g., ad blocking). Independently of whether such regulatory requirements materialize, what level of breakage would be tolerated by users in order to achieve better guarantees about their protection from PIPs, and in what contexts? How do browsers move beyond prior voluntary standards, and how should new standards be defined? These questions could serve as a starting point for future research in this area and could help guide the development of the next generation of browsers that are better at empowering users to regain control over PIPs.

6 Conclusion

In this work, we used a mixed-methods approach to answer fundamental questions about the understanding, preferences, and expectations of users towards potentially intrusive practices (PIPs) online. By focusing on user-centered perspectives, we provide new insight within a broader scope than the prior art – we included broader categories of invasive practices beyond advertisement-related tracking, specific techniques for fingerprinting, or specific examples of other tracking technologies seen in prior literature. We surveyed responses across categories of websites, introducing new contexts which we show to have an impact on user preferences. We also observed relationships between preferences to opt out, categories of PIPs, and different categories of popular and unpopular websites. Most users strongly preferred to opt out of PIPs generally, but were unable to express their preferences even though they expected to be able to, and many became resigned. We show that most users also want to be notified about the presence or absence of PIPs, but cannot reliably determine whether they are present or not and instead rely on signals which are potentially misleading.

During our exploration, we uncovered pervasive misunderstandings and misconceptions around PIPs. This allowed us to suggest concrete design changes which have the potential to address gaps in current browser settings, which we discuss in § 5.1. We also char-

acterized the relationship between accuracy and user burden for a variety of alternative settings models, comparing them to what are offered by popular browsers. Our results show that the settings and defaults that users would need to have the necessary awareness of and control over PIPs are currently missing. One-size-fits-all settings are inaccurate, and the patchwork of settings offered by individual websites is unsatisfactory. Browsers which support deny by default and support granular PIP settings that differentiate between different categories of websites offer the prospect of significantly enhancing user control without imposing undue burden on users. While actions can already be taken to opt out of a limited set of PIPs on websites, particularly using plug-ins [8], users would benefit from a more systematic and usable approach to controlling PIPs through the provision of standardized settings made available in their browser. In addition, we want to draw regulatory attention to intentional breakage, namely websites that intentionally break when users attempt to opt out, as a way of discouraging them from doing so. In principle, browsers (at least browsers that do not benefit from PIPs) could act as neutral actors that empower users to effectively restrict PIPs with which they are not comfortable – in contrast to websites, which generally benefit from the PIPs they host or implement.

Acknowledgments

This study was supported in part by grants from DARPA and AFRL under the Brandeis project on Personalized Privacy Assistants (FA8750-15-2-0277), by grants from the National Science Foundation Secure and Trustworthy Computing program (CNS-15-13957, CNS-1801316, CNS-1914486) and by an unrestricted grant from Mozilla. The authors would like to also thank Dr. Steven Engelhardt (Mozilla) and Prof. Alessandro Acquisti (CMU) for their helpful suggestions on this work. The US Government is authorized to reproduce and distribute reprints for Governmental purposes not withstanding any copyright notice. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the NSF, DARPA, AFRL or the US Government.

References

- [1] Ruba Abu-Salma and Benjamin Livshits. Evaluating the end-user experience of private browsing mode. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, CHI '20, page 1–12, New York, NY, USA, 2020. Association for Computing Machinery.
- [2] Jagdish Prasad Achara, Javier Parra-Arnau, and Claude Castelluccia. Mytrackingchoices: Pacifying the ad-block war by enforcing user privacy preferences, 2016.
- [3] Alessandro Acquisti. Nudging privacy: The behavioral economics of personal information. *IEEE Security and Privacy*, 7(6):82–85, 2009.
- [4] Alessandro Acquisti, Idris Adjerid, Rebecca Balebako, Laura Brandimarte, Lorrie Faith Cranor, Saranga Komanduri, Pedro Giovanni Leon, Norman Sadeh, Florian Schaub, Manya Sleeper, Yang Wang, and Shomir Wilson. Nudges for privacy and security: Understanding and assisting users' choices online. *ACM Comput. Surv.*, 50(3), August 2017.
- [5] Alessandro Acquisti, Curtis Taylor, and Liad Wagman. The economics of privacy. *Journal of Economic Literature*, 54(2):442–92, June 2016.
- [6] Lalit Agarwal, Nisheeth Shrivastava, Sharad Jaiswal, and Saurabh Panjwani. Do not embarrass: Re-examining user concerns for online tracking and advertising. In *Proceedings of the Ninth Symposium on Usable Privacy and Security*, SOUPS '13, New York, NY, USA, 2013. Association for Computing Machinery.
- [7] Amazon. Alexa top sites. <https://www.alexametrics.com/topsites>, 2020.
- [8] Vinayshekhar Bannihatti Kumar, Roger Iyengar, Namita Nisal, Yuanyuan Feng, Hana Habib, Peter Story, Sushain Cherivirala, Margaret Hagan, Lorrie Cranor, Shomir Wilson, Florian Schaub, and Norman Sadeh. Finding a choice in a haystack: Automatic extraction of opt-out statements from privacy policy text. In *Proceedings of The Web Conference 2020*, WWW '20, page 1943–1954, New York, NY, USA, 2020. Association for Computing Machinery.
- [9] Catherine Barrett. Emerging trends from the first year of eu gdpr enforcement. *Scitech Lawyer*, 16(3):22–25,35, Spring 2020.
- [10] Lemi Baruh, Ekin Secinti, and Zeynep Cemalcilar. Online Privacy Concerns and Privacy Management: A Meta-Analytical Review. *Journal of Communication*, 67(1):26–53, 01 2017.
- [11] Douglas Bates, Martin Mächler, Ben Bolker, and Steve Walker. Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, 67(1):1–48, 2015.
- [12] P. Beatty, I. Reay, S. Dick, and J. Miller. P3p adoption on e-commerce web sites: A survey and analysis. *IEEE Internet Computing*, 11(2):65–71, 2007.
- [13] Annika Bergström. Online privacy concerns: A broad approach to understanding the concerns of different groups for different uses. *Computers in Human Behavior*, 53:419–426, 2015.
- [14] Dan Bouhnik and Golan Carmi. Interface application comprehensive analysis of ghostery. *International Journal of Computer Systems*, 5(3), 03 2018.
- [15] Renato Bruni and Gianpiero Bianchi. Website categorization: A formal approach and robustness analysis in the case of e-commerce detection. *Expert Systems with Applications*, 142:113001, 2020.
- [16] Bill Budington, Alexei Miagkov, Katarzyna Szymielewicz, and Jason Kelley. Do not track. <https://www.eff.org/issues/do-not-track>, 2016.
- [17] Dave Camp. Firefox now available with enhanced tracking protection by default plus updates to facebook container, firefox monitor and lockwise. <https://blog.mozilla.org/blog/2019/06/04/firefox-now-available-with-enhanced-tracking-protection-by-default/>, Jun 2019.
- [18] Michelangelo Ceci and Donato Malerba. Classifying web documents in a hierarchy of categories: a comprehensive study. *Journal of Intelligent Information Systems*, 28(1):37–78, 2007.
- [19] Hongliang Chen, Christopher E. Beaudoin, and Traci Hong. Securing online privacy: An empirical test on internet scam victimization, online privacy concerns, and privacy protection behaviors. *Computers in Human Behavior*, 70:291 – 302, 2017.
- [20] Lorrie Cranor and Rigo Wenning. Platform for privacy preferences (p3p) project. <https://www.w3.org/P3P/>, Feb 2018.
- [21] Lorrie Faith Cranor. Necessary but not sufficient: Standardized mechanisms for privacy notice and choice. *JTHTL*, 10:273–308, 2012.
- [22] Tobias Dienlin and Sabine Trepte. Is the privacy paradox a relic of the past? an in-depth analysis of privacy attitudes and privacy behaviors. *European Journal of Social Psychology*, 45(3):285–297, 2015.
- [23] Disconnect. Take back your privacy. <https://disconnect.me/>, 2020.
- [24] Serge Egelman and Eyal Peer. The myth of the average user: Improving privacy and security systems through individualization. In *Proceedings of the 2015 New Security Paradigms Workshop*, NSPW '15, page 16–28, New York, NY, USA, 2015. Association for Computing Machinery.
- [25] Steven Englehardt and Arvind Narayanan. Online tracking: A 1-million-site measurement and analysis. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, CCS '16, page 1388–1401, New York, NY, USA, 2016. Association for Computing Machinery.
- [26] Cori Faklaris, Laura Dabbish, and Jason I. Hong. A self-report measure of end-user security attitudes (sa-6). In *Proceedings of the Fifteenth USENIX Conference on Usable Privacy and Security*, SOUPS '19, page 61–77, USA, 2019. USENIX Association.
- [27] Franz Faul, Edgar Erdfelder, Axel Buchner, and Albert-Georg Lang. Statistical power analyses using g*power 3.1: Tests for correlation and regression analyses. *Behavior research methods*, 41:1149–60, 11 2009.
- [28] Electronic Frontier Foundation. Privacy badger automatically learns to block invisible trackers. <https://privacybadger.org/>, 2020.
- [29] Mozilla Foundation. Firefox - protect your life online with privacy-first products. <https://www.mozilla.org/en-US/firefox/>, 2020.
- [30] Xianyi Gao, Yulong Yang, Huiqing Fu, Janne Lindqvist, and Yang Wang. Private browsing: An inquiry on usability and

- privacy protection. In *Proceedings of the 13th Workshop on Privacy in the Electronic Society, WPES '14*, page 97–106, New York, NY, USA, 2014. Association for Computing Machinery.
- [31] Barney G Glaser and Anselm L Strauss. *Discovery of grounded theory: Strategies for qualitative research*. Routledge, 2017.
- [32] Cliqz International GmbH. Ghostery makes the web cleaner, faster and safer! <https://www.ghostery.com/>, Feb 2020.
- [33] Google. Choose your privacy settings. <https://support.google.com/chrome/answer/114836>, 2021.
- [34] Peiqing Guan and Wei Zhou. Business analytics generated data brokerage: Law, ethical and social issues. In Robin Doss, Selwyn Piramuthu, and Wei Zhou, editors, *Future Network Systems and Security*, pages 167–175, Cham, 2017. Springer International Publishing.
- [35] Hana Habib, Sarah Pearman, Jiamin Wang, Yixin Zou, Alessandro Acquisti, Lorrie Faith Cranor, Norman Sadeh, and Florian Schaub. “it’s a scavenger hunt”: Usability of websites’ opt-out and data deletion choices. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems, CHI '20*, page 1–12, New York, NY, USA, 2020. Association for Computing Machinery.
- [36] Hana Habib, Yixin Zou, Aditi Jannu, Neha Sridhar, Chelse Swoopes, Alessandro Acquisti, Lorrie Faith Cranor, Norman Sadeh, and Florian Schaub. An empirical analysis of data deletion and opt-out choices on 150 websites. In *Fifteenth Symposium on Usable Privacy and Security (SOUPS 2019)*, Santa Clara, CA, 2019. USENIX Association.
- [37] Brave Incorporated. Brave: Secure, fast & private web browser with adblocker. <https://brave.com/>, 2020.
- [38] Ruogu Kang, Stephanie Brown, Laura Dabbish, and Sara Kiesler. Privacy attitudes of mechanical turk workers and the u.s. public. In *10th Symposium On Usable Privacy and Security (SOUPS 2014)*, pages 37–49, Menlo Park, CA, July 2014. USENIX Association.
- [39] Ruogu Kang, Laura Dabbish, Nathaniel Fruchter, and Sara Kiesler. “my data just goes everywhere:” user mental models of the internet and implications for privacy and security. In *Eleventh Symposium On Usable Privacy and Security (SOUPS 2015)*, pages 39–52, Ottawa, July 2015. USENIX Association.
- [40] Soroush Karami, Panagiotis Illia, Konstantinos Solomos, and Jason Polakis. Carnus: Exploring the privacy threats of browser extension fingerprinting. In *Proceedings of the Symposium on Network and Distributed System Security (NDSS)*, 2020.
- [41] Eunjin Kim and Byungtae Lee. E-service quality competition through personalization under consumer privacy concerns. *Electronic Commerce Research and Applications*, 8(4):182 – 190, 2009. Special Issue: Economics and Electronic Commerce.
- [42] Linda Lee, David Fifield, Nathan Malkin, Ganesh Iyer, Serge Egelman, and David Wagner. A usability evaluation of tor launcher. *Proceedings on Privacy Enhancing Technologies*, 2017(3):90 – 109, 2017.
- [43] Pedro Leon, Blase Ur, Richard Shay, Yang Wang, Rebecca Balebako, and Lorrie Cranor. Why johnny can’t opt out: A usability evaluation of tools to limit online behavioral advertising. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '12*, page 589–598, New York, NY, USA, 2012. Association for Computing Machinery.
- [44] Jialiu Lin, Shahriyar Amini, Jason I Hong, Norman Sadeh, Janne Lindqvist, and Joy Zhang. Expectation and purpose: understanding users’ mental models of mobile app privacy through crowdsourcing. In *Proceedings of the 2012 ACM conference on ubiquitous computing*, pages 501–510, 2012.
- [45] Jialiu Lin, Bin Liu, Norman Sadeh, and Jason I Hong. Modeling users’ mobile app privacy preferences: Restoring usability in a sea of permission settings. In *10th Symposium On Usable Privacy and Security (SOUPS 2014)*, pages 199–212, 2014.
- [46] Awio Web Services LLC. Web browser market share. <http://www.w3counter.com/globalstats.php?year=2021&month=1>, Jan 2021.
- [47] Arunesh Mathur, Jessica Vitak, Arvind Narayanan, and Marshini Chetty. Characterizing the use of browser-based blocking extensions to prevent online tracking. In *Fourteenth Symposium on Usable Privacy and Security (SOUPS 2018)*, pages 103–116, Baltimore, MD, August 2018. USENIX Association.
- [48] Aleecia McDonald and Jon M Peha. Track gap: Policy implications of user expectations for the ‘do not track’ internet privacy feature. In *39th Research Conference on Communication, Information and Internet Policy, (TPRC 2011)*. Elsevier, 2011.
- [49] Nora McDonald, Sarita Schoenebeck, and Andrea Forte. Reliability and Inter-rater Reliability in Qualitative Research: Norms and Guidelines for CSCW and HCI Practice. *Proceedings of the ACM Human-Computer Interaction*, pages 1–23, August 2019.
- [50] 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. *Proceedings on Privacy Enhancing Technologies*, 2016(2):135 – 154, 2016.
- [51] G. Merzdovnik, M. Huber, D. Buhov, N. Nikiforakis, S. Neuner, M. Schmiedecker, and E. Weippl. Block me if you can: A large-scale study of tracker-blocking tools. In *2017 IEEE European Symposium on Security and Privacy (EuroS P)*, pages 319–333, 2017.
- [52] Mozilla. Enhanced tracking protection in firefox. <https://support.mozilla.org/en-US/kb/enhanced-tracking-protection-firefox-desktop>, 2021.
- [53] Pardis Emami Naeini, Sruti Bhagavatula, Hana Habib, Martin Degeling, Lujo Bauer, Lorrie Faith Cranor, and Norman Sadeh. Privacy expectations and preferences in an iot world. In *Thirteenth Symposium on Usable Privacy and Security (SOUPS 2017)*, pages 399–412, Santa Clara, CA, July 2017. USENIX Association.
- [54] Rishab Nithyanand, Sheharbano Khattak, Mobin Javed, Narseo Vallina-Rodriguez, Marjan Falahrestegar, Julia E. Powles, Emiliano De Cristofaro, Hamed Haddadi, and Steven J. Murdoch. Ad-blocking and counter blocking: A slice of the arms race. *CoRR*, abs/1605.05077, 2016.
- [55] Eyal Pe’er, Serge Egelman, Marian Harbach, Nathan Malkin, Arunesh Mathur, and Alisa Frik. Nudge me right: Personalizing online nudges to people’s decision-making styles. *SSRN Electronic Journal*, 01 2019.

- [56] The Tor Project. The tor project: Privacy & freedom online. <https://www.torproject.org/>, 2020.
- [57] Emilee Rader. Awareness of behavioral tracking and information privacy concern in facebook and google. In *10th Symposium On Usable Privacy and Security (SOUPS 2014)*, pages 51–67, 2014.
- [58] Ashwini Rao, Florian Schaub, Norman Sadeh, Alessandro Acquisti, and Ruogu Kang. Expecting the unexpected: Understanding mismatched privacy expectations online. In *Twelfth Symposium on Usable Privacy and Security (SOUPS 2016)*, pages 77–96, Denver, CO, June 2016. USENIX Association.
- [59] Iskander Sanchez-Rola, Matteo Dell’Amico, Platon Kotzias, Davide Balzarotti, Leyla Bilge, Pierre-Antoine Vervier, and Igor Santos. Can i opt out yet? gdpr and the global illusion of cookie control. In *Proceedings of the 2019 ACM Asia Conference on Computer and Communications Security, Asia CCS ’19*, page 340–351, New York, NY, USA, 2019. Association for Computing Machinery.
- [60] Florian Schaub, Rebecca Balebako, Adam L. Durity, and Lorrie Faith Cranor. A design space for effective privacy notices. In *Eleventh Symposium On Usable Privacy and Security (SOUPS 2015)*, pages 1–17, Ottawa, July 2015. USENIX Association.
- [61] Sebastian Schelter and Jérôme Kunegis. On the ubiquity of web tracking: Insights from a billion-page web crawl, 2016.
- [62] Michael Simon. Apple is removing the do not track toggle from safari, but for a good reason. <https://www.macworld.com/article/3338152/apple-safari-removing-do-not-track.html>, Feb 2019.
- [63] Daniel Smullen, Yuanyuan Feng, Shikun Aerin Zhang, and Norman Sadeh. The best of both worlds: Mitigating trade-offs between accuracy and user burden in capturing mobile app privacy preferences. *Proceedings on Privacy Enhancing Technologies*, 2020(1):195 – 215, 01 Jan. 2020.
- [64] Daniel Solove. A taxonomy of privacy. *University of Pennsylvania Law Review*, 154:477, 2005-2006.
- [65] Aditya K Sood and Richard J Enbody. Malvertising—exploiting web advertising. *Computer Fraud & Security*, 2011(4):11–16, 2011.
- [66] Janice Y. Tsai, Serge Egelman, Lorrie Cranor, and Alessandro Acquisti. The effect of online privacy information on purchasing behavior: An experimental study. *Information Systems Research*, 22(2):254–268, 2011.
- [67] R. Upathilake, Y. Li, and A. Matrawy. A classification of web browser fingerprinting techniques. In *7th International Conference on New Technologies, Mobility and Security (NTMS)*, pages 1–5, 2015.
- [68] Christine Utz, Martin Degeling, Sascha Fahl, Florian Schaub, and Thorsten Holz. (un) informed consent: Studying gdpr consent notices in the field. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*, pages 973–990, 2019.
- [69] Paul Voigt and Axel Von dem Bussche. The eu general data protection regulation (gdpr). *A Practical Guide, 1st Ed., Cham: Springer International Publishing*, 2017.
- [70] Diane Walker and Florence Myrick. Grounded theory: An exploration of process and procedure. *Qualitative Health Research*, 16(4):547–559, 2006. PMID: 16513996.
- [71] R. Wang, S. Chen, and X. Wang. Signing me onto your accounts through facebook and google: A traffic-guided security study of commercially deployed single-sign-on web services. In *2012 IEEE Symposium on Security and Privacy*, pages 365–379, 2012.
- [72] Yang Wang, Pedro Giovanni Leon, Alessandro Acquisti, Lorrie Faith Cranor, Alain Forget, and Norman Sadeh. A field trial of privacy nudges for facebook. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI ’14, page 2367–2376, New York, NY, USA, 2014. Association for Computing Machinery.
- [73] Logan Warberg, Alessandro Acquisti, and Douglas Sicker. Can privacy nudges be tailored to individuals’ decision making and personality traits? In *Proceedings of the 18th ACM Workshop on Privacy in the Electronic Society, WPES ’19*, page 175–197, New York, NY, USA, 2019. Association for Computing Machinery.
- [74] Gabriel Weinberg. Duckduckgo: Privacy, simplified. <https://duckduckgo.com/>, 2020.
- [75] Pamela J. Wisniewski, Bart P. Knijnenburg, and Heather Richter Lipford. Making privacy personal: Profiling social network users to inform privacy education and nudging. *International Journal of Human-Computer Studies*, 98:95 – 108, 2017.
- [76] Yuxi Wu, Panya Gupta, Miranda Wei, Yasemin Acar, Sascha Fahl, and Blase Ur. Your secrets are safe: How browsers’ explanations impact misconceptions about private browsing mode. In *Proceedings of the 2018 World Wide Web Conference, WWW ’18*, page 217–226, Republic and Canton of Geneva, CHE, 2018. International World Wide Web Conferences Steering Committee.
- [77] Yaxing Yao, Davide Lo Re, and Yang Wang. Folk models of online behavioral advertising. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing*, pages 1957–1969, 2017.
- [78] Maciej Zawadzinski. What is intelligent tracking prevention (itp)? versions 1.0 - 2.3 explained. <https://clearcode.cc/blog/intelligent-tracking-prevention/>, Apr 2020.
- [79] Shikun Zhang, Yuanyuan Feng, Lujo Bauer, Lorrie Faith Cranor, Anupam Das, and Norman Sadeh. Did you know this camera tracks your mood? understanding privacy expectations and preferences in the age of video analytics. *Proceedings on Privacy Enhancing Technologies*, 2021(2):282–304, 2021.
- [80] Shikun Zhang, Yuanyuan Feng, Anupam Das, Lujo Bauer, Lorrie Faith Cranor, and Norman Sadeh. Understanding people’s privacy attitudes towards video analytics technologies. *Proceedings of FTC PrivacyCon*, pages 1–18, 2020.
- [81] Yixin Zou, Kevin Roundy, Acar Tamersoy, Saurabh Shintre, Johann Roturier, and Florian Schaub. Examining the adoption and abandonment of security, privacy, and identity theft protection practices. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, CHI ’20, page 1–15, New York, NY, USA, 2020. Association for Computing Machinery.

Appendix

Surveys and Accompanying Texts

For the purpose of sharing our survey instruments, as well as sharing the full descriptions of all PIP included in the study as well as their accompanying risks, benefits, and scenario text which were used in the surveys, we have made our study artifacts public.

Both of our surveys provided descriptions of PIP and accompanying risks and benefits at the beginning and also throughout surveys, accessible via a prominently placed button. The text which was provided to participants containing the descriptions of each PIP are seen in Table 6. The accompanying risks and benefits are seen in Table 5. All surveys begin with a standard consent and screening form which has not been included.

Survey 1 - Qualitative

Section 1

In the following section, you will be asked to provide examples of websites which you routinely browse, based on a number of categories. The questions concerning each category will be presented in random order. The categories are as follows:

- News and Information
- Entertainment and Games
- Shopping
- Travel
- Finance
- Adult
- Health and Wellbeing
- Social Media and Blogging

Later in the survey, we will be asking you questions which use the examples you provide us to set the context.

Section 2

If you are unable to provide 2 examples, or are uncomfortable with providing 2 examples for a category, you may proceed and examples will be provided for you. **Please note that if you do not provide 2 examples for at least 4 out of the 8 categories, you will**

be automatically withdrawn from the study. *[The website categories which follow are presented in random order.]*

Take a moment to think of two **News and Information** pages you have visited, or which you browse routinely. These are websites which can include news papers, online journals, Wikis, and any other source of news or information. If it helps, check your browsing history and see if you can find good examples. Enter the names of the two websites you think of into the fields below. **Please only enter the names of the websites into the fields below.** *[Participant is presented with free text entry fields.]*

Take a moment to think of two **Entertainment and Games** pages you have visited, or which you browse routinely. These are websites concerning digital, print, online, and other forms of media and entertainment, including videogames, movies, gambling, and more. If it helps, check your browsing history and see if you can find good examples. Enter the names of the two websites you think of into the fields below. **Please only enter the names of the websites into the fields below.** *[Participant is presented with free text entry fields.]*

Take a moment to think of two **Shopping** pages you have visited, or which you browse routinely. These are websites where you can purchase goods and services online, and browse for items you wish to purchase. If it helps, check your browsing history and see if you can find good examples. Enter the names of the two websites you think of into the fields below. **Please only enter the names of the websites into the fields below.** *[Participant is presented with free text entry fields.]*

Take a moment to think of two **Travel** pages you have visited, or which you browse routinely. These are websites concerning booking travel and accommodations, travel planning, reviews, hotels, and more. If it helps, check your browsing history and see if you can find good examples. Enter the names of the two websites you think of into the fields below. **Please only enter the names of the websites into the fields below.** *[Participant is presented with free text entry fields.]*

Take a moment to think of two **Finance** pages you have visited, or which you browse routinely. These are websites which include trading, online banking, finan-

Table 5. Risks and benefits associated with each PIP which were presented as part of surveys.

PIP Name	Risks	Benefits
Identity/Sign-In Services	-This could be used to track you across many websites that may not be related -Allows inference of personal details that may be used for purposes other than logging in	+Don't need to remember as many passwords +Don't need to re-enter personal information or your account username and password with every new website you log in to
Targeted Advertising	-Data collected by advertisers may be used in ways you didn't anticipate, and for purposes other than advertisements	+Ads you are shown may be more relevant to your interests
Behavioral Profiling	-Facts may be inferred about you which are sensitive, or may make you feel uncomfortable -In some jurisdictions, profiles can be bought and sold and you have no rights to them	+May enable websites to improve products and services that they offer to you
Session Replay	-May reveal sensitive information, or information in a sensitive context	+May enable websites to improve products and services that they offer to you
Reporting and Analytics	-May reveal personal information, or information in a sensitive context	+May enable websites to improve products and services that they offer to you
Fingerprinting and Deanonimization	-Can prevent you from remaining anonymous, by identifying you even when you've taken steps to hide your identity (e.g., after you've cleared cookies or used the privacy mode in a browser)	+May enable websites to offer better security features, which can protect your account and account information
"Nag" Screens	-Can prevent you from accessing content, even in the middle of reading it	+May help websites ensure that their business meets regulatory requirements in some jurisdictions +May help to ensure that the website earns enough revenue to continue operating
Crypto-Mining	-Can negatively affect the performance of your device, which can also disrupt your browsing experience	+May enable websites to improve your browsing experience +May enable websites to remove ads or give you access to premium content

cial advice, market-related information, and more. If it helps, check your browsing history and see if you can find good examples. Enter the names of the two websites you think of into the fields below. **Please only enter the names of the websites into the fields below.** *[Participant is presented with free text entry fields.]*

Take a moment to think of two **Adult** pages you have visited, or which you browse routinely. These are websites which include sexually (or otherwise) explicit materials, including videos, photos, and other material not intended for consumption by minors. If it helps, check your browsing history and see if you can find good examples. Enter the names of the two websites you think of into the fields below. **Please only enter the names of the websites into the fields below.** *[Participant is presented with free text entry fields.]*

Take a moment to think of two **Health and Wellbeing** pages you have visited, or which you browse routinely. These are websites which concern medical, spiritual, dietary, and other forms of advice and discussion for the betterment of your physical, mental, and spiritual

health. If it helps, check your browsing history and see if you can find good examples. Enter the names of the two websites you think of into the fields below. **Please only enter the names of the websites into the fields below.** *[Participant is presented with free text entry fields.]*

Take a moment to think of two **Social Media and Blogging** pages you have visited, or which you browse routinely. These are websites which belong to social media networks, blogs, or other forms of online social interaction. If it helps, check your browsing history and see if you can find good examples. Enter the names of the two websites you think of into the fields below. **Please only enter the names of the websites into the fields below.** *[Participant is presented with free text entry fields.]*

Section 3

The next part of the survey is intended to collect information about your thoughts and experiences with a

Table 6. Descriptions used for PIP in survey instruments. Participants were surveyed about only one.

PIP	Definition Provided To Participants
Identity/Sign-In Services	Identity/Sign-In Services help you log in to websites without relying on passwords specific to these websites. Examples are “Log in with Google”, “Log in with Facebook”, and “Sign in with your Apple ID”. These services save you the effort of creating and remembering passwords for individual websites. Because they see the websites you access, these services might be able to infer details about you, such as your interests, education, income, and more. This information could be used for purposes that go beyond helping you log in.
Targeted Advertising	Targeted Advertising uses information collected about you to tailor the advertisements that are shown to you on a particular website.
Behavioral Profiling	Behavioral Profiling collects information about who you are, your interests, and the things you do, to categorize you into specific categories (or profiles). For example, a website might try to determine your age, whether you are an “impulse buyer,” your political beliefs, and potentially much more. Sometimes, the profiles can be incorrect. The use of Behavioral Profiling does not necessarily mean that you will be subjected to advertisements, but it does mean that information may be collected and inferred about you.
Reporting and Analytics	Reporting and Analytics monitors what is happening as you are browsing websites, and generates technical information for the website developers. Often, this includes information about the state of your device, browser, and may also include technical information about what happened during your interaction with a particular website. This can help websites improve their products and services, but can potentially reveal sensitive information.
Session Replay	Session Replay creates detailed logs that record the actions you take while browsing a particular website and sends these logs to the website owners. This means that website owners can observe and replay exactly what you did and what you saw. Note that this is not a feature that enables you, as the person browsing a website, to replay what you did. Sometimes, sensitive information can be found in these recordings because it wasn’t properly removed.
Fingerprinting and Deanonymization	Fingerprinting and Deanonymization is a technique which ensures that the website you are browsing always recognizes you, even if you are not signed in. Fingerprinting and Deanonymization also enables websites to detect whether a device that the website doesn’t recognize is interacting with the website. This can be useful for a variety of reasons, such as detecting when someone tries to access an account on a new or unrecognized device. This technique does not mean that the website is using biometrics (i.e. a fingerprint scanner) to identify you, and the technique has nothing to do with physical fingerprints. Rather, Fingerprinting and Deanonymization refers to ways that your device can be picked out and recognized among others.
“Nag” Screens	Nag Screens can force you to see a popup, to watch an ad, to prevent you from viewing content, or otherwise to do something that disrupts your normal browsing experience. Sometimes “Nag” Screens appear when you’re using an ad-blocker, or because the website needs you to interact with something, such as giving consent where required by law.
Crypto-Mining	Crypto-Mining uses your device to generate digital cash, such as Bitcoin, during the time you spend browsing a particular website. Generally this digital cash is sent to the owners of the websites, but in some circumstances you may get a share. Some websites use Crypto-Mining as a way of using your device to make money for the website instead of (or in addition to) advertisements. Since it uses your device’s processing power to work, Crypto-Mining uses electricity or battery power on your device, and can affect device performance when you browse websites that employ Crypto-Mining.

particular web technology. Please read the text on the following page carefully. After, you will be asked a series of questions. *[Participant is presented with the description of one PIP.]*

Prior to this survey, had you ever encountered any examples of *[PIP]*? *[Participant may choose between: Yes/No]*

What do you think the **risks** might be for you when you browse a website with *[PIP]*? *[Participant is presented with a free text entry field.]*

What do you think the **benefits** might be for you when you browse a website with *[PIP]*? *[Participant is presented with a free text entry field.]*

Here are some examples of concrete risks and benefits

associated with *[PIP]*: *[Participant is presented with the risks and benefits for the PIP.]*

Section 4

In this section, you will be asked about *[PIP]* in a variety of scenarios. You can hover over the (i) symbol to remind you of the definition of *[PIP]*. Opting out of *[PIP]* means that: *[Participant is presented with the specific PIP scenario.]*

[This form of question repeats for all of the specific websites, and website categories that the user provided in

Table 7. Opt out scenarios for each PIP, which were provided as part of surveys.

Tracker Name	Opt-Out Scenario (Specific Websites)	Opt-Out Scenario (All Websites)
Identity/Sign-In Services	Imagine that you are given a new setting in your browser that enables you to block Identity/Sign-In Services on specific websites you choose ("opting out" of Identity/Sign-In Services on these websites), requiring you to log in to these specific websites manually instead. This also requires you to log in to these specific websites separately. When enabled, the buttons and links to use Identity/Sign-In Services on the specific websites you opt out from are removed from the websites you opt out of, and the ability for these services to collect data is also removed on these websites. Assume that you will be able to reverse this setting on any website, at any time.	Imagine that you are given a new setting in your browser that enables you to block Identity/Sign-In Services on all websites ("opting out" of Identity/Sign-In Services), requiring you to log in manually instead on all websites. This also requires you to log in to all websites separately. When enabled, all the buttons and links to use Identity/Sign-In Services on websites are removed, and the ability for these services to collect data is also removed. Assume that you will be able to undo this setting at any time.
Targeted Advertising	Imagine that you are given a new setting in your browser that allows you to block Targeted Advertising on specific websites you choose ("opting out" of Targeted Advertising on these websites). When enabled, ads which use Targeted Advertising are blocked on websites which you opt out of. Ordinary ads which do not use Targeted Advertising are not affected by this setting. By default, you are still shown Targeted Advertising on websites which you are not opted out of. Assume that you will be able to reverse this setting on any website, at any time.	Imagine that you are given a new setting in your browser that enables you to block all Targeted Advertising ads ("opting out" of Targeted Advertising). When enabled, ads which use Targeted Advertising are blocked on all websites. Ordinary ads which do not use Targeted Advertising are not affected by this setting. Assume that you will be able to undo this setting at any time.
Behavioral Profiling	Imagine that you are given a new setting in your browser that enables you to block Behavioral Profiling from occurring on specific websites you choose ("opting out" of Behavioral Profiling on these websites). On the websites you opt out from, your browser hides your identity and blocks any information your browser might send in the background while you are browsing. This ensures the specific websites you opt out from cannot perform Behavioral Profiling on you. Assume that you will be able to reverse this setting on any website, at any time.	Imagine that you are given a new setting in your browser that enables you to block all websites from performing Behavioral Profiling ("opting out" of Behavioral Profiling). When enabled, the setting hides your identity and blocks any information your browser might send in the background while you are browsing. Assume that you will be able to undo this setting at any time.
Session Replay	Imagine you are given a new setting in your browser that enables you to block Session Replay from occurring on specific websites you choose ("opting out" of Session Replay on these websites), preventing the websites you opt out from collecting what is needed for Session Replay to occur. By default, on websites you have not opted out from, Session Replay will still occur normally. Assume that you will be able to reverse this setting on any website, at any time.	Imagine you are given a new setting in your browser that enables you to block Session Replay from occurring on all websites ("opting out" of Session Replay), preventing all websites from collecting what is needed for Session Replay to occur. Assume that you will be able to undo this setting at any time.
Reporting and Analytics	Imagine that you are given a new setting in your browser that enables you to block specific websites you choose from performing Reporting and Analytics ("opting out" of Reporting and Analytics on these websites). When enabled, your browser sends misleading signals to the websites that you opt out from, preventing the Reporting and Analytics mechanisms on websites from working there. By default, Reporting and Analytics will still work as it would normally on websites you do not choose to opt out from. Assume that you will be able to reverse this setting on any website, at any time.	Imagine that you are given a new setting in your browser that enables you to block all websites from performing Reporting and Analytics ("opting out" of Reporting and Analytics). When enabled, your browser sends misleading signals to all websites, to prevent the Reporting and Analytics mechanisms from working anywhere. Assume that you will be able to undo this setting at any time.
Fingerprinting and Deanonymization	Imagine that you are given a new setting in your browser that enables you to block Fingerprinting and Deanonymization from occurring on specific websites you choose ("opting out" of Fingerprinting and Deanonymization on these websites). When enabled, the setting sends misleading signals to the websites you opt out from, which prevents Fingerprinting and Deanonymization from taking place on those websites. By default, websites which you have not opted out from will still allow the Fingerprinting and Deanonymization to take place as they would normally. Assume that you will be able to reverse this setting on any website, at any time.	Imagine that you are given a new setting in your browser that enables you to block Fingerprinting and Deanonymization from occurring on all websites ("opting out" of Fingerprinting and Deanonymization). When enabled, the setting sends misleading signals to all websites, which prevents Fingerprinting and Deanonymization from occurring. Assume that you will be able to undo this setting at any time.
"Nag" Screens	Imagine that you are given a new setting in your browser that enables you to block "Nag" Screens on specific websites you choose ("opting out" of "Nag" Screens on these websites). When enabled, your browser blocks "Nag" Screens on specific websites, removing them from the contents of websites you opt out from. By default, on websites you have not opted out from, "Nag" Screens are still shown as they would normally be. Assume that you will be able to reverse this setting on any website, at any time.	Imagine that you are given a new setting in your browser that enables you to block "Nag" Screens ("opting out" of "Nag" Screens). When enabled, your browser blocks "Nag" Screens everywhere, removing them from all websites. Assume that you will be able to undo this setting at any time.
Crypto-Mining	Imagine you are given a new setting in your browser that enables you to block Crypto-Mining on specific websites you choose ("opting out" of Crypto-Mining on these websites), preventing Crypto-Mining from taking place on these websites in your browser. By default, on websites which you have not opted out from, Crypto-Mining is still allowed. Assume that you will be able to reverse this setting on any website, at any time.	Imagine you are given a new setting in your browser that enables you to block Crypto-Mining on all websites ("opting out" of Crypto-Mining), preventing any Crypto-Mining from taking place in your browser. Assume that you will be able to undo this setting at any time.

the priming exercise in the first part of the survey, in random order:]

Consider [specific user-provided website], which is a [website category] website.

If you had a single one-click setting which enabled you to opt out of [PIP] on [specific user-provided website], how likely would you be to use it? [Participant is presented with 4-point Likert scale response options ranging from Very Unlikely to Very Likely with no neutral response.]

Consider all the [website category] websites across the entire internet, which includes [specific user-provided website] and [specific user-provided website].

If you had a single one-click setting that enabled you to opt out of [PIP] on all [website category] websites, how likely would you be to use it? [Participant is presented with 4-point Likert scale response options ranging from Very Unlikely to Very Likely with no neutral response.]

[Participant is presented with a randomized attention check question, which includes a reCAPTCHA test.]

Section 5

What benefits do you think that companies which have [PIP] on their website get from [PIP]? [Participant is presented with a free text entry field.]

Are you aware of anything you can do to enable or disable [PIP] while browsing? Please explain. [Participant is presented with a free text entry field.]

Have you ever tried to enable or disable [PIP]? [Participant may choose between: Yes/No]

Why or why not? [Participant is presented with a free text entry field.]

[If the participant answered yes:] How did you know if you succeeded or failed? Would you want to be informed about the presence or absence of [PIP] on the websites you browse? [Participant is presented with 5-point Likert scale response options ranging from Definitely Yes to Definitely not with a neutral response of I don't know.]

[Participant is presented with the post-survey.]

Survey 2 - Quantitative

Section 1

The next part of the survey is intended to collect your thoughts and experiences with a particular web technology.

Please read the description of the technology on the following page carefully. You will be asked a series of questions which depend on you having read the description.

[Participant is presented with the PIP description.] Here are some examples of concrete risks and benefits for [PIP].

Please take note of these risks and benefits and consider them carefully as you progress through the rest of the survey. [Participant is presented with the list of PIP risks and benefits.]

Throughout the survey, you can click on the following button located at the top of each page, to see a reminder of the definition of [PIP] and the risks and benefits associated with it. [Participant is presented with a button labeled with the name of the PIP, which is present throughout the survey on the top of each page.]

Section 2

In this section, you will be asked about [PIP] in a variety of scenarios.

Please carefully consider the definition of [PIP] and the associated risks and benefits you just saw when answering the questions which follow.

[The following section repeats for all the specific websites participants were asked about, across all website categories, in random order.]

[Participant is presented with a screenshot of a specific website, along with the name of the website, their logo, the date they were established, the country they are based in.]

[website name] is a [website category] website, established [date], based in [location].

Please take a moment to familiarize yourself with the website if you aren't already familiar with it.

[Participant is presented with the scenario text for opting out of a specific website.]

Consider *[website]*, which is a *[website category]* website. How likely would you be to use the setting described above to **opt out of [PIP] on [website]**? *[Participant is presented with 4-point Likert scale response options ranging from Very Unlikely to Very Likely with no neutral response.]*

[Participant is presented with the scenario text for opting out of a website category.]

Consider all the *[website category]* websites across the internet, which includes the two websites you saw a moment ago, *[specific website]* and *[specific website]*, and many others.

How likely would you be to use the setting described above to opt out of *[PIP]* for all *[website category]* websites?

This setting would not affect your separate choice to opt out (or to not opt out) for specific websites. *[Participant is presented with 4-point Likert scale response options ranging from Very Unlikely to Very Likely with no neutral response.]*

Section 3

[Participant is presented with the scenario text for opting out of PIP on every website.]

Note that this would apply to every website you visit, no matter what category it belongs to.

How likely would you be to use the setting described above, to opt out of *[PIP]* **on every website you visit**? *[Participant is presented with 4-point Likert scale response options ranging from Very Unlikely to Very Likely with no neutral response.]*

Section 4

Please answer the following questions about how often you would like to be notified about *[PIP]* on different categories of websites. *[The participant is presented with a matrix of questions from all website categories in randomized order.]*

On *[website category]* websites, how often would you like to be notified about *[PIP]*? *[Participant is presented with 5 response options: Notify every time I visit, Notify me only once per week, Notify me only once per month, Notify me only the first time I visit, Never notify me.]*
[Participant is presented with the post-survey.]

Grounded Analysis Codebook

Overarching Themes

The codebook contains 26 codes in 7 categories of codes which represent overarching themes seen in the responses. “Trends” refers to a category of codes which were generated in second-cycle coding, which had specific relevance to trends in responses seen after first-cycle coding. “Understanding” is a category based around questions concerning what practices participants seemed to understand, or misunderstand. “Bad Assumptions” is a category of responses created in second-cycle coding which was intended to identify specific assumptions that participants were making in their responses that were at times based on misunderstandings, lack of knowledge, or misperceptions. “Opposition” is a category which reflects attitudes, actions, and concerns participants expressed in opposition to intrusive practices. “Acceptance” highlights reasons, experiences, and expressions of ambivalence or ignorance towards practices which led to accepting them in certain circumstances. “Experience” is a category which highlights specific experiences, incidents, and their circumstances which participants shared, as well as expressions of lacking experience. Finally, “Miscellaneous” was a category with only one code, used to highlight responses which were selected for removal from the dataset due to poor quality or survey abuse which was not detected by automated measures. It is worth noting that some codes were not mutually exclusive; many of the responses contained multiple layers of meanings and thus were assigned several codes simultaneously.

Trends

security_thinking: participant expresses evidence of thinking that is directly related to security, protection from security threats, protecting accounts and preventing fraud/scams (10 instances in 10 responses)

profiling_mentions_ads: participant explicitly seems to be making a connection between behavioral profiling and advertisements, targeted or otherwise (3 instances in 3 responses)

breakage: participant explicitly mentions parts of a website not functioning correctly (7 instances in 7 responses)

Understanding

understanding_demonstrates_knowledge: participant expresses factual or operational knowledge of

the technology and/or ramifications of their interactions with it (371 instances in 160 responses)

understanding_vague: participant seems to express a vague or incomplete understanding of the technology or their interactions with it, such that it is difficult to gauge their level of understanding or expertise (123 instances in 83 responses)

understanding_misconception: participant seems to demonstrate a lack of knowledge about the technology and/or ramifications of their interactions with it, either by expressing factual inaccuracies, or other errors such as mixed-up terminology (110 instances in 74 responses)

Assumptions

bad_assumption: participant seems to be making an incorrect assumption (129 instances in 99 responses)

adblocker_effectiveness: participant seems to be making a bad assumption, specifically about the effectiveness of ad blocking tools (21 instances in 19 responses)

incognito_mode: participant seems to be making an assumption, specifically about the effectiveness of incognito mode/private browsing mode or similar features offered by private browsers, Tor, VPNs, general privacy extensions which are not ad-blockers, clearing cookies/history (58 instances in 51 responses)

antivirus: participant seems to be making an assumption about the effectiveness of antivirus tools or firewalls in blocking privacy threats (3 instances in 3 responses)

has_control: participant seems to be making an assumption about having control over a setting which does not actually exist, or over a variable which they do not actually have control over (60 instances in 62 responses)

safe_browsing: participant seems to be making an assumption about being protected based on their own special browsing behavior, which makes them safe (15 instances in 13 responses)

malware_risk: participant seems to be making an assumption about the risk of being infected with malware (12 instances in 12 responses)

concerned_ads: participant is explicitly concerned with advertisements, either thinking that this is the way they can tell there is a problem, or that they are safe (16 instances in 16 responses)

unconcerned: participant seems to be totally unconcerned with any privacy or security risk that may

come about as a result of this practice (7 instances in 7 responses)

Opposition

opposition_action: a specific action or mitigation strategy that a participant employs to oppose an intrusive practice (83 instances in 80 responses)

opposition_disable_attempt: participants experiences with disabling/attempting to disable a practice (54 instances in 51 responses)

opposition_concern: participants expressing a specific concern that they were attempting to address/mitigate (138 instances in 113 responses)

Acceptance

acceptance_approval: reasons why participants seem to express explicit or tacit approval of a practice; they like it, and they don't believe that there are negatives/risks for them (50 instances in 44 responses)

acceptance_ambivalence: reasons why participants seem to express explicit or tacit acceptance toward a practice; they recognize it is/might be bad or intrusive, but it does not bother them, will not get in the way, etc. (34 instances in 33 responses)

acceptance_ignorance: reasons where participants express ignorance about a practice and/or the ramifications of their interactions with it, suggesting that they are okay with the practice because they do not understand it or know enough about it to form an opinion (108 instances in 107 responses)

Experience

experience_positive: participants express a positive experience when interacting with a practice, including acknowledging that they received a benefit (9 instances in 8 responses)

experience_negative: participants express a negative experience when interacting with a practice, including fears of repercussions, "creep factor" and other concerns or harms that they directly or indirectly experienced (29 instances in 25 responses)

experience_neutral: participants express some form of experience with interacting with a practice, but without obvious or apparent risks or benefits; they just acknowledge that there was some kind of experience without making a judgment about it (65 instances in 65 responses)

experience_lacking: participants expressing a lack of experience and/or ignorance about whether they actually had an experience with a practice (9 instances in 8 responses)

Table 8. Breakdown of self-reported demographics from our surveys. Note that the quantitative Survey 2 had additional demographics collected, which were not collected during the qualitative Survey 1.

		Survey 1	Survey 2
Total Responses		223	1069
Rejections	Survey Abuse	24	48
	Poor Quality	13	0
	Rejection Rate	17%	4%
Gender	Male	65%	57%
	Female	35%	42%
	Other	0%	1%
Age Range	18 to 24	9%	6%
	25 to 44	69%	69%
	45 to 64	19%	21%
	≥ 65	3%	4%
Education Level	< High School	<1%	<1%
	High School	14%	11%
	Some College	14%	18%
	2-year Associates	11%	13%
	4-year Bachelor's	48%	45%
	Advanced Degree	12%	13%
City Size	Rural Area	10%	12%
	Town or Suburb	41%	37%
	City	31%	32%
	Large City	18%	19%
Marital Status	Never Married	51%	42%
	Married	37%	47%
	Divorced	8%	6%
	Other	4%	5%
SA-6 Score [26]	Mean	3.7	3.8
	Median	3.8	3.8
	Std.	0.86	0.83
Education Field	STEM	-	41%
	Non-STEM	-	59%
Employment Field	STEM	-	50%
	Non-STEM	-	50%
Preferred Browser	Chrome	-	80%
	Firefox	-	13%
	Safari	-	3%
	Edge	-	1%
	IE	-	1%
	Other	-	2%
Looked at Settings	This year	-	24%
	This month	-	42%
	This week	-	31%
	Never	-	3%
Changed Settings	This year	-	33%
	This month	-	43%
	This week	-	18%
	Never	-	6%
At-Risk Group	Yes	-	19%
	No	-	81%

Table 9. Regression table for opt-out likelihood for all PIPs. Factor levels with no data points are marked with \emptyset , and odds ratios which did not converge due to large standard error are marked NC. P-values and odds ratios (Z-test) are shown with respect to the intercept. All factors are included (except gender and SA-6), however factors with statistically significant p-values are darkened. Intercept: AgeRange [18-24], EducationLevel [Associates], CitySize [City], MaritalStatus [Divorced], EmploymentStatus [Employed], EmploymentField [Non-STEM], PrivacySettings_LastLooked [Past month], PrivacySettings_LastChanged [Past month], Browser [Chrome], PrivacySurveys_PastYear [6-9], Privacy_AtRisk [FALSE], website_category [ADULT].

Factors	Behav. Profiling (n=113)		Reporting (n=113)		Session Replay (n=99)		Targeted Ads (n=103)	
	Odds Ratios	p	Odds Ratios	p	Odds Ratios	p	Odds Ratios	p
(Intercept)	22.86	0.009	2.39	0.481	43.54	0.043	1.25	0.816
AgeRange [25-44]	1.09	0.887	0.83	0.834	2.32	0.567	3.8	0.016
AgeRange [45-64]	1.56	0.505	3.53	0.169	1.93	0.660	4.41	0.021
AgeRange [65+]	5.34	0.096	4.8	0.339	NC	0.648	7.87	0.045
EducationLevel [Bachelors]	0.38	0.043	1.79	0.258	0.55	0.427	0.86	0.659
EducationLevel [PhD]	\emptyset	\emptyset	19.53	0.019	0.14	0.319	1.61	0.770
EducationLevel [High School]	0.41	0.150	2.29	0.192	0.35	0.250	0.62	0.333
EducationLevel [<High School]	\emptyset	\emptyset	\emptyset	\emptyset	0.12	0.295	\emptyset	\emptyset
EducationLevel [Masters]	0.30	0.053	1.45	0.540	0.34	0.261	1.2	0.782
EducationLevel [JD, MD]	1.01	0.995	NC	0.950	1.18	0.944	5.1	0.249
EducationLevel [Some College]	0.56	0.303	2.58	0.080	1.03	0.976	1.05	0.920
EducationField [STEM]	0.72	0.396	4.09	0.007	0.71	0.558	0.45	0.047
CitySize [Large City]	2.84	0.018	2.08	0.114	1.29	0.664	1.24	0.564
CitySize [Rural Area]	1.81	0.199	1.5	0.435	1.27	0.688	0.84	0.735
CitySize [Town or Suburb]	2.35	0.022	1.05	0.910	3.93	0.006	1.18	0.661
MaritalStatus [Married]	1.45	0.599	0.88	0.856	0.28	0.109	3.17	0.059
MaritalStatus [Never married]	0.85	0.829	1.44	0.612	0.37	0.206	1.72	0.363
MaritalStatus [Prefer not to disclose]	0.16	0.242	0.74	0.851	0.35	0.452	93.95	0.527
MaritalStatus [Separated]	3.14	0.304	1.42	0.824	4.16	0.423	20.74	0.002
MaritalStatus [Widowed]	2.22	0.611	3.85	0.503	0.13	0.114	\emptyset	\emptyset
EmploymentStatus [Student]	\emptyset	\emptyset	0.2	0.026	0.04	0.163	1.85	0.550
EmploymentStatus [Unemployed]	2.32	0.132	0.16	0.003	0.44	0.451	0.64	0.388
EmploymentStatus [Prefer not to answer]	1.19	0.920	4.01	0.419	\emptyset	\emptyset	\emptyset	\emptyset
EmploymentField [STEM]	0.79	0.538	0.33	0.031	1.1	0.869	1.46	0.349
PrivacySettings_LastLooked [Past week]	3.20	0.015	1	0.992	2.58	0.126	2.09	0.033
PrivacySettings_LastLooked [Past year]	1.07	0.872	1.13	0.797	1.21	0.719	2.64	0.042
PrivacySettings_LastLooked [Never]	4.92	0.084	99.44	0.007	0.55	0.572	0.99	0.994
PrivacySettings_LastChanged [Past week]	0.09	<0.001	1.94	0.238	0.88	0.842	1.3	0.562
PrivacySettings_LastChanged [Past year]	0.31	0.004	1.11	0.797	1.3	0.580	1.11	0.758
PrivacySettings_LastChanged [Never]	0.05	<0.001	0.67	0.526	1.49	0.721	1.93	0.597
Browser [Edge]	\emptyset	\emptyset	0.87	0.932	\emptyset	\emptyset	\emptyset	\emptyset
Browser [Firefox]	1.92	0.118	1.94	0.114	0.47	0.182	3.75	0.011
Browser [IE]	5.19	0.262	0.03	0.020	\emptyset	\emptyset	2.25	0.476
Browser [Other]	4.37	0.224	15.99	0.007	NC	0.949	2.36	0.297
Browser [Safari]	1.32	0.805	1.22	0.807	0.02	0.024	10	0.011
PrivacySurveys_PastYear [<5]	0.68	0.337	1.2	0.704	0.4	0.135	1.33	0.450
PrivacySurveys_PastYear [>10]	1.91	0.414	2.2	0.268	0.72	0.685	2.32	0.199
PrivacySurveys_PastYear [0]	0.71	0.429	0.97	0.941	0.33	0.093	0.35	0.012
Privacy_AtRisk [TRUE]	2.66	0.028	3.05	0.007	3.22	0.016	1.32	0.484
website_category [FINANCE]	0.60	0.082	0.29	<0.001	0.69	0.298	0.34	0.001
website_category [GAMES]	0.14	<0.001	0.17	<0.001	0.12	<0.001	0.12	<0.001
website_category [HEALTH]	0.20	<0.001	0.11	<0.001	0.12	<0.001	0.1	<0.001
website_category [NEWS]	0.27	<0.001	0.22	<0.001	0.1	<0.001	0.3	<0.001
website_category [SHOPPING]	0.15	<0.001	0.13	<0.001	0.23	<0.001	0.08	<0.001
website_category [SOCIAL]	1.05	0.877	0.64	0.135	0.88	0.720	0.53	0.05
website_category [TRAVEL]	0.46	0.008	0.39	0.001	0.53	0.065	0.28	<0.001

Factors	Crypto-Mining (n=102)		Identity Sign-In (n=133)		Fingerprinting (n=121)		Nag Screens (n=104)	
	Odds Ratios	p	Odds Ratios	p	Odds Ratios	p	Odds Ratios	p
(Intercept)	1423.54	0.012	642.25	<0.001	27.39	0.001	141.75	<0.001
AgeRange [25-44]	0.32	0.337	0.13	0.009	0.57	0.299	0.48	0.303
AgeRange [45-64]	0.44	0.494	0.11	0.009	0.56	0.331	0.45	0.280
AgeRange [65+]	4.38	0.343	0.09	0.063	1.6	0.670	0.27	0.205
EducationLevel [Bachelors]	1.2	0.721	1.38	0.488	1.7	0.182	0.51	0.157
EducationLevel [PhD]	0.08	0.043	1.99	0.613	1.55	0.648	∅	∅
EducationLevel [High School]	0.3	0.036	1.04	0.953	0.64	0.398	0.26	0.041
EducationLevel [<High School]	∅	∅	NC	0.421	∅	∅	∅	∅
EducationLevel [Masters]	2.51	0.159	1.79	0.310	5.77	0.002	1.04	0.958
EducationLevel [JD, MD]	0.17	0.047	7.84	0.092	1.94	0.544	2.95	0.554
EducationLevel [Some College]	0.52	0.236	0.98	0.973	3.01	0.020	0.57	0.348
EducationField [STEM]	0.57	0.152	1.68	0.188	2.47	0.118	0.71	0.517
CitySize [Large City]	0.65	0.332	0.69	0.342	0.47	0.049	0.62	0.310
CitySize [Rural Area]	3.67	0.032	0.78	0.590	0.43	0.110	0.39	0.091
CitySize [Town or Suburb]	0.36	0.003	0.89	0.722	0.97	0.926	0.54	0.146
MaritalStatus [Married]	0.01	0.096	0.41	0.115	0.27	0.011	0.6	0.393
MaritalStatus [Never married]	0.02	0.129	0.22	0.013	0.41	0.094	0.66	0.497
MaritalStatus [Prefer not to disclose]	0	0.019	5.01	0.369	0.54	0.519	NC	0.989
MaritalStatus [Separated]	0.03	0.244	∅	∅	∅	∅	746.3	0.530
MaritalStatus [Widowed]	0.2	0.612	0.03	0.063	∅	∅	1.36	0.801
EmploymentStatus [Student]	83.24	0.007	0.21	0.180	13.42	0.041	3.37	0.200
EmploymentStatus [Unemployed]	0.83	0.729	1.65	0.225	1.28	0.686	1.47	0.478
EmploymentStatus [Prefer not to answer]	∅	∅	1.12	0.933	2.62	0.361	0.15	0.214
EmploymentField [STEM]	4.93	<0.001	1.08	0.830	0.5	0.237	1.18	0.763
PrivacySettings_LastLooked [Past week]	0.25	0.002	1.87	0.127	1.42	0.309	0.34	0.015
PrivacySettings_LastLooked [Past year]	0.69	0.400	2.03	0.086	0.79	0.516	0.71	0.479
PrivacySettings_LastLooked [Never]	1.89	0.643	3.36	0.145	0.22	0.172	0.03	0.014
PrivacySettings_LastChanged [Past week]	1.33	0.554	0.51	0.165	0.63	0.301	1.75	0.306
PrivacySettings_LastChanged [Past year]	0.93	0.853	0.38	0.010	0.62	0.195	0.71	0.393
PrivacySettings_LastChanged [Never]	0.04	0.006	0.07	<0.001	1.31	0.757	3.69	0.282
Browser [Edge]	8.93	0.021	0.88	0.911	0.27	0.393	∅	∅
Browser [Firefox]	2.83	0.013	1.83	0.178	1.78	0.182	2.85	0.022
Browser [IE]	∅	∅	27.42	0.080	3.79	0.257	13.27	0.159
Browser [Other]	0	0.444	0.22	0.018	0.25	0.387	7.75	0.195
Browser [Safari]	NC	0.502	0.54	0.719	0.22	0.076	0.48	0.460
PrivacySurveys_PastYear [<5]	1.14	0.782	0.39	0.020	1.43	0.463	0.6	0.297
PrivacySurveys_PastYear [>10]	2.06	0.237	0.99	0.987	1.74	0.409	2.35	0.149
PrivacySurveys_PastYear [0]	2.27	0.124	0.46	0.067	0.91	0.857	0.61	0.344
Privacy_AtRisk [TRUE]	2.35	0.029	0.79	0.586	1.53	0.358	0.84	0.693
website_category [FINANCE]	1.42	0.307	0.25	<0.001	0.09	<0.001	0.21	<0.001
website_category [GAMES]	0.9	0.740	0.16	<0.001	0.14	<0.001	0.56	0.036
website_category [HEALTH]	0.33	<0.001	0.19	<0.001	0.13	<0.001	0.22	<0.001
website_category [NEWS]	0.5	0.029	0.16	<0.001	0.32	<0.001	0.63	0.096
website_category [SHOPPING]	0.81	0.513	0.16	<0.001	0.16	<0.001	0.28	<0.001
website_category [SOCIAL]	1.84	0.083	0.23	<0.001	0.39	<0.001	1.05	0.876
website_category [TRAVEL]	1.33	0.397	0.34	<0.001	0.28	<0.001	0.5	0.012