Generalizable Active Privacy Choice: Designing a Graphical User Interface for Global Privacy Control

Sebastian Zimmeck, Eliza Kuller, Chunyue Ma, Bella Tassone, Joe Champeau*
{szimmeck,ekuller,cma01,itassone,jchampeau}@wesleyan.edu
Wesleyan University
United States

ABSTRACT

The California Consumer Privacy Act and other privacy laws give people a right to opt out of the sale and sharing of personal information. In combination with privacy preference signals, especially, Global Privacy Control (GPC), such rights have the potential to empower people to assert control over their data. However, many laws prohibit opt out settings being turned on by default. The resulting usability challenges for people to exercise their rights motivate generalizable active privacy choice — an interface design principle to make opt out settings usable without defaults. It is based on the idea of generalizing one individual opt out choice towards a larger set of choices. For example, people may apply an opt out choice on one site towards a larger set of sites. We explore generalizable active privacy choice in the context of GPC.

We design and implement nine privacy choice schemes in a browser extension and explore them in a usability study with 410 participants. We find that generalizability features tend to decrease opt out utility slightly. However, at the same time, they increase opt out efficiency and make opting out less disruptive, which was more important to most participants. For the least disruptive scheme, selecting website categories to opt out from, 98% of participants expressed not feeling disrupted, a 40% point increase over the baseline schemes. 83% of participants understood the meaning of GPC. They also made their opt out choices with intent and, thus, in a legally relevant manner. To help people exercise their opt out rights via GPC our results support the adoption of a generalizable active privacy choice interface in web browsers.

KEYWORDS

Opt Out, Privacy Choice, Privacy Rights, Privacy Laws, Global Privacy Control, GPC, Do Not Sell, Do Not Track, DNT, Privacy Preference Signals, Usability, CCPA, CPRA, Web Privacy

1 INTRODUCTION

Data collection and sharing are central to many companies' business models. People pay with their data for news, videos, and other content on the web. This trade — data for content — has fueled the development of an extensive web tracking ecosystem [46]. A study on the frequency and reach of trackers on over 21 million pages of 350,000 unique websites found that 95% contain third party requests to potential trackers and 78% attempt to transfer data elements that are either user identifiers or can be used as such [78]. Common tracking techniques include third party cookies [20], browser fingerprinting [55], and tracking pixels [32]. Web tracking is generally non-transparent, and people have no effective way of preventing it. New privacy laws aim to change this situation. The General Data Protection Regulation (GDPR) [21] sparked a flurry of privacy law-making activity around the world. Stateside, California enacted the California Consumer Privacy Act (CCPA) [9], which was extended by the California Privacy Rights Act (CPRA). Various other states, e.g., Colorado [22], Connecticut [70], Montana [43], Oregon [58], and Texas [44], recently enacted privacy laws as well.

Per the CCPA, California consumers have a right to opt out from the sale and sharing of personal information. In order for an opt out choice to be valid it has to "[c]learly represent a consumer's intent and be free of defaults constraining or presupposing that intent."[1] Similarly, Colorado law prohibits to "adopt a mechanism that is a default setting" but rather requires a mechanism that "clearly represents the consumer's affirmative, freely given, and unambiguous choice to opt out of the processing of personal data."[2] Per Connecticut law, consumers need "to make an affirmative, freely given and unambiguous choice."[3] In order to bring opt out rights to life under these and similar laws we need usable privacy choice interfaces that are not reliant on default settings.[4]

Default settings became a major obstacle for the adoption of Do Not Track (DNT) [75]. DNT could be set to one of two states — enabled or disabled — and whichever was set by default resulted in a choice for an individual [19]. However, arguably such choice can be at odds with an individual's autonomy and, thus, be irrelevant from a legal perspective. While default settings are based on the premises that there is a choice that fits the majority of people and that real choice is practically infeasible [19], we make the assumptions here that people make their own choices and that the usability problem can be resolved. Thus, we transform the problem of default settings into the problem of designing usable choice interfaces.

Legal compliance and usability are the two main requirements for enabling people to effectively and efficiently opt out. To satisfy the former a privacy choice interface must enable an individual to show their intent by making an *active* choice, e.g., by clicking on an opt out link. To make an active choice usable in the context of browsing the web, we can *generalize* an individual's choice. Notably,

*Eliza Kuller and Chunyue Ma graduated in Spring 2023 and 2022, respectively, and are no longer at Wesleyan University. Corresponding author: Sebastian Zimmeck.

Proceedings on Privacy Enhancing Technologies 2024(1), 258–279

https://doi.org/10.56553/popets-2024-0015

[1] California Civil Code §1798.185(a)(19)(A)(iii).
[2] Colorado Privacy Act §6-1-1313(2)(c).
[3] Connecticut Data Privacy Act §6(e)(1)(A)(ii)(II).
[4] While California and Colorado consumers can exercise their choices via privacy-by-design products or services that are marketed to include opt out functionality [13, 68], such implicit choices may not be recognized by all laws.

California, Colorado, Connecticut, Montana, Oregon, and Texas explicitly recognize universal opt outs. Such universal settings aid usability and prevent the plight of cookie banners, which many people just click away without caring about their choices [25].

The principle of generalizable active privacy choice may be applied to different types of privacy choices. Use cases are opting out from web tracking, cookie consent, or app privacy settings. In this paper we focus on opting out from web tracking using Global Privacy Control (GPC) [23]. GPC is a privacy preference signal by which web users can transmit an opt out choice to a website mediated by their browser. We see GPC as an important use case as the signal is increasingly adopted in the web ecosystem and enforced by the Office of the California Attorney General [69]. We use generalizable active privacy choice to explore the design of graphical user interfaces for opting out with GPC. We make the following contributions:

(1) We introduce the principle of generalizable active privacy choice to make privacy choice interfaces usable without default settings. We design and implement nine privacy choice schemes exploring various generalizability dimensions. (§3)
(2) Evaluating our schemes in a usability study with 410 participants we observe that generalizability features decrease the perceived level of browsing disruption. 98% of the participants of the least disruptive scheme expressed that they did not feel disrupted. (§4 and §5)
(3) We provide recommendations for regulators, browser vendors, and publishers. Our results support the integration of a generalizable active privacy choice interface for GPC in the web browser. (§6)

2 BACKGROUND AND RELATED WORK

Generalizable active privacy choice is based on notice and choice.

2.1 Notifying People of their Privacy Choices

Notice and choice is a fundamental building block of many privacy laws around the world. We focus on notice and choice for opting out. If people are not aware of their opt out rights, they would be deprived of exercising them [34, 63, 65]. Notices can be categorized according to various dimensions, e.g., when they are shown or through which channel [62]. Short-form notices, which we explore here, are particularly useful when shown in context, e.g., when a website requests location access. Similarly, icons can concisely convey privacy-related information [14, 28]. Apple and Google now require apps in their app stores to display privacy labels. Such labels have the potential to improve people's understanding and control of apps' privacy practices [79]. While some labels were shown to be inaccurate and misleading, especially in less popular apps [39, 40], concise privacy notices are a promising approach to raise awareness when displayed in a salient way [18].

2.2 Designing Usable Privacy Choice Interfaces

Without usable privacy choice interfaces people's privacy rights would not amount to much. For the usability evaluation of our privacy choice schemes we consider a broad spectrum of factors [25]: user needs and sentiment (§5.1), effort and ability (§5.2), awareness, comprehension, and intent (§5.3), and nudging patterns and decision reversal (§5.4). With generalizable active privacy choice, we aim to avoid two types of effectively unusable privacy settings, which are at the opposite extremes of the user awareness dimension: (1) hidden privacy settings and (2) excessively shown privacy settings. Hiding privacy settings makes exercising privacy choices akin to a scavenger hunt [14]. A recent study demonstrated this difficulty for Do Not Sell links on data broker sites [48]. Surfacing a standardized opt out banner in addition to a link would increase user awareness [57, 63] However, banners may also negatively impact usability. Cookie banners serve as a cautionary tale [24, 25, 42].[5] Thus, we are interested in evaluating the extent to which banner choices can be generalized to achieve better usability.

People are generally less concerned about tracking on first party sites compared to tracking on third party sites [11]. Recent results further suggest that there is value in privacy choice settings distinguishing between website categories [65]. Learning people's privacy preferences seems also a good way to reduce the complexity of privacy choices [36, 53, 64]. These findings motivate the design of our privacy choice schemes (§3.3). Generally, control, or at least the feeling of having control, was shown to increase satisfaction with privacy interfaces [11, 35, 38, 52]. On the other hand, inconsistent privacy interface implementations make it confusing and difficult for people to make informed privacy choices [27, 45]. Thus, it is important to design the schemes based on the mental models of the audience. As people rely on folk models in their comprehension of web tracking [77] and experience expectational mismatches [61], our designs are based on the former and try to avoid the latter.

2.3 Avoiding Behavioral Nudges

People should make *their* choices. Thus, we want to avoid behavioral nudges. The design and architecture of choice interfaces heavily influence people's decisions [71, 73]. Thus, publishers hold considerable power over access to the data of their users. For example, with simple design changes they can affect the choices of a substantial share of people to consent to the use of cookies [3]. Other design changes, such as displaying notifications in the lower left part of the screen, can nudge people towards more interaction [73]. Designs with negative framing (e.g., "Deny cookies and degrade your experience on this site") result in significantly lower cookie opt out rates compared to positive framing (e.g., "Accept cookies to improve your experience on this site") [47]. A comparison of interface designs showed that people prefer neutral and text-based communications with visual iconography over purely visual designs [29].

2.4 Eliminating Dark Patterns

Dark patterns are user interface designs that obstruct an individual's autonomy in the decision-making process. The use of dark patterns and default settings could nudge people away from selecting privacy-protective options [26]. Dark patterns can be considered a weaponization of behavioral research to serve the aims of the surveillance economy [54]. By exploiting people's cognitive biases [49], online services influence people to purchase goods and subscriptions, spend more time on-site, or accept the harvesting

[5] Apart from the limited usability of current cookie consent implementations, many sites use questionable practices, such as registering positive consent even though people did not make any choice [50]. Automating cookie consent interface discovery as well as consent decisions seem viable paths for future improvements [37].

of their data [4, 56]. While one study found an inverse correlation between dark pattern recognition and likelihood to be influenced, interestingly, study participants' level of awareness did not play a significant role in predicting their ability to resist manipulative designs implying that active interventions are needed to combat dark patterns [4]. Thus, regulation should not only require consent but also clarify how this consent has to be obtained to ensure that people can make free and informed choices [74]. Helping people to make such choices is the broader goal that we hope to further with generalizable active privacy choice.

2.5 Choosing the Right Layer of Abstraction

Participants in our study interacted with a median of 78 unique sites per week. This large number of sites makes site-by-site opt outs generally challenging [48, 57]. Site-by-site opt outs naturally beget some level of interruption to user activity. Thus, the preferable layer of abstraction for privacy choice interfaces appears to be the platform (e.g., the web browser) and not their individual applications (e.g., the websites) (§5.1.2). Standardizing privacy choice interfaces in the browser rather than being left to individual sites would also have the advantage of providing a uniform interface to support notification, control, and could help mitigate dark patterns [65]. Different from websites, individual mobile app opt outs are more usable due to the lower number of apps people use.[6]

2.6 The Return of Privacy Preference Signals

Privacy preference signals are digital representations to express if and how people agree to their data being processed [31]. As they must be adopted by both senders and recipients, adoption remains an unsolved coordination problem [31]. The first major privacy preference signal, the Platform for Privacy Preferences Project (P3P) [15, 16], enabled people to delegate privacy choices to user agents that could automatically react to websites' privacy practices bases on the sites' machine-readable privacy policies. Then, DNT [75] was developed as a binary signal for people to express their opt out from tracking per the California Online Privacy Protection Act (CalOPPA) [8]. DNT adoption remains low as CalOPPA does not require DNT signal recipients to actually comply with received DNT signals but only to disclose whether they comply.[7] Now, an increasing number of new privacy laws makes privacy preference signals a corner stone of their opt out regimes [80].

We focus here on GPC, which is a privacy preference signal developed by a coalition of privacy organizations, publishers, browser vendors, extension developers, and academics for helping people to exercise their Do Not Sell and Do Not Share rights [23, 81]. The CCPA and various other new privacy laws require certain recipients of GPC signals to respect those as valid opt out expressions.[8] While the privacy choice schemes we implement are evaluated in the context of GPC, they also inform choice implementation for other privacy preference signals and settings more broadly. Privacy preference signals work well in conjunction with the platform-layer abstraction that is preferred for privacy choices. Advanced Data Protection Control (ADPC) is a privacy preference signal similar to P3P and focused on enabling cookie consent and other privacy choices under the GDPR and ePrivacy Directive [1, 33]. The future will tell if the coexistence of multiple privacy preference signals creates ambiguity as people may transmit more than one signal [30].

2.7 Reframing the Default Problem

We start our inquiry with the default problem. Privacy-friendly default settings protect individuals' privacy and personal data more effectively [2]. On the other hand, privacy-invasive default settings have a bias towards data sharing [76]. Managing privacy online can be complex and often people do not change defaults or use granular privacy settings [76]. Thus, default settings — whether privacy-protective or privacy-invasive — may not be representative of a person's intention. A person may not even know of a default setting. This argument forms the basis for why DNT signals are ignored by most sites receiving them, especially, after Microsoft turned on DNT by default on its Internet Explorer 10 [19]. Turning on DNT by default had the effect that sites and ad networks were given a reason for ignoring DNT signals, further eroding the already weak legal basis of DNT. However, any solution that does not rely on defaults has to contend with the resulting usability challenges. The default problem becomes a usability problem.

3 ACTIVITY AND GENERALIZABILITY

Various privacy laws require people to make their privacy choices generally without relying on default settings.[9] People must actively engage in a choice (§3.1). This requirement creates usability challenges, especially, when it comes to site-by-site opt outs [48, 57]. Thus, a choice for one site should be generalizable towards a larger set of sites (§3.2).

3.1 Making Active Choices

Making an active privacy choice indicates people's intent to change their privacy settings. While the activity does not necessarily need to express their intent explicitly, it must allow for its inference. For example, if people turn on universal opt out controls in their browsers instead of setting opt out choices for each individual site they visit, it can be inferred that they want to opt out from all sites they visit. Similarly, if people are aware of a particular privacy feature in a browser and how to change it, it can be inferred that they want to use it when they use the browser, even if the feature is turned on by default. This idea of an implicit act that allows the inference of intent towards a particular privacy choice is expressed in the CCPA [68]: "The consumer exercises their choice by affirmatively choosing the privacy control [...] including when utilizing privacy-by-design products or services." It is a further dimension of active choice whether people make a prompted or an unprompted choice:

- **Install Time Prompts**: People can make a choice at install time, for example, during the setup of a new browser indicating their general choice for all websites they visit in the future.

[6] US consumers used an average of 46 apps each month in the first half of 2021 [10].
[7] California Business and Professions Code §22575(a)(5).
[8] Notably, the Office of the California Attorney General brought an enforcement action and entered into a settlement agreement with fashion retailer Sephora for failing to disclose the selling of personal information and not processing opt outs via global privacy controls [69].
[9] See, for example, the California Civil Code §1798.185(a)(19)(A)(iii), Colorado Privacy Act §6-1-1313(2)(c), and Connecticut Data Privacy Act §6(e)(1)(A)(ii)(II).

- **Run Time Prompts**: People can make a choice at run time, for example, in a browser at the time of visiting an individual website and being prompted for their choice on that site.
- **Unprompted Settings**: People can access their privacy settings at any time and make their choices.

3.2 Generalizability of Choices

Once a user has made a privacy choice, it can be generalized along multiple dimensions.

3.2.1 Vertical and Horizontal Generalizability. Vertical generalizability refers to the layer of abstraction in terms of its depth. For example, a privacy choice can be generalized from the operating system layer to the application layer so that one choice covers multiple applications. On the other hand, horizontal generalizability refers to the breadth of a privacy choice within a layer of abstraction. For example, a privacy choice can be generalized in a web browser from one site to a larger set of sites. Settings based on horizontal generalizability can have breadth across browsers on different devices. Cookie consent fatigue, as a consequence of the deluge of cookie banners on websites [12], could be addressed with horizontal generalizability. We focus our inquiry here on horizontal generalizability.

3.2.2 Direct and Indirect Generalizability. Direct generalizability means that people can generalize privacy choices explicitly. For example, a browser could ask people to make their privacy choices for one site with an option to apply their choices towards all future sites they visit. Indirect generalizability, on the other hand, refers to privacy choices that are not directly generalized from people's choices. For example, a browser could ask people to select categories of websites from which they want to opt out. Then, if a visited site belongs to one of the selected categories, people would be opted out from tracking on that site. Another example of indirect generalizability would be the use of an automated agent that makes privacy choices for an individual based on learned preferences from the individual's previous choices.

3.2.3 Generalizability and Individualizability. Generalizability of a privacy choice is based on the idea of deepening or broadening the initial scope of a choice towards a larger set of choices. The opposite of a generalizable privacy choice is an individualizable privacy choice. Individualizability means to establish a general rule first and later add exceptions to it for a smaller set of privacy choices. For example, people could universally opt out from all sites they visit and then refine their opt out to exclude certain sites, in other words, individualize certain sites on which they want to remain opted in. They could also do the opposite: opt in generally and opt out for certain sites.

3.2.4 Usability beyond Generalizability. We find that generalizability can improve usability (§5). However, generalizability of a privacy choice is not a sufficient condition to achieve maximum usability. Additional usability considerations may be required. There may be also counteracting factors that reduce the effectiveness of a generalizable privacy choice design. Generalizability is not a comprehensive solution that, once implemented, solves all usability problems. Rather, it is one design principle that can be applied together with others.

3.3 Privacy Choice Schemes

We study how various dimensions of active choice and generalizability impact the usability of GPC. To that end, we designed a set of privacy choice schemes (**schemes**), which implement various active choice and generalizability features, in a browser extension that sends GPC signals based on people's choices.[10] Our set of schemes is not meant to be comprehensive but rather representative of core dimensions of generalizable active privacy choice. The user interface for each of our schemes is shown in Appendix A.3. The schemes we designed and evaluated are the following:

- ***SB-Base***: Baseline scheme with opt out banner. Participants are prompted for a choice via an opt out banner on each new site they visit. *SB-Base* requires from participants the highest level of activity among all our schemes. It does not have any generalizability feature and, thus, serves as the base treatment in our evaluation.[11]
- ***S0-Snooze***: Extended baseline scheme with opt out banner and snooze feature. Participants are prompted for a choice via an opt out banner on each new site they visit. They can snooze the banner for 12 hours a time, during which no GPC signals are sent to newly visited sites. The snooze feature is not a generalizability feature. Rather, it is used here to study the extent to which participants intentionally made an opt out choice as opposed to just clicking away the opt out banners they were presented without caring about their choices (§5.3).[12]
- ***S1-Apply-all***: Banner scheme with apply-all feature. Participants are prompted for a choice via an opt out banner on each new site they visit with the addition of an apply-all feature. This feature allows participants to generalize their choice for the current site at any time towards all future sites they visit. The apply-all feature in *S1-Apply-all* is one possible implementation of direct, horizontal generalizability.
- ***S2-Snooze+Apply-all***: Banner scheme with apply-all and snooze features. This scheme is a combination of *S0-Snooze* and *S1-Apply-all*. The opt out banner contains both the apply-all as well as snooze features to compare the relative effect of each as well as participants' intent to opt out.
- ***S3-Profile***: Privacy profile category scheme. Participants are prompted at install time to choose a privacy profile: high, medium, or low privacy-sensitivity. For the high and low privacy-sensitivity profiles the extension opted participants out of data sharing on all and no sites, respectively. Medium privacy-sensitivity opted

[10] Our browser extension with scheme implementations is available at https://github.com/privacy-tech-lab/gpc-privacy-choice. For real-world use cases the schemes should be implemented directly in a browser. Install features of the schemes would be shown in the browser at install time while banner features would be shown upon visiting websites (Table 2).

[11] *SB-Base* is not a scheme in the strict sense because it lacks a generalizability feature.

[12] *S0-Snooze* is not a scheme in the strict sense because it lacks a generalizability feature and allows inactivity. A modified active and generalizable version of the scheme could be to send the most recent choice — GPC being turned on or off — to all newly visited sites until the end of the snooze period, at which time choice prompts start again.

participants out on sites whose domain was included in the Disconnect Tracker Protection lists [17]. *S3-Profile* is a form of indirect, horizontal generalizability utilizing profiles as abstraction layer for determining opt out choices.

- ***S4-Website***: Website category scheme. Participants are prompted at install time to select website categories that they would like to opt out from. Their choices will determine from which sites they will be opted out. *S4-Website* also makes use of indirect, horizontal generalizability, although, in a more granular form than *S3-Profile*. Website categories are based on Disconnect's Tracker Protection lists [17].
- ***S5-Learn***: Learning scheme. Participants are prompted for a choice via an opt out banner on the first ten new sites they visit. Their choices are then used to learn the privacy profile per *S3-Profile* that best matches their choices, which they were then told. Participants were assigned profiles depending on their number of opt outs on the initial ten sites: eight or more opt outs lead to high privacy-sensitivity, four to seven opt outs to medium privacy-sensitivity, three or fewer opt outs to low privacy-sensitivity. The learning feature of *S5-Learn* allowed us to investigate indirect, horizontal generalizability in the form of an agent. Setting the learning threshold at ten sites is intended to balance accuracy, dependent on the number of choices to extrapolate from, and keeping the learning period short. Higher thresholds run the risk that participants visiting only a small set of sites would leave the learning period only late in the study.
- ***S6-Universal***: Universal category scheme. *S6-Universal* is the simplest of all schemes we study. Participants are prompted at install time whether they would like to opt out from all sites they visit or not. They may change this universal preference through the settings page. *S6-Universal* represents the most extreme form of indirect, horizontal generalizability. It also allows us to explore individualizability in that participants can exempt individual sites from their general choice if they wish to do so.
- ***S7-Data***: Data category scheme. Participants are prompted at install time to select the categories of data that should not be shared or sold. *S7-Data* is distinct from *S3-Profile* and *S4-Website* in that it lets us evaluate indirect, horizontal generalizability as it occurs along the dimension of disclosed data categories.

For all schemes except *S7-Data* participants could also adjust their opt out settings for each domain individually via a domain list (Appendix A.3, Figure 19). As participants browsed, each visited first party domain was appended to the domain list, on which participants could change whether it should receive GPC signals. For *S5-Learn*, the domain list feature became available after the learning period was finished. Schemes *SB-Base*, *S0-Snooze*, *S1-Apply-all*, and *S2-Snooze+Apply-all* required participants to make at least one of their privacy choices on an individual site via a choice banner (**banner schemes**, as shown in Appendix A.3, Figure 13). Schemes *S3-Profile*, *S4-Website*, *S6-Universal*, and *S7-Data* immediately generalized privacy choices based on categories (**category schemes**, as shown in Appendix A.3, Figures 14, 15, 17, and 18). *S5-Learn* is a mixed banner-category scheme as participants were shown opt out banners during the learning period and assigned a privacy profile afterwards (Appendix A.3, Figure 16).

The design space for the set of privacy choice schemes we cover here is motivated by existing real-world challenges. Currently, GPC is only available in browsers' and extensions' settings menus rather than being surfaced via choice prompts upon site visits or at installation time. GPC is further only available for a limited set of browsers and extensions. There, consequently, exists an imperative to find usable choice interfaces. *SB-Base* and *S6-Universal* are extensions of current opt out interfaces on websites with opt out links and browsers that support DNT, respectively. *S3-Profile*, *S4-Website*, and *S7-Data* are motivated by the notion that opt out interfaces should not be all-or-nothing choices but rather allow for more nuanced choices. As people may not want to make lots of individual choices, *S1-Apply-all* and *S5-Learn* provide an exploration into reducing choices made by directly asking people and, in case of *S5-Learn*, supplementing people's choices with choices by an automated agent. *S0-Snooze* and *S2-Snooze+Apply-all* are motivated by banner fatigue and used here to evaluate participants' intent to opt out.

3.4 Use Cases

We are evaluating generalizable active privacy choice in the context of opting out from web tracking via GPC. However, many other use cases exist. First, other privacy preference signals could make use of generalizable active privacy choice. Given the prevailing cookie banner fatigue, it could also be used for cookie consent interfaces. Cookie banner opt outs do not allow people to generalize their choices beyond the site on which they make their choices. Another use case could be to withdraw consent or object to processing of personal data per the GDPR. For example, the objection to processing of personal data for direct marketing purposes could be generalized.[13] Setting browser permissions, for example, whether to allow sites access to an individual's location, use of the microphone, or video could be another use case.

4 EXPERIMENTAL SETUP

We evaluated generalizable active privacy choice in a usability study with 410 participants. To that end, we implemented the nine schemes — *SB-Base* to *S7-Data* — in a browser extension for Google Chrome and other Chromium-based browsers.[14]

4.1 GPC Privacy Choice Browser Extension

We used our extension to inject opt out banners, settings, and other scheme features into the browsers and websites that our study participants used and visited. In conjunction with a backend database our extension collected participants' browsing history as well as the following extension interaction data:

- **Site Interaction History** covers privacy choices made via GPC privacy choice banners on *SB-Base*, *S0-Snooze*, *S1-Apply-all*, *S2-Snooze+Apply-all*, and *S5-Learn*. It also covers any individual site choices made using the domain list, which was available for all schemes except for *S7-Data*.

[13] GDPR Art. 21(3).

[14] Our browser extension with scheme implementations is available at https://github.com/privacy-tech-lab/gpc-privacy-choice.

- **Privacy Configuration Interaction History** covers the privacy choices participants made in *S3-Profile* to *S7-Data*, e.g., for *S3-Profile*, the initial privacy profile and any later modifications.
- **Snooze Interaction History** covers schemes with a snooze button, i.e., *S0-Snooze* and *S2-Snooze+Apply-all*. It recorded when participants chose to use the snooze button.[15]

In order to evaluate the usability of the different schemes in a uniform manner we kept the language and user interface design uniform unless a scheme inherently required a deviation, e.g., *S0-Snooze* and *S2-Snooze+Apply-all* required a snooze button while the remaining schemes did not. We also aimed to keep the language in all schemes uniform and neutral to avoid nudging participants to make a particular privacy choice. For example, the choice banner uses "Opt in" and "Opt out" instead of "Yes" and "No" with regard to allowing tracking because the latter tends to elicit "Yes" answers as most people prefer to agree [41].

4.2 Study Procedure and Eligibility Criteria

We recruited participants for our study on the crowd-working platform Prolific [60]. Our study consisted of three parts: (1) a sign-up survey, (2) browser extension use, and (3) an exit survey. Upon sign-up we informed participants of who we are and how they could contact us. We explained that the purpose of our study is to find out whether people understand what GPC is and how they would use it. We provided participants a complete list of data categories we would collect from them. We explained that the data would be stored at our institution and its service providers using current best practices, that it would not be disclosed except in aggregate form, and that we would retain a copy after the study for record-keeping purposes. We informed participants that they could have a particular piece of data deleted at any time and withdraw from our study at any time for any reason. We received approval for our study from our institution's IRB. Prolific also confirmed that our planned study complied with their policies.

The eligibility criteria to participate in our study were: (1) use of a Chromium browser as default browser on a laptop or desktop computer, (2) US residency, (3) fluency in English, (4) 100% approval rate for previous tasks on Prolific, (5) completion of at least 30 previous tasks on Prolific, and (6) a minimum age of 18 years. To ensure that criteria (2) – (6) were met we relied on the information provided by Prolific. For each participant, Prolific also provided demographic data. We initially signed up around 65 participants per scheme. We later excluded participants from our dataset if they did not follow our instructions, including if they did not install our browser extension, we did not receive sufficient browsing data, or they did not participate in our exit survey.

We asked participants to install our browser extension citing our purpose as studying the usability of a new privacy feature in web browsers called Global Privacy Control (GPC). Participants installed our extension via the Chrome Web Store using an activation password we provided during sign-up, which we used to prevent non-participants from interfering with the study. Upon installation, participants received an explanation of GPC.[16]

[15] A detailed list of all categories of data we collected from the participants is shown in Appendix A.1. Some categories of data we evaluated did not lead to relevant findings for our purposes and, consequently, are not discussed any further.

[16] The GPC explanation is shown in Appendix A.2 and is based on earlier work [81].

Age Range		Sex		Race/Ethnicity		Student		Employment	
18–24	16%	Male	51%	White	73%	Yes	20%	Full-Time	47%
25–34	35%	Female	49%	Asian	10%	No	80%	Unpaid work	17%
35–44	24%	Other	<1%	Black	9%			Unemployed	15%
45–54	12%			Mixed	7%			Part-Time	11%
55–64	10%			Other	2%			Recently Hired	2%
>65	4%							Other	9%

Table 1: Participant demographics. 410 participants signed up for our study with a mean of 46 participants per scheme: *SB*: 36, *S0*: 46, *S1*: 43, *S2*: 47, *S3*: 47, *S4*: 49, *S5*: 40, *S6*: 62, *S7*: 40. Some participants did not provide data for all categories: <1% for Age Range, 0% for Sex, <1% for Race/Ethnicity, 10% for Student status, and 17% for Employment status. Percentages are adjusted to account for any omissions.

Participants were also prompted for their Prolific IDs, which are pseudonyms by which we identified participants. We used the Prolific ID to associate each participant's browsing data with their answers to the exit survey, where we also asked for this ID. For schemes with an initial privacy choice configuration, e.g., *S3-Profile*, participants were prompted to set their configuration. Once we started receiving data, we informed each participant via a message on Prolific that the extension was properly configured and that they should browse the web as usual for a week.

Our extension sent the interaction data from each participant to a Firestore database. The inflow of data was monitored daily. We inquired with participants who had insufficient browsing activity and prompted them, as necessary, to make normal use of their browser. If a participant had less than thirty browsing entries within the first three days of the study, we contacted the participant on the fourth day via a Prolific message. If we did not receive a response or the lack of activity continued, we did not include the participant's data in the dataset. Participants whose data we included in the dataset had our extension running for a median of 7.1 days with a standard deviation of 2.2 days.

We collected data one scheme at a time. After finishing the data collection for a scheme, we gave participants of that scheme a week to fill out an exit survey in which we asked them for their opinions on web privacy and the usability of their opt out experience. The survey questions were the same for participants across all nine schemes.[17] Participants spent a median of eight minutes completing the exit survey. We included an attention-check question that was correctly answered by 410 participants. We excluded from our dataset the data of one participant who answered the attention-check question incorrectly. Each participant could only participate once. For their participation we paid each participant $10, the amount recommended by Prolific for our study.

4.3 Sample Representativeness

Our participant sample is partially representative of the US population as to participant demographics and technologies used.

4.3.1 Demographics. Comparing our participant sample (Table 1) to the 2020 American Community Survey and Census data [5, 7], we find our sample to be largely representative of the US population in terms of student status and sex. Regarding ethnicity, however, we observe differences: Black participants (sample 9%; population

[17] The complete set of survey questions is shown in Appendix A.2.

Scheme	SB-Base	S0-Snooze	S1-Apply-all	S2-Snooze+Apply-all	S3-Profile	S4-Website	S5-Learn	S6-Universal	S7-Data
Generalizability	x (Baseline)	x (Baseline)	Horizontal Direct	Horizontal Direct	Horizontal Indirect	Horizontal Indirect	Horizontal Direct (During Learning) Indirect (After Learning)	Horizontal Direct Individualizability	Horizontal Indirect
Scheme Type	Banner Scheme	Banner Scheme	Banner Scheme	Banner Scheme	Category Scheme	Category Scheme	Banner (During Learning) Category (After Learning)	Category Scheme	Category Scheme
Banner Timing	Every New Site	Every New Site	Every New Site	Every New Site	x	x	First 10 New Sites	x	x
Install Time Prompt	x	x	x	x	Privacy Profile Choice	Website Category Choice	x	Universal Choice	Data Category Choice
Run Time Prompt	Site Choice	Site Choice Snooze Banner	Site Choice Apply-all Choice	Site Choice Snooze Banner Apply-all Choice	x x x	x x x	Site Choice	x x x	x x x
Unprompted Settings	Domain List	Domain List	Domain List Apply-all Choice	Domain List Apply-all Choice	Domain List Privacy Profile Choice	Domain List Website Category Choice	Domain List Privacy Profile Choice	Domain List Universal Choice	Data Category Choice

Table 2: Schemes and their features. Study participants who were assigned the baseline scheme — *SB-Base* — were prompted for their GPC choice via a banner on every new website they visited. *S0-Snooze* is our extended baseline scheme that allowed participants to snooze banners for 12 hours a time. All other schemes tested various types of active choices (§3.1) and their generalizability (§3.2). All schemes, except *S7-Data*, allowed participants to make GPC choices for individual sites via a domain list on the settings page. For *S1-Apply-all* and *S2-Snooze+Apply-all* the domain list also had an apply-all feature. Features not available for a scheme are denoted by an x.

13% [5]) are underrepresented while Asian (sample 10%; population 6% [5]) and White (sample 73%; population 70% [5]) participants are overrepresented. Furthermore, the percentage of participants who were unemployed was noticeably higher (sample: 15%; population: 3% [6]), which may be a natural consequence of our recruitment of participants on a crowd-working platform.

4.3.2 Technologies. Comparing our participant sample to data provided by Statcounter [66, 67]), an online analytics resource, we note that our sample skews towards a larger share of Windows participants (sample 75%; population 66% [67]) and a smaller share of macOS participants (sample 19%; population 25% [67]). Further, as participation in our study was restricted to participants running a Chromium-based browser on a laptop or desktop computer as their default browser (§4.2) most participants were using the Chrome browser. In fact, the percentage was even higher than Chrome's market share among Chromium-based browsers (sample 90%; population 81% [66]), in particular, at the expense of the underrepresented Edge browser (sample 5%; population 19% [66]).[18] Interestingly, the share of participants using Brave (5%) significantly exceeded the browser's market share, for which Statcounter does not provide a figure due to its small user base. We note a high degree of diversity in each participant's browsing history. Overall, participants visited a median of 78 unique sites per week, and 90% visited at least 20 unique sites. Participant activity along this metric did not differ substantially across schemes.[19]

4.4 Limitations

Our study is subject to various limitations. While we believe that our findings provide an indicator for how people would perceive the different schemes we discuss here, our participant sample is relatively small and not fully representative of the US population in terms of demographics and technology use. Further, the participants in our sample may be less privacy-conscious than the average person on the web as it was a condition for participating in our study to allow the collection of browsing history and other data. On the other hand, the outsized prevalence of Brave, as a privacy-protective browser, may be an indicator that at least 5% participants in our sample care about online privacy. In any case, the privacy leanings of our participants one way or the other could have influenced the study results.

It should also be noted that our usability study is focused on the user interface interactions with the schemes and not on the effects of opting out, such as seeing generic vs personalized as. As GPC is not yet broadly adopted by sites and its enforcement is in its initial stages, participants' choices did not affect their browsing experience in a major way. Our explanation of GPC upon extension installation also noted that whether or not a site respects GPC depends on local law. While we do not have any evidence that participants' behavior could have been influenced by residing in a state that has not yet adopted GPC, we cannot exclude it. If residency affected participants' behavior, we would expect such behavior to be distributed equally across schemes.

Finally, our implementation of *S7-Data* does not send GPC signals because it would require detecting which data categories sites are sharing. Such detection goes beyond our work here.[20] *S7-Data* were made aware that their extension would not send GPC signals. However, since the sending of GPC signals is not transparent, the user interface is the same independently of whether GPC signals are sent or not. Indeed, when asked in the exit survey about the confidence with which participants felt that their opt out choices were honored, there was no statistically significant difference between the answer distribution of *S7-Data* participants and those of all other schemes. This result suggests that cases of GPC being non-functional impacted neither participants' perception of GPC nor their browsing behavior to a meaningful extent.

5 USABILITY EVALUATION

At the core of our inquiry stands the development of a usable GPC privacy choice interface for the web that does not rely on default settings. Thus, we evaluate how active choices on one site can be generalized towards larger sets of sites in a usable way. Our evaluation is based on (1) browsing history and extension interaction data

[18] The percentages for each browser are adjusted in relation to its market share among Chromium-based browsers.

[19] A detailed breakdown of participant browsing statistics can be found in Appendix A.4.

[20] However, such functionality is generally feasible. For example, the Privacy Pioneer browser extension detects which data categories sites are sharing [59].

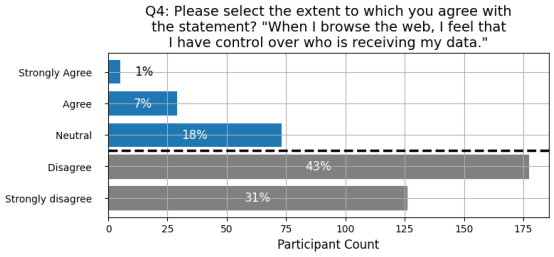

Figure 1: A majority of the 410 participants in our study across schemes expressed a lack of control over who is receiving their data.

collected from our study participants' use of our browser extension and (2) their responses in our exit survey.[21] Table 2 shows the schemes and their features that we implemented in our browser extension. For our usability evaluation we consider a broad spectrum of factors [25]: user needs and sentiment (§5.1), effort and ability (§5.2), awareness, comprehension, and intent (§5.3), and nudging patterns and decision reversal (§5.4).

5.1 User Needs and Sentiment

Many people would like to have more control over what marketers can learn about them online [72]. Indeed, 74% of participants in our study agreed or strongly agreed that they do not have control over their data when they browse the web (Figure 1). The perceived lack of and desire for more control motivate the need for an efficient and effective privacy choice interface. To that end, our results suggest that generalizability slightly decreases opt out utility (§5.1.1) but increases opt out efficiency (§5.1.2) and makes opting out less disruptive (§5.1.3), which is more important to many participants than opt out utility (§5.1.4).

5.1.1 Generalizability Decreased Opt out Utility Slightly. Overall, participants across schemes expressed that they could make their opt out choices the way they wanted (Figure 2). However, we do observe statistical variation between schemes overall (Kruskal-Wallis test, $p \approx .00033$). In particular, there are significant differences (Corrected Dunn test, $p < .05$) for individual pairwise comparisons between the distributions of our baseline schemes, which do not include any generalizability features, and schemes that include such.[22] Most notably, there are significant differences between *SB-Base*, and, individually, *S2-Snooze+Apply-all* and *S6-Universal* as well as between *S0-Snooze* and, individually, *S2-Snooze+Apply-all*, *S3-Profile*, *S4-Website*, *S5-Learn*, and *S6-Universal*.[23] Participants who were assigned a baseline scheme have a higher rate of strong agreement on being able to making their choices the way they wanted. This finding is plausible as these schemes leave less room for mental and temporal disconnect. They required participants to make their privacy choices for the individual sites they visited at the time they visited them.

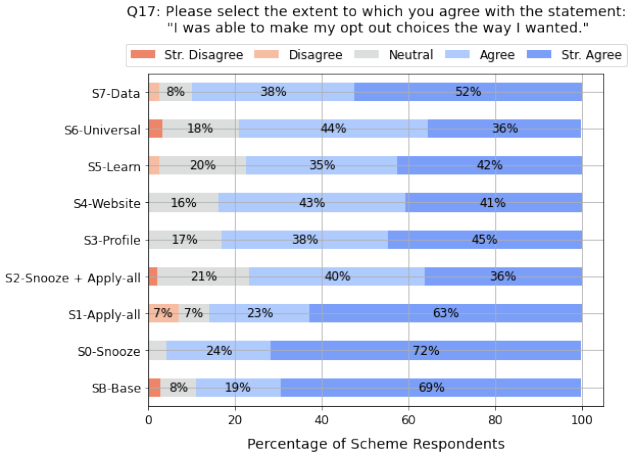

Figure 2: *SB-Base*, *S0-Snooze* and *S1-Apply-all* have relatively higher rates of participants who strongly agreed that they could make their choices the way they wanted. However, the differences are a matter of degree as most participants across schemes agreed or strongly agreed that they could make their opt out choices the way they wanted.

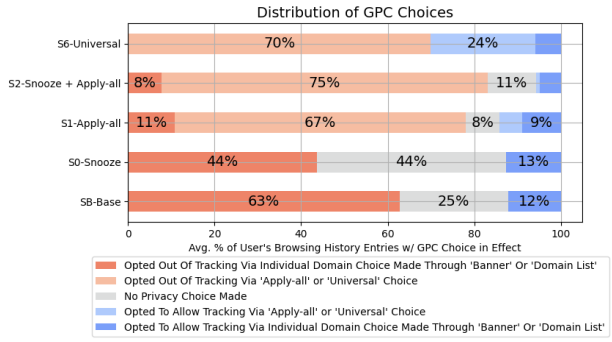

Figure 3: Participants' GPC privacy choices in schemes with direct generalizability compared to the baseline schemes. The latter required participants to make choices for individual sites. *S0-Snooze* participants could also snooze banners for 12 hours a time. Sites without privacy choices being made are sites for which participants snoozed the banner or sites they were visiting for the first time without the apply-all or universal choice enabled, in which case we had not yet recorded their choice.

Comparable to our baseline schemes — *SB-Base* and *S0-Snooze* — *S1-Apply-all* also exhibits a higher rate of strong agreement among participants on being able to make their opt out choices the way they wanted. While this higher rate is not statistically significant, the trend towards strong agreement in all three schemes is noticeable. A possible explanation could be that all three schemes are perceived similarly due to their nature as banner schemes. All of them require an individual choice unless the snooze (*S0-Snooze*) or apply-all (*S1-Apply-all*) features are used. Interestingly, the fourth banner scheme — *S2-Snooze+Apply-all* — does not exhibit the trend of

[21] A detailed list of all data categories collected from participants is shown in Appendix A.1. The complete set of survey questions is shown in Appendix A.2.

[22] For all pairwise comparisons between the nine schemes in this paper, we use the post-hoc Dunn test and apply the Benjamini-Hochberg correction for multiple testing. The Benjamini-Hochberg correction was designed to reduce the high false discovery rate, that is, the chance of a rejected null hypothesis being a false positive, associated with doing many comparisons in sequence [51]. Applying the Benjamini-Hochberg correction generally produces more stringent p-values and fewer null hypothesis rejections.

[23] The full set of p-values, corrected and uncorrected, is shown in Appendix A.5, Table 3.

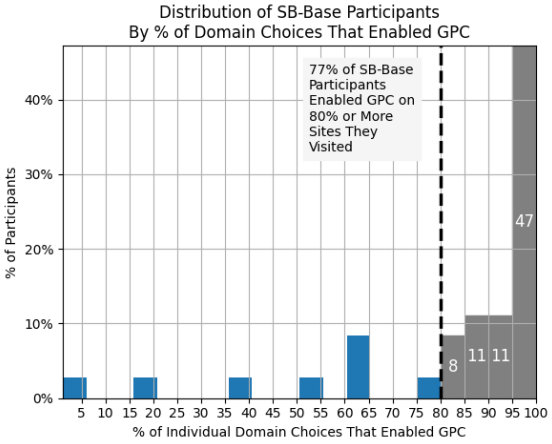

Figure 4: The opt out choices of participants in the *SB-Base* scheme, which required a choice on every newly visited site, were fairly homogeneous with most participants opting out on most sites.

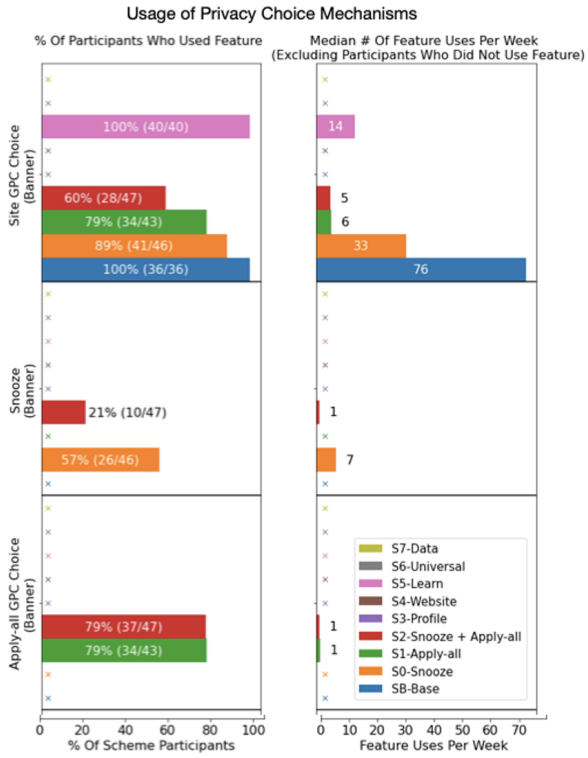

Figure 5: Percentage of participants in each scheme making use of available scheme features at least once (left) and median number of feature uses normalized to a one week period (right). Percentages and number of uses exclude required interactions. Features not available for a scheme are denoted by an x.

strong agreement. Perhaps, the combined impact of the snooze and apply-all features aligns *S2-Snooze+Apply-all* with the trend we see for the other schemes with generalizability features.

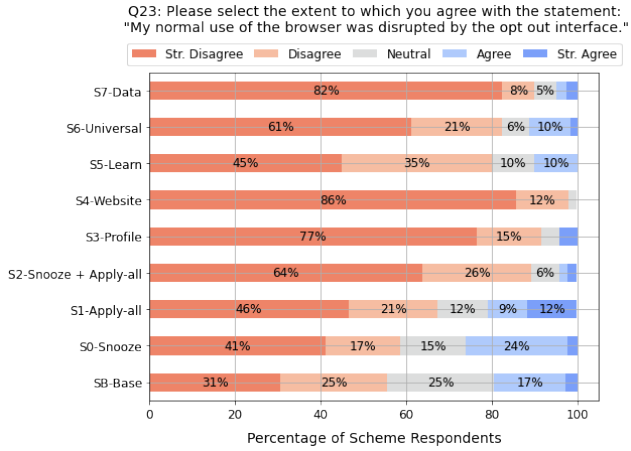

Figure 6: The perceived level of disruption is lower for schemes with generalizability features than for baseline schemes and, in a less pronounced trend, for banner schemes overall. For the least disruptive scheme — *S4-Website* — 98% participants reported that they did not feel disrupted, a 40% point increase over the baseline schemes.

5.1.2 Generalizability Increased Opt out Efficiency. When given the option to generalize their privacy choices directly, *S1-Apply-all* participants made use of it for 67% of their choices, *S2-Snooze+Apply-all* participants for 75% of their choices, and *S6-Universal* participants for 70% of their choices (Figure 3). Despite the lack of generalizability features, we observe the same opt out trend for *SB-Base*. The majority of *SB-Base* participants opted out on most sites they visited. Specifically, 77% of *SB-Base* participants chose to enable GPC on 80% or more sites they visited (Figure 4). Thus, instead of requiring them to opt out on every new site individually, a scheme with generalizable choices would have been more efficient for most *SB-Base* participants. Medians of 6 banner interactions for *S1-Apply-all* participants and 5 for *S2-Snooze+Apply-all* participants compare favorably to the 76 for *SB-Base* participants (Figure 5).

Providing generalizability features does not mean that there is no room for fine-grained privacy choices — quite the contrary. A domain list to fine-tune privacy choices was of value to a significant minority of participants and complemented their use of generalizability features. Depending on the scheme, between 2% and 33% of participants made privacy choices for a median of 1 to 10 individual sites using the domain list. While participants rarely used individual choice features broadly unless they had to, such features can prove useful to add nuance to the overall choice configuration. Providing individualizability features in addition to generalizability features does not increase the level of browsing disruption as people can also choose to not use them.

5.1.3 Generalizability Made Opting Out Less Disruptive. The degree of agreement on the disruption of normal browser use is scheme-dependent and statistically significant (Kruskal–Wallis test, p<.001). Overall, participants perceived schemes with generalizability features as less disruptive when compared to baseline schemes. Figure 6 shows participants' exit survey responses when asked

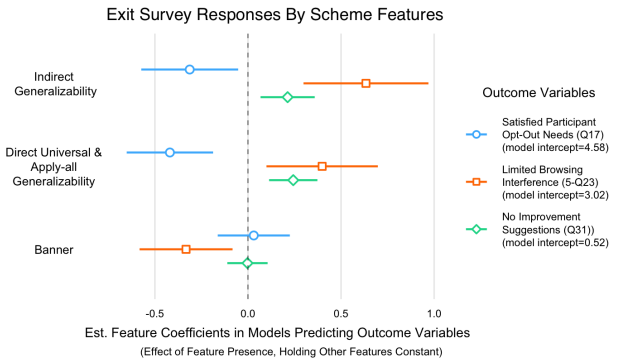

Figure 7: Generalizability features decreased participants' opt out utility, i.e., the sense of being able to make choices the way they want, but increased their sense of no browsing disruption. Illustrated are the 95% confidence intervals of the coefficients on variables representing the presence of the specified scheme features. Q17 and Q23 responses are based on a 5-point Likert scale while responses to Q31 are based on a Boolean variable representing whether a participant had no improvement suggestions. Q23 responses are inverted, 5-Q23, so that, as for the other two outcome variables, positive values represent positive participant perceptions. Replacing Q31 responses with responses to Q15, which asked participants for change suggestions, leads to a similar result.

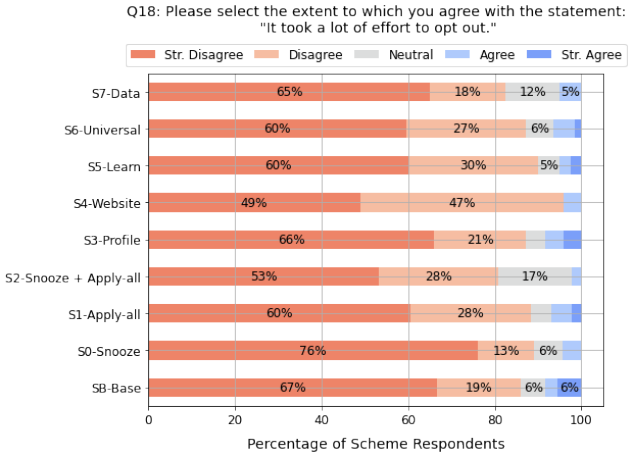

Figure 8: Participants' responses on the effort required to opt out. Across schemes most participants perceived the effort as low.

about the level of disruption they experienced. In particular, participants found schemes with indirect generalizability — *S3-Profile*, *S4-Website*, and *S7-Data* — to be the least disruptive. This result is plausible because participants assigned those schemes only needed to make an initial choice while further interactions, for example, privacy profile changes or manual fine-grained choices at the site-level via the domain list, were entirely optional.

Interestingly *S5-Learn*, a scheme that features indirect generalizability, differs significantly from each of the other schemes that feature indirect generalizability (Corrected Dunn test, p<.05).[24] This result is especially noteworthy as the only difference between the *S3-Profile* and *S5-Learn* schemes is the learning period. This period is the only time when *S5-Learn* participants could not generalize their choices. It lasted a scant median of 7.85 hours per participant out of the week-long study period. However, it could be that requiring participants to interact with choice banners for this period led to an increase in perceived browsing disruption and superseded the non-disruptive nature of indirect generalizability.

Participants experienced schemes without generalizability — *SB-Base* and *S0-Snooze* — as the most disruptive. A statistically significant difference (Corrected Dunn test, p<.05) exists between the distributions of each of *SB-Base* and *S0-Snooze* and, individually, *S2-Snooze+Apply-all*, *S3-Profile*, *S4-Website*, *S6-Universal*, and *S7-Data*. The *S1-Apply-all* scheme was also perceived as relatively disruptive. A statistically significant difference (Corrected Dunn test, p<.05) exists between the distributions of *S1-Apply-all* and, individually, *S2-Snooze+Apply-all*, *S3-Profile*, *S4-Website*, and *S7-Data*. Thus, banner schemes are generally perceived as more disruptive than category schemes. Despite being a banner scheme, *S2-Snooze+Apply-all* was perceived as less disruptive when compared to *S1-Apply-all*. A reason could be the increased disruption-reducing feature usage. While usage rates for the apply-all feature in both schemes did not differ with 79% each (Figure 5), 91% of participants in *S2-Snooze+Apply-all* made use of either the snooze or apply-all feature.

5.1.4 Less Disruption Was More Important than Opt out Utility. While participants assigned to schemes with generalizability features were less likely to strongly agree with the statement "I was able to make my opt out choices the way I wanted" (Figure 2), they were more likely to strongly disagree with the statement "My normal use of the browser was disrupted by the opt out interface" (Figure 6). These diverging relationships can be illustrated through linear regression models (Figure 7). The presence of a generalizability feature in a scheme significantly decreased the expected browsing disruption experienced by the participants of the scheme. At the same time, these schemes had lower predicted responses as to their utility. However, the presence of generalizability features decreased the likelihood that participants would suggest improvements. Thus, it appears that participants found it generally more important to have less browsing disruption than engaging more deeply with the opt out choices.

The preference for making use of generalizability features instead of accepting browsing disruption can be witnessed in the significant disparity in median feature uses between schemes *SB-Base* and *S1-Apply-all* (Figure 5). Apart from the generalizability feature both schemes are identical. However, while *SB-Base* participants made a median of 76 individual site choices, *S1-Apply-all* participants only made 6 opting to generalize their choices instead. Only 21% of participants across the two schemes with apply-all feature — *S1-Apply-all* and *S2-Snooze+Apply-all* — chose to exclusively make their privacy choices via banners. 68% used the apply-all feature to make their privacy choices on more than 80% of the domains they visited. Thus, while a number of participants seem to value making specific adjustments to their GPC settings, more preferred an overall less disruptive experience by generalizing their choices.

[24] The full set of p-values, corrected and uncorrected, is shown in Appendix A.5, Table 4.

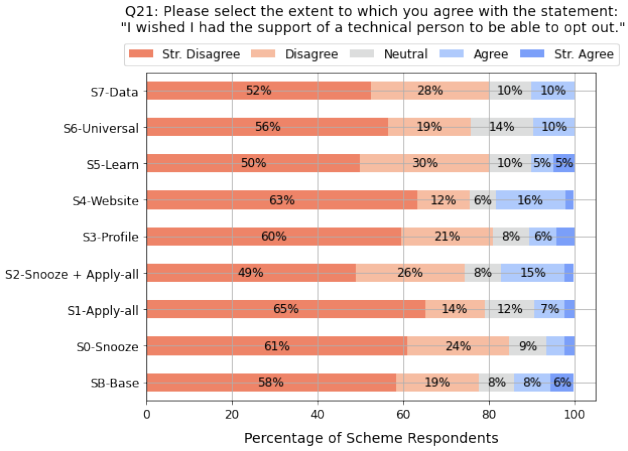

Figure 9: The extent to which participants felt they needed assistance in making their privacy choices.

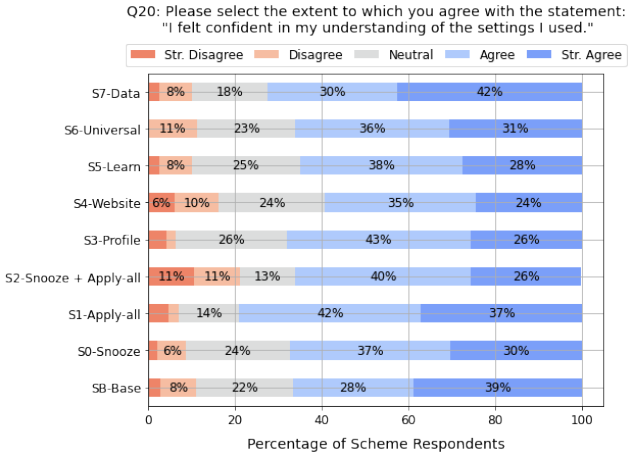

Figure 10: Participants' perceptions of how well they understood the various privacy settings.

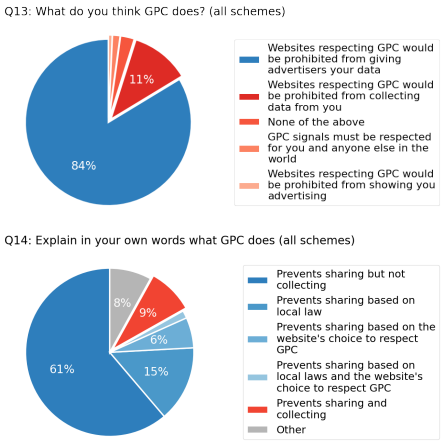

Figure 11: Participants' understanding of GPC across schemes as evaluated by responses to a multiple-choice question (Q13) and a free-response question (Q14). Correct responses are shown in blue, incorrect responses are shown in red, and responses to Q14 that could not be coded as either correct or incorrect are shown in grey.

5.2 Effort and Ability

Generally, banner schemes — *SB-Base*, *S0-Snooze*, *S1-Apply-all*, and *S2-Snooze+Apply-all* — required more interactions than category schemes — *S3-Profile*, *S4-Website*, *S6-Universal*, and *S7-Data*. Scheme *SB-Base* had the highest interaction rate by far (Appendix A.4, Figure 21). Participants had to make a privacy choice every time they visited a new site. This rate decreased in *S0-Snooze* as participants had the option to snooze banners for 12 hours a time. It dropped even more substantially in *S1-Apply-all* and *S2-Snooze+Apply-all* as participants had the apply-all feature available. Category schemes, on the other hand, required only minimal interactions.

Independently of their scheme assignment, most participants tolerated the effort that their assigned scheme required from them to make their opt out choices. In our exit survey participants across all schemes generally expressed that it did not take them a lot of effort to opt out (Figure 8). Across schemes, 87% disagreed or strongly disagreed that it took them a lot of effort to opt out. This result does not necessarily mean that all schemes had parity in terms of invested effort. Participants were asked to rate their experience against a threshold of "a lot of effort." The question did not ask participants to compare effort. However, our results suggest that generally people would not feel overburdened by making their opt out choices via any of the discussed schemes.

5.3 Awareness, Comprehension, and Intent

We evaluated how well participants understood their assigned schemes and the features they used. A majority of participants across schemes expressed that they did not need the help of a technical person (Figure 9). Most expressed confidence in their comprehension (Figure 10). These measures remain relatively similar across schemes indicating that participants were also able to comprehend the more abstract schemes that implement indirect generalizability, such as *S4-Website*. In their improvement suggestions some participants expressed a desire for an indicator that the choice mechanism is actively working. 24% of all participants across schemes suggested in their responses to Q15 or Q31 that more transparency or clarity would have improved their opt out experience. These responses suggest that ensuring people are aware of not just the features at their disposal but also of their activity could aid in enhancing confidence in privacy choice interfaces.

A lack of participants' awareness or understanding of their privacy choices could undermine their legal validity. When they installed our browser extension, we presented all participants with a standardized explanation of GPC.[25] In the exit survey we showed them this explanation again to evaluate their understanding of GPC. Our results indicate that the majority of participants understood GPC. We found no statistically significant differences between schemes. 84% of the 410 participants across all schemes provided

[25]The GPC explanation is shown in Appendix A.2 and is based on earlier work [81].

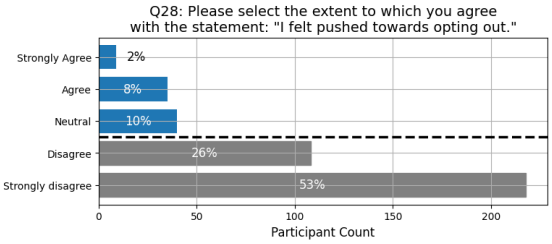

Figure 12: A majority of the 410 participants in our study across schemes expressed that they did not feel pushed towards opting out.

the correct answer to a multiple choice question that sending GPC signals under California law would prevent the selling and sharing of data, yet, would still allow first party data collection and advertising (Figure 11, Q13). The correct answer to the question was rephrased to not resemble the GPC explanation we showed participants earlier. The correct responses to the multiple choice question were confirmed by an 83% rate of participants' correct free-form responses to the question of explaining GPC in their own words (Figure 11, Q14).

A closer look at the snooze feature usage also indicates that it was the intent of most participants to express a privacy choice as opposed to just getting rid of the choice banner. 57% of participants assigned the *S0-Snooze* scheme made use of the snooze feature (Figure 5). However, the snooze rate for participants assigned the *S2-Snooze+Apply-all* scheme was only 21% while 79% of *S2-Snooze+Apply-all* participants made use of the apply-all feature instead. This rate is the same at which participants in *S1-Apply-all* — who did not have the snooze feature available — made use of the apply-all feature. This result suggests that participants actually meant to make a legally effective privacy choice using the apply-all feature.[26] 52% agreed or strongly agreed with feeling confident that their opt out choices were honored (Q22).

5.4 Nudging Patterns and Decision Reversal

Privacy choice interfaces should be as neutral as possible. A neutral design supports the legal validity of people's choices as it would not be subject to the objection of nudging them towards a choice they otherwise would not make. Thus, we phrased our explanation of GPC as well as any other user interface language as neutral as possible.[27] We also kept buttons for enabling and disabling GPC in the same size and in the same style. We then asked participants in our exit survey whether they felt pushed towards opting out. To counter the natural tendency of answering in the affirmative [41] we asked the question such that participants had to disagree with the statement "I felt pushed towards opting out." 79% of the participants disagreed or strongly disagreed suggesting that privacy choice interfaces can be designed such that people do not feel nudged (Figure 12). Participants' sense of not being nudged appears consistent across schemes as we could not establish statistically significant differences between them.

[26] We have no indicator that the intent to make a privacy choice would be substantially different for other schemes.
[27] The GPC explanation is shown in Appendix A.2 and the user interface language in Appendix A.3.

Participants across schemes generally felt that they were able to correct errors with ease as well as to change their opt out choices without much difficulty. For the former, 60% of participants cited not encountering any errors. 27% agreed or strongly agreed that they were able to correct errors with ease (Q26). Likewise, 37% of participants disagreed or strongly disagreed when asked if they found it difficult to change their opt out choices, with another 40% reporting not having had the need to make any changes (Q27).

6 DISCUSSION

To help people exercising their opt out rights we have three recommendations: regulators should require publishers to honor GPC signals sent via a generalizable active privacy choice interface (§6.1), browser vendors should integrate GPC in their browsers via a generalizable active privacy choice interface (§6.2), and publishers should honor opt outs via GPC (§6.3).

6.1 Regulators Should Require Publishers to Honor GPC Signals Sent via a Generalizable Active Privacy Choice Interface

A majority of study participants expressed a lack of control over who is receiving their data (Figure 1). People want more control [72]. The right to opt out is important for control because it prevents data from entering the online ad ecosystem in the first place. Once it does, data will be disseminated downstream to ad networks and other third parties becoming more difficult to control. Thus, data breaches, data misuses, and other types of data violations can be reduced by broadening the availability of opt out rights and making choice interfaces easier to use. However, given an opt out regime, there are usability challenges to facilitate choice without default settings. To help people make their choices as efficiently and effectively as possible regulators should require publishers to honor GPC signals sent via a generalizable active privacy choice interface.

As our results demonstrate, GPC signals can be attributed legal meaning. By sending GPC signals the majority of participants understood what they were declaring to the sites they visited (§5.3). They understood that — to the extent recognized by their jurisdiction — sending a GPC signal would prevent a site from selling and sharing their data while it would still be allowed to collect data and advertise to them. There was no statistically significant difference between schemes. Comparing the usage rates of the snooze and apply-all features further indicates that participants generally had the intention to opt out when they turned on GPC as opposed to just silencing the choice banner. They did not feel pushed either; the decision to opt out was theirs (§5.4).

Our results further indicate that most participants across schemes would opt out from most sites they visit. This is true independently of whether they made individual site-by-site choices in the baseline schemes or generalized their choices in schemes with generalizability features (Figures 3 and 4). In both groups most participants opted out on most sites. When available, many participants made use of generalizability features (Figure 5). Thus, instead of requiring people to opt out on sites individually, generalizable active privacy choice would provide a more efficient opt out basis for most people. Generalizability features allowed most participants to arrive at their final opt out configuration with fewer interactions and less friction

(§5.1.2). A supplemental domain list would still allow people to make more fine-grained choices, if they so desire.

People should be allowed to opt out by selecting and using privacy-protective technologies, for example, a browser that is marketed with privacy features. As recognized by California law [68], the selection and use of such privacy-protective technologies are indicators that an individual wants to opt out, unless known otherwise. Thus, if people are notified that GPC is being turned on by default, and it is explained how they can adjust their setting, for example, by a prompt at install time, it is clear that they want to use GPC. This situation is similar to the *S6-Universal* scheme. People are made aware of a browser's GPC setting and by using the browser they confirm that they are intending to send GPC signals accordingly.

6.2 Browser Vendors Should Integrate GPC in their Browsers via a Generalizable Active Privacy Choice Interface

Regulators are responsible for determining whether GPC signals are considered valid opt outs. They also provide the broad interface requirements, e.g., whether choices can be generalized. Browser vendors, on the other hand, are responsible for designing and implementing the interfaces for their browsers per those requirements. Generalizability features for GPC should be implemented at the browser-layer and not at the site-layer because the former provides horizontal generalizability. Standardizing privacy choice interfaces in the browser rather than being left to individual sites would have the advantage of providing a uniform interface to support notification, control, and mitigating dark patterns [65].

Our results do not point to a winning scheme. All schemes had high levels of utility (Figure 2) and tolerable browsing disruption (Figure 6). A substantial percentage of participants was willing to engage with opt out banners confirming previous results [57]. While prompting participants for their choice at runtime and in the context of a site visit slightly increased opt out utility rates (§5.1.1), decreasing the level of browsing disruption, which was achieved to the greatest degree by bannerless schemes with indirect generalizability (§5.1.3), seems more important (§5.1.4). Thus, the implementation of an indirect scheme is preferable.

Browser vendors could implement a scheme like *S6-Universal* at browser install time. They should consider the trade-off between disruption and utility for their user base. The level of non-intrusiveness of a generalizability feature deserves the highest priority in making this trade-off. If browser vendors are concerned about browsing disruption, they could describe its GPC setting on its download page or via prompts at install time. To ensure that people understand what they are declaring browsers should display an explanation of GPC. While they should ensure the usability of their GPC interface, browser vendors do not need to concern themselves with the applicability or meaning of GPC as the browser is just the conduit of the signal [23]. It would be sufficient to explain that turning on GPC will have the effect of opting out the user to the extent GPC signal recipients are required to honor it under applicable law.[28]

Browser vendors should also consider that integrating GPC can have an impact on the fingerprinting surface of their browser. However, as GPC is a binary setting, the impact will be small and can be mitigated by turning on GPC by default for privacy-protective browsers.

6.3 Publishers Should Honor Opt Outs via GPC

Sites should be able to identify GPC signals and pass them downstream to the third party sites they integrate. Consent management platforms provide key-turn implementations that can help with these tasks. As a number of participants expressed a desire for an indication that the choice mechanism is actively working (§5.3), it would be helpful to implement GPC's `.well-known` resource or another feedback mechanism to notify people that a site is compliant with GPC [23]. We urge publishers that rely on targeted ads to redefine their business models and embrace the future of privacy-preserving ad serving. As previous work has shown [81], many people are fine with ads as long as they are not privacy-invasive. But change requires initiative on part of the publishers.

7 CONCLUSIONS

Increasingly, privacy laws in the US provide people a right to opt out. However, these laws require an intentional choice and generally prohibit opt out settings being turned on by default. Generalizable active privacy choice is an interface design principle for mitigating the usability challenges originating from this legal requirement. Its premises are (1) an active privacy choice that (2) can be generalized towards a larger set of choices. To help people exercising their opt out rights on the web our results support its adoption in form of a browser-layer interface in combination with GPC. As GPC adoption increases it would be interesting to study the real-world deployment of GPC interfaces, how they explain GPC, and how people's experience of the web changes depending on their choices. A more fundamental inquiry would be a critical examination of the choice regime in the US. After all, we currently do not allow opting out by default but permit a de-facto opt in by default, which seems not what most people want.

ACKNOWLEDGMENTS

We kindly thank our anonymous shepherd and reviewers for their valuable feedback. Their work helped us to develop our arguments and make this paper better. We are grateful to the Alfred P. Sloan Foundation (Program in Universal Access to Knowledge) and to the National Science Foundation (Award #2055196) for their support of this research. We also thank Wesleyan University, its Department of Mathematics and Computer Science, and the Anil Fernando Endowment for their additional support. This work would surely not have been possible without our collaborators in the GPC coalition. Conclusions reached or positions taken are our own and not necessarily those of our financial supporters, its trustees, officers, or staff. In loving memory of Rita Zimmeck, mother of the first author.

AVAILABILITY OF ARTIFACTS

Our browser extension with scheme implementations is available at https://github.com/privacy-tech-lab/gpc-privacy-choice.

[28] The GPC explanation is shown in Appendix A.2 and is based on earlier work [81].

REFERENCES

[1] Advanced Data Protection Control (ADPC). 2022. https://www.dataprotectioncontrol.org/. Accessed: October 22, 2023.

[2] Jef Ausloos, Els Kindt, Eva Lievens, Peggy Valcke, and Jos Dumortier. 2013. Guidelines for Privacy-Friendly Default Settings. http://dx.doi.org/10.2139/ssrn.2220454. *SSRN Electronic Journal* 2013, 12 (02 2013), 34.

[3] Jan M. Bauer, Regitze Bergstrøm, and Rune Foss-Madsen. 2021. Are you sure, you want a cookie? – The effects of choice architecture on users' decisions about sharing private online data. *Computers in Human Behavior* 120 (2021), 106729. https://doi.org/10.1016/j.chb.2021.106729

[4] Kerstin Bongard-Blanchy, Arianna Rossi, Salvador Rivas, Sophie Doublet, Vincent Koenig, and Gabriele Lenzini. 2021. "I Am Definitely Manipulated, Even When I Am Aware of It. It's Ridiculous!" - Dark Patterns from the End-User Perspective. In *DIS '21: Designing Interactive Systems Conference 2021* (Virtual Event, USA) *(DIS '21)*. Association for Computing Machinery, New York, NY, USA. https://doi.org/10.1145/3461778.3462086

[5] U.S. Census Bureau. 2020. American Community Survey, 2020, Demographics and Housing Estimates. https://data.census.gov/cedsci/table?q=United%20States&g=0100000US&tid=ACSDP5Y2020.DP05. Accessed: October 22, 2023.

[6] U.S. Census Bureau. 2020. American Community Survey, 2020, Selected Economic Characteristics. https://data.census.gov/cedsci/table?q=United%20States&t=Employment&g=0100000US&tid=ACSDP5Y2020.DP03. Accessed: October 22, 2023.

[7] U.S. Census Bureau. 2020. School Enrollment in the United States: October 2020 - Detailed Tables, Table 1. https://www.census.gov/data/tables/2020/demo/school-enrollment/2020-cps.html. Accessed: October 22, 2023.

[8] California Legislative Information. 2003. California Online Privacy Protection Act. https://leginfo.legislature.ca.gov/faces/codes_displayText.xhtml?division=8.&chapter=22.&lawCode=BPC. Accessed: October 22, 2023.

[9] California State Legislature. 2018. Assembly Bill No. 375. https://leginfo.legislature.ca.gov/faces/codes_displayText.xhtml?division=3.&part=4.&lawCode=CIV&title=1.81.5. Accessed: October 22, 2023.

[10] Stephanie Chan. 2021. U.S. Consumers Used an Average of 46 Apps Each Month in the First Half of 2021. https://sensortower.com/blog/apps-used-per-us-smartphone. Accessed: October 22, 2023.

[11] Farah Chanchary and Sonia Chiasson. 2015. User Perceptions of Sharing, Advertising, and Tracking. In *Eleventh Symposium On Usable Privacy and Security (SOUPS 2015)*. USENIX Association, Ottawa, 53–67. https://www.usenix.org/conference/soups2015/proceedings/presentation/chanchary

[12] Hanbyul Choi, Jonghwa Park, and Yoonhyuk Jung. 2018. The role of privacy fatigue in online privacy behavior. *Computers in Human Behavior* 81 (2018), 42–51. https://doi.org/10.1016/j.chb.2017.12.001

[13] Colorado Department of Law Consumer Protection Section. 2023. Colorado Privacy Act Rules 4 CCR 904-3. https://coag.gov/app/uploads/2023/03/FINAL-CLEAN-2023.03.15-Official-CPA-Rules.pdf. accessed: October 22, 2023.

[14] Lorrie Faith Cranor. 2021. Informing California Privacy Regulations with Evidence from Research. *Commun. ACM* 64, 3 (Feb. 2021), 29–32. https://doi.org/10.1145/3447253

[15] Lorrie Faith Cranor, Brooks Dobbs, Serge Egelman, Giles Hogben, Jack Humphrey, Marc Langheinrich, Massimo Marchiori, Martin Presler-Marshall, Joseph M. Reagle, Matthias Schunter, David A. Stampley, and Rigo Wenning. 2006. The Platform for Privacy Preferences 1.1 (P3P1.1) Specification. https://www.w3.org/TR/P3P11/.

[16] Lorrie Faith Cranor, Marc Langheinrich, Massimo Marchiori, Martin Presler-Marshall, and Joseph M. Reagle. 2002. The Platform for Privacy Preferences 1.0 (P3P1.0) Specification. https://www.w3.org/TR/P3P/.

[17] Disconnect. 2023. Tracker Protection lists. https://github.com/disconnectme/disconnect-tracking-protection. Accessed: October 22, 2023.

[18] Nico Ebert, Kurt Alexander Ackermann, and Björn Scheppler. 2021. Bolder is Better: Raising User Awareness through Salient and Concise Privacy Notices. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, Article 67, 12 pages. https://doi.org/10.1145/3411764.3445516

[19] Serge Egelman. 2012. It's The Users, Stupid! Towards User-Centered Privacy Standards by Considering Default Settings. https://www.w3.org/2012/dnt-ws/position-papers/26.pdf.

[20] Steven Englehardt, Dillon Reisman, Christian Eubank, Peter Zimmerman, Jonathan Mayer, Arvind Narayanan, and Edward W. Felten. 2015. Cookies That Give You Away: The Surveillance Implications of Web Tracking. In *Proceedings of the 24th International Conference on World Wide Web* (Florence, Italy) *(WWW '15)*. International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, CHE, 289–299. https://doi.org/10.1145/2736277.2741679

[21] European Parliament and Council of the European Union. 2016. Regulation (EU) 2016/679. https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:32016R0679&from=EN. Accessed: October 22, 2023.

[22] General Assembly of the State of Colorado. 2021. Colorado Privacy Act. https://leg.colorado.gov/sites/default/files/documents/2021A/bills/2021a_190_rer.pdf. Accessed: October 22, 2023.

[23] Global Privacy Control Group. 2022. Global Privacy Control (GPC). https://privacycg.github.io/gpc-spec/. Accessed: October 22, 2023.

[24] Louise Gröndahl. 2020. Public knowledge of digital cookies: Exploring the design of cookie consent forms. https://www.diva-portal.org/smash/get/diva2:1470723/FULLTEXT01.pdf.

[25] Hana Habib, Megan Li, Ellie Young, and Lorrie Cranor. 2022. "Okay, Whatever": An Evaluation of Cookie Consent Interfaces. In *CHI '22: Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) *(CHI '22)*. Association for Computing Machinery, New York, NY, USA, Article 621, 27 pages. https://doi.org/10.1145/3491102.3501985

[26] Hana Habib, Sarah Pearman, Jiamin Wang, Yixin Zou, Alessandro Acquisti, Lorrie Faith Cranor, Norman Sadeh, and Florian Schaub. 2020. "It's a Scavenger Hunt": Usability of Websites' Opt-Out and Data Deletion Choices. In *CHI* (Honolulu, HI, USA). Association for Computing Machinery, New York, NY, USA, 12 pages. https://doi.org/10.1145/3313831.3376511

[27] Hana Habib, Yixin Zou, Aditi Jannu, Neha Sridhar, Chelse Swoopes, Alessandro Acquisti, Lorrie Faith Cranor, Norman Sadeh, and Florian Schaub. 2019. An Empirical Analysis of Data Deletion and Opt-Out Choices on 150 Websites. In *SOUPS*. USENIX Association, Santa Clara, CA. https://www.usenix.org/conference/soups2019/presentation/habib

[28] Hana Habib, Yixin Zou, Yaxing Yao, Alessandro Acquisti, Lorrie Cranor, Joel Reidenberg, Norman Sadeh, and Florian Schaub. 2021. Toggles, Dollar Signs, and Triangles: How to (In)Effectively Convey Privacy Choices with Icons and Link Texts. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) *(CHI '21)*. Association for Computing Machinery, New York, NY, USA, Article 63, 25 pages. https://doi.org/10.1145/3411764.3445387

[29] Margaret Hagen. 2016. User-Centered Privacy Communication Design. In *Twelfth Symposium on Usable Privacy and Security (SOUPS 2016)*. USENIX Association, Denver, CO. https://www.usenix.org/conference/soups2016/workshop-program/wfpn/presentation/hagan

[30] Maximilian Hils, Daniel W. Woods, and Rainer Böhme. 2021. Conflicting Privacy Preference Signals in the Wild. https://arxiv.org/pdf/2109.14286.pdf. Accessed: October 22, 2023.

[31] Maximilian Hils, Daniel W. Woods, and Rainer Böhme. 2021. Privacy Preference Signals: Past, Present and Future. *Proceedings on Privacy Enhancing Technologies* 2021, 4 (07 2021), 249–269. https://doi.org/10.2478/popets-2021-0069

[32] Hang Hu, Peng Peng, and Gang Wang. 2019. Characterizing Pixel Tracking through the Lens of Disposable Email Services. In *2019 IEEE Symposium on Security and Privacy (SP)*. Institute of Electrical and Electronics Engineers, Piscataway, NJ, USA, 365–379. https://doi.org/10.1109/SP.2019.00033

[33] Soheil Human, Harshvardhan J. Pandit, Victor Morel, Cristiana Santos, Martin Degeling, Arianna Rossi, Wilhelmina Botes, Vitor Jesus, and Irene Kamara. 2022. Data Protection and Consenting Communication Mechanisms: Current Open Proposals and Challenges. https://harshp.com/research/publications/051-Consenting-Communication-Mechanisms. Accessed: October 22, 2023.

[34] Garrett A. Johnson, Scott K. Shriver, and Shaoyin Du. 2020. Consumer Privacy Choice in Online Advertising: Who Opts Out and at What Cost to Industry? *Marketing Science* 39, 1 (2020), 33–51. https://doi.org/10.1287/mksc.2019.1198

[35] Bailey Kacsmar, Miti Mazmudar Kyle Tilbury, and Florian Kerschbaum. 2022. Caring about Sharing: User Perceptions of Multiparty Data Sharing. In *31st USENIX Security Symposium (USENIX Security 22)*. USENIX Association, Boston, MA, 899–916. https://www.usenix.org/conference/usenixsecurity22/presentation/kacsmar

[36] Rishabh Khandelwal, Thomas Linden, Hamza Harkous, and Kassem Fawaz. 2021. PriSEC: A Privacy Settings Enforcement Controller. In *30th USENIX Security Symposium (USENIX Security 21)*. USENIX Association, Berkeley, CA, USA, 465–482. https://www.usenix.org/conference/usenixsecurity21/presentation/khandelwal

[37] Rishabh Khandelwal, Asmit Nayak, Hamza Harkous, and Kassem Fawaz. 2022. CookieEnforcer: Automated Cookie Notice Analysis and Enforcement. https://doi.org/10.48550/ARXIV.2204.04221

[38] Agnieszka Kitkowska, Mark Warner, Yefim Shulman, Erik Wästlund, and Leonardo A. Martucci. 2020. Enhancing Privacy through the Visual Design of Privacy Notices: Exploring the Interplay of Curiosity, Control and Affect. In *SOUPS*. USENIX Association, Boston, MA, USA, 437–456. https://www.usenix.org/conference/soups2020/presentation/kitkowska

[39] Simon Koch, Malte Wessels, Benjamin Altpeter, Madita Olvermann, and Martin Johns. 2022. Keeping Privacy Labels Honest. In *Proceedings on Privacy Enhancing Technologies Symposium (PoPETS 2022)*. PoPETS 2022, Sydney, Australia and Online, 486–506.

[40] Konrad Kollnig, Anastasia Shuba, Max Van Kleek, Reuben Binns, and Nigel Shadbolt. 2022. Goodbye Tracking? Impact of IOS App Tracking Transparency and Privacy Labels. In *2022 ACM Conference on Fairness, Accountability, and Transparency* (Seoul, Republic of Korea) *(FAccT '22)*. Association for Computing Machinery, New York, NY, USA, 508–520. https://doi.org/10.1145/3531146.3533116

[41] Jon A. Krosnick. 1999. SURVEY RESEARCH. *Annual Review of Psychology* 50 (1999), 552–553. https://doi.org/10.1146/annurev.psych.50.1.537

[42] Oksana Kulyk, Annika Hilt, Nina Gerber, and Melanie Volkamer. 2018. "This Website Uses Cookies": Users' Perceptions and Reactions to the Cookie Disclaimer. https://www.ndss-symposium.org/wp-content/uploads/2018/06/eurousec2018_12_Kulyk_paper.pdf.
[43] Legislature of the State of Montana. 2023. Consumer Data Privacy Act. https://leg.mt.gov/bills/2023/billpdf/SB0384.pdf. Accessed: October 22, 2023.
[44] Legislature of the State of Texas. 2023. Texas Data Privacy and Security Act. https://capitol.texas.gov/tlodocs/88R/billtext/pdf/HB00004F.pdf. Accessed: October 22, 2023.
[45] Pedro Leon, Blase Ur, Richard Shay, Yang Wang, Rebecca Balebako, and Lorrie Cranor. 2012. Why Johnny Can't Opt out: A Usability Evaluation of Tools to Limit Online Behavioral Advertising. In *CHI* (Austin, Texas, USA). Association for Computing Machinery, New York, NY, USA, 10 pages. https://doi.org/10.1145/2207676.2207759
[46] Adam Lerner, Anna Kornfeld Simpson, Tadayoshi Kohno, and Franziska Roesner. 2016. Internet Jones and the Raiders of the Lost Trackers: An Archaeological Study of Web Tracking from 1996 to 2016. In *USENIX Security*. USENIX Association, Austin, TX, 997–1013. https://www.usenix.org/conference/usenixsecurity16/technical-sessions/presentation/lerner
[47] Eryn Ma and Eleanor Birrell. 2022. Prospective Consent: The Effect of Framing on Cookie Consent Decisions. In *Extended Abstracts of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) *(CHI EA '22)*. Association for Computing Machinery, New York, NY, USA, Article 400, 6 pages. https://doi.org/10.1145/3491101.3519687
[48] Maureen Mahoney. 2020. California Consumer Privacy Act: Are Consumer's Digital rights protected? https://advocacy.consumerreports.org/wp-content/uploads/2020/09/CR_CCPA-Are-Consumers-Digital-Rights-Protected_092020_vf.pdf. accessed: October 22, 2023.
[49] Arunesh Mathur, Gunes Acar, Michael J. Friedman, Elena Lucherini, Jonathan Mayer, Marshini Chetty, and Arvind Narayanan. 2019. Dark Patterns at Scale: Findings from a Crawl of 11K Shopping Websites. *Proc. ACM Hum.-Comput. Interact.* 3, CSCW, Article 81 (Nov. 2019), 32 pages. https://doi.org/10.1145/3359183
[50] Célestin Matte, Nataliia Bielova, and Cristiana Santos. 2020. Do Cookie Banners Respect My Choice? Measuring Legal Compliance of Banners from IAB Europe's Transparency and Consent Framework. In *S&P*. Institute of Electrical and Electronics Engineers S&P 2022, Virtual Event, USA, 791–809. https://arxiv.org/pdf/1911.09964.pdf
[51] John H. McDonald. 2014. *Handbook of Biological Statistics* (3rd ed.). Sparky House Publishing, Baltimore, Maryland, Chapter 6.1, 257–263.
[52] William Melicher, Mahmood Sharif, Joshua Tan, Lujo Bauer, Mihai Christodorescu, and Pedro Giovanni Leon. 2016. (Do Not) Track Me Sometimes: Users' Contextual Preferences for Web Tracking. *PETS* 2016, 2 (2016). https://doi.org/10.1515/popets-2016-0009
[53] Toru Nakamura, Shinsaku Kiyomoto, Welderufael B. Tesfay, and Jetzabel M. Serna. 2016. Personalised Privacy by Default Preferences - Experiment and Analysis. In *Proceedings of the 2nd International Conference on Information Systems Security and Privacy, ICISSP 2016, Rome, Italy, February 19-21, 2016*, Olivier Camp, Steven Furnell, and Paolo Mori (Eds.). SciTePress, Setúbal, Portugal, 53–62. https://doi.org/10.5220/0005681100530062
[54] Arvind Narayanan, Arunesh Mathur, Marshini Chetty, and Mihir Kshirsagar. 2020. Dark Patterns: Past, Present, and Future: The Evolution of Tricky User Interfaces. *Queue* 18, 2 (April 2020), 67–92. https://doi.org/10.1145/3400899.3400901
[55] Nick Nikiforakis, Alexandros Kapravelos, Wouter Joosen, Christopher Kruegel, Frank Piessens, and Giovanni Vigna. 2013. Cookieless monster: Exploring the ecosystem of web-based device fingerprinting. In *S&P*. IEEE, San Francisco, CA, 541–555. http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=6547132
[56] Midas Nouwens, Ilaria Liccardi, Michael Veale, David Karger, and Lalana Kagal. 2020. Dark Patterns after the GDPR: Scraping Consent Pop-Ups and Demonstrating Their Influence. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 1–13. https://doi.org/10.1145/3313831.3376321
[57] Sean O'Connor, Ryan Nurwono, Aden Siebel, and Eleanor Birrell. 2020. (Un)clear and (In)conspicuous: The right to opt-out of sale under CCPA. *CoRR* abs/2009.07884 (2020), 59–72. https://arxiv.org/abs/2009.07884
[58] Oregon Legislative Assembly. 2023. Oregon Consumer Privacy Act. https://olis.oregonlegislature.gov/liz/2023R1/Downloads/MeasureDocument/SB619/Enrolled. Accessed: October 22, 2023.
[59] privacy-tech-lab. 2023. Privacy Pioneer. https://github.com/privacy-tech-lab/privacy-pioneer. Accessed: October 22, 2023.
[60] Prolific. 2023. Quickly find research participants you can trust. https://www.prolific.co/. Accessed: October 22, 2023.
[61] Ashwini Rao, Florian Schaub, Norman Sadeh, Alessandro Acquisti, and Ruogu Kang. 2016. Expecting the Unexpected: Understanding Mismatched Privacy Expectations Online. In *SOUPS*. USENIX Association, Denver, CO, 77–96. https://www.usenix.org/conference/soups2016/technical-sessions/presentation/rao
[62] Florian Schaub, Rebecca Balebako, Adam L. Durity, and Lorrie Faith Cranor. 2015. A Design Space for Effective Privacy Notices. In *SOUPS*. USENIX Association, Ottawa, 1–17. https://www.usenix.org/conference/soups2015/proceedings/presentation/schaub
[63] Aden Siebel and Eleanor Birrell. 2022. The Impact of Visibility on the Right to Opt-out of Sale under CCPA. https://arxiv.org/pdf/2206.10545.pdf.
[64] Daniel Smullen, Yuanyuan Feng, Shikun Aerin Zhang, and Norman Sadeh. 2020. 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 (2020), 195–215. https://doi.org/10.2478/popets-2020-0011
[65] Daniel Smullen, Yaxing Yao, Yuanyuan Feng, Norman Sadeh, Arthur Edelstein, and Rebecca Weiss. 2021. Managing Potentially Intrusive Practices in the Browser: A User-Centered Perspective. *Proceedings on Privacy Enhancing Technologies* 2021, 4 (2021), 500–527. https://doi.org/10.2478/popets-2021-0082
[66] Statcounter. 2022. Desktop Browser Market Share United States of America, July 2022. https://gs.statcounter.com/browser-market-share/desktop/united-states-of-america. Accessed: June, 2022.
[67] Statcounter. 2022. Desktop Operating System Market Share United States Of America, July 2022. https://gs.statcounter.com/os-market-share/desktop/united-states-of-america. Accessed: June, 2022.
[68] State of California Department of Justice. 2020. California Consumer Privacy Act (CCPA) Final Statement of Reasons, Appendix E. https://oag.ca.gov/sites/all/files/agweb/pdfs/privacy/ccpa-fsor-appendix-e.pdf. accessed: October 22, 2023.
[69] State of California Department of Justice. 2022. Attorney General Bonta Announces Settlement with Sephora as Part of Ongoing Enforcement of California Consumer Privacy Act. https://oag.ca.gov/news/press-releases/attorney-general-bonta-announces-settlement-sephora-part-ongoing-enforcement. Accessed: October 22, 2023.
[70] State of Connecticut General Assembly. 2022. Connecticut Data Privacy Act. https://www.cga.ct.gov/2022/amd/S/pdf/2022SB-00006-R00SA-AMD.pdf. Accessed: October 22, 2023.
[71] Richard H. Thaler and Cass R. Sunstei. 2009. Nudge: Improving Decisions About Health, Wealth, and Happiness.
[72] Joseph Turow, Yphtach Lelkes, Nora A. Draper, and Ari Ezra Waldman. 2023. Americans Can't Consent to Companies' Use of their data. https://www.asc.upenn.edu/sites/default/files/2023-02/Americans_Can%27t_Consent.pdf
[73] Christine Utz, Martin Degeling, Sascha Fahl, Florian Schaub, and Thorsten Holz. 2019. (Un)Informed Consent: Studying GDPR Consent Notices in the Field. In *CCS* (London, United Kingdom) *(CCS '19)*. Association for Computing Machinery, New York, NY, USA, 973–990. https://doi.org/10.1145/3319535.3354212
[74] Christine Utz, Martin Degeling, Sascha Fahl, Florian Schaub, and Thorsten Holz. 2019. (Un)informed Consent: Studying GDPR Consent Notices in the Field. In *CCS*. Association for Computing Machinery, New York, NY, USA, 973–990.
[75] W3C. 2019. Tracking Preference Expression (DNT). https://www.w3.org/TR/tracking-dnt/. Accessed: October 22, 2023.
[76] Jason Watson, Heather Richter Lipford, and Andrew Besmer. 2015. Mapping User Preference to Privacy Default Settings. *ACM Trans. Comput.-Hum. Interact.* 22, 6, Article 32 (Nov. 2015), 20 pages. https://doi.org/10.1145/2811257
[77] Yaxing Yao, Davide Lo Re, and Yang Wang. 2017. Folk Models of Online Behavioral Advertising. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing* (Portland, Oregon, USA) *(CSCW '17)*. Association for Computing Machinery, New York, NY, USA, 1957–1969. https://doi.org/10.1145/2998181.2998316
[78] Zhonghao Yu, Sam Macbeth, Konark Modi, and Josep M. Pujol. 2016. Tracking the Trackers. In *Proceedings of the 25th International Conference on World Wide Web* (Montréal, Québec, Canada) *(WWW '16)*. International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, CHE, 121–132. https://doi.org/10.1145/2872427.2883028
[79] Shikun Zhang, Yuanyuan Feng, Yaxing Yao, Lorrie Faith Cranor, and Norman Sadeh. 2022. How Usable Are iOS App Privacy Labels?. In *PoPETS 2022*. PoPETS 2022, Sydney, Australia and Virtual, 204–228.
[80] Sebastian Zimmeck and Kuba Alicki. 2020. Standardizing and Implementing Do Not Sell. https://dl.acm.org/doi/abs/10.1145/3411497.3420224. In *WPES*. ACM, Virtual Event, USA, 15–20.
[81] Sebastian Zimmeck, Oliver Wang, Kuba Alicki, Jocelyn Wang, and Sophie Eng. 2023. Usability and Enforceability of Global Privacy Control. In *23rd Privacy Enhancing Technologies Symposium (PETS 2023)*, Vol. 2. Sciendo, Lausanne, Switzerland, 265–281. https://doi.org/10.56553/popets-2023-0052

A APPENDIX

A.1 Categories of Data Collected from Study Participants

- General

- Prolific ID
- Browser (e.g., Google Chrome)
- Whether HTTP cookies are enabled (true or false)
- Whether Do Not Track is enabled (true or false)
- Whether JavaScript is enabled (true or false)
- Operating system language (e.g., en-US)
- Geographic location (e.g., latitude/longitude coordinates)
- Whether Local Storage is enabled in the browser (true or false)
- Whether Session Storage is enabled (true or false)
- Operating system (e.g., macOS)
- Browser plugins (e.g., Chrome PDF Plugin)
- Browser's rendering engine (e.g., WebKit)
- Time zone
- UI scheme
- User agent of the browser

- Browser History
  - Website URLs visited (e.g., https://www.cnbc.com/finance/)
  - Timestamp of when a site was visited
  - Selected GPC status for the current site (true or false)
  - Whether GPC is globally enabled for all sites visited
  - Referer
  - Tab IDs of the different tabs in the browser
  - Whether ads on websites are clicked
  - The website URL that the browser redirects to when an ad is clicked
  - The website URLs of the ad networks integrated in the sites visited
- Third Party Requests (only the first 50 are registered alongside a summary)
  - The domain and third party associated with the request
  - The categories linked to the third party per Disconnect's Tracker Protection lists
  - The site the request was sent to
  - A sum of the total number of requests from each third party
  - Timestamp
- Site Interactions (records per-website changes to GPC settings)
  - The website URLs whose settings are being changed
  - The origin of the change (either the settings page or a GPC banner for schemes *SB-Base*, *S0-Snooze*, *S1-Apply-all*, and *S2-Snooze+Apply-all*)
  - The previous and new settings
  - Whether or not the setting was applied to all future sites (for schemes *S1-Apply-all* and *S2-Snooze+Apply-all*)
  - Timestamp
- Privacy Configuration Interactions (non-website-specific privacy choices made for schemes *S3-Profile*, *S4-Website*, *S5-Learn*, *S6-Universal*, and *S7-Data*)
  - The type of setting (scheme-dependent)
  - The previous and new settings
  - Timestamp
- Snooze Interactions (applicable to schemes *S0-Snooze* and *S2-Snooze+Apply-all*)
  - The URL of the site that was snoozed
  - Timestamp
- Ad Interactions
  - The source of the ad
  - The URL navigated to after the interaction
  - The reason for the event being flagged as an ad interaction
  - Timestamp

A.2 Exit Survey Questionnaire

- **Q1** Timestamp [Recorded automatically when participant submits survey]
- **Q2** Are you comfortable visiting a website that collects each of the following data from you to show you relevant ads? Assume that the site will *not* share your data with any other company. Select all that applies or "None" if you prefer no data collection. [Checkboxes; answer required]
  ∘ Phone number ∘ Email address ∘ GPS location (within 20 feet of your actual location) ∘ Zip code ∘ Browsing history ∘ Age ∘ Ethnicity/Race ∘ Income ∘ Gender ∘ None
- **Q3** Are you comfortable visiting a website that shares each of your following data with advertisers? Assume that some of the advertisers will share your data with other advertisers and data brokers. Select all that applies or "None" if you prefer no data sharing. [Checkboxes; answer required]
  ∘ Phone number ∘ Email address ∘ GPS location (within 20 feet of your actual location) ∘ Zip code ∘ Browsing history ∘ Age ∘ Ethnicity/Race ∘ Income ∘ Gender ∘ None
- **Q4** Please select the extent to which you agree with the statement? "When I browse the web, I feel that I have control over who is receiving my data." [Multiple choice; answer required]
  Strongly Disagree ∘ 1 ∘ 2 ∘ 3 ∘ 4 ∘ 5 Strongly Agree
- **Q5** Do you currently have paid subscriptions for news content on the Internet? (e.g., NYTimes, Reddit Premium, ...) [Multiple choice; answer required]
  ∘ Yes ∘ No
- **Q6** Which paid news content subscriptions do you have? Please enter all subscriptions separated by a comma. [Long answer text; answer required if participants have paid subscriptions]
- **Q7** What are the reasons for why you do not have paid subscriptions for news content? [Long answer text; answer required if participants do not have paid subscriptions]
- **Q8** Please select the extent to which you agree with the statement? "For news content that is of interest to me, I am happy to pay a subscription fee of $5 per month." [Multiple choice; answer required]
  Strongly Disagree ∘ 1 ∘ 2 ∘ 3 ∘ 4 ∘ 5 Strongly Agree
- **Q9** Are you using any social media services? (e.g., Facebook, Instagram, TikTok, Snapchat, ...) [Multiple choice; answer required]
  ∘ Yes ∘ No

- **Q10** For any social media service you are using, have you ever changed its privacy settings? [Multiple choice; answer required if participants use social media]
  - ○ Yes ○ No
- **Q11** What did you change in the social media privacy settings? [Long answer text; answer required if participants changed a privacy setting]
- **Q12** What are the reasons for why you did not change any social media privacy settings? [Long answer text; answer required if participants did not change a privacy setting]
- **Q13** If you recall, in our browser extension you were shown the following explanation on what Global Privacy Control (GPC) does [shown above]. Based on the explanation, what do you think is true? [Multiple choice; answer required; the correct answer is shown in bold]
  - ○ Websites respecting GPC would be prohibited from collecting data from you
  - ○ **Websites respecting GPC would be prohibited from giving advertisers your data**
  - ○ Websites respecting GPC would be prohibited from showing you advertising
  - ○ GPC signals must be respected for you and anyone else in the world
  - ○ None of the above
- **Q14** Please explain in your own words what GPC does (please assume that it applies in your state of residence). [Long answer text; answer required]
- **Q15** If you could change one thing about GPC, what would it be? Why? [Long answer text; answer required]
- **Q16** Is there anything you find exciting about GPC? What is it? Why is it exciting? [Long answer text; answer required]
- **Q17** Please select the extent to which you agree with the statement: "I was able to make my opt out choices the way I wanted." [Multiple choice; answer required]
  Strongly Disagree ○ 1 ○ 2 ○ 3 ○ 4 ○ 5 Strongly Agree
- **Q18** Please select the extent to which you agree with the statement: "It took a lot of effort to opt out." [Multiple choice; answer required]
  Strongly Disagree ○ 1 ○ 2 ○ 3 ○ 4 ○ 5 Strongly Agree
- **Q19** Please select the extent to which you agree with the statement: "I was able to find all the settings I was looking for." [Multiple choice; answer required]
  Strongly Disagree ○ 1 ○ 2 ○ 3 ○ 4 ○ 5 Strongly Agree
- **Q20** Please select the extent to which you agree with the statement: "I felt confident in my understanding of the settings I used." [Multiple choice; answer required]
  Strongly Disagree ○ 1 ○ 2 ○ 3 ○ 4 ○ 5 Strongly Agree
- **Q21** Please select the extent to which you agree with the statement: "I wished I had the support of a technical person to be able to opt out." [Multiple choice; answer required]
  Strongly Disagree ○ 1 ○ 2 ○ 3 ○ 4 ○ 5 Strongly Agree
- **Q22** Please select the extent to which you agree with the statement: "I felt confident that my opt out choices were honored." [Multiple choice; answer required]
  Strongly Disagree ○ 1 ○ 2 ○ 3 ○ 4 ○ 5 Strongly Agree
- **Q23** Please select the extent to which you agree with the statement: "My normal use of the browser was disrupted by the opt out interface." [Multiple choice; answer required]
  Strongly Disagree ○ 1 ○ 2 ○ 3 ○ 4 ○ 5 Strongly Agree
- **Q24** Please select the extent to which you agree with the statement: "The number of options for making my opt out choices was insufficient." [Multiple choice; answer required]
  Strongly Disagree ○ 1 ○ 2 ○ 3 ○ 4 ○ 5 Strongly Agree
- **Q25** Please select the extent to which you agree with the statement: "The granularity of options for making my opt out choices was sufficient." [Multiple choice; answer required]
  Strongly Disagree ○ 1 ○ 2 ○ 3 ○ 4 ○ 5 Strongly Agree
- **Q26** Please select the extent to which you agree with the statement: "I was able to correct errors with ease." [Multiple choice; answer required]
  - ○ Strongly disagree ○ Disagree ○ Neutral ○ Agree ○ Strongly agree ○ No error occurred
- **Q27** Please select the extent to which you agree with the statement: "I found it difficult to change my opt out choices." [Multiple choice; answer required]
  - ○ Strongly disagree ○ Disagree ○ Neutral ○ Agree ○ Strongly agree ○ No change was necessary
- **Q28** Please select the extent to which you agree with the statement: "I felt pushed towards opting out." [Multiple choice; answer required]
  Strongly Disagree ○ 1 ○ 2 ○ 3 ○ 4 ○ 5 Strongly Agree
- **Q29** Please select the extent to which you agree with the statement: "I would have liked to further personalize my opt out choices." [Multiple choice; answer required]
  Strongly Disagree ○ 1 ○ 2 ○ 3 ○ 4 ○ 5 Strongly Agree
- **Q30** It is important that you pay attention to this study. Please select "Neutral". [Multiple choice; answer required]
  - ○ Strongly agree ○ Agree ○ **Neutral** ○ Disagree ○ Strongly disagree
- **Q31** Do you have improvement suggestions for GPC or the opt out experience you had? If so, please let us know. [Long answer text; answer required]
- Please enter your Prolific ID. [Short answer text; answer required]

A.3 Scheme User Interfaces

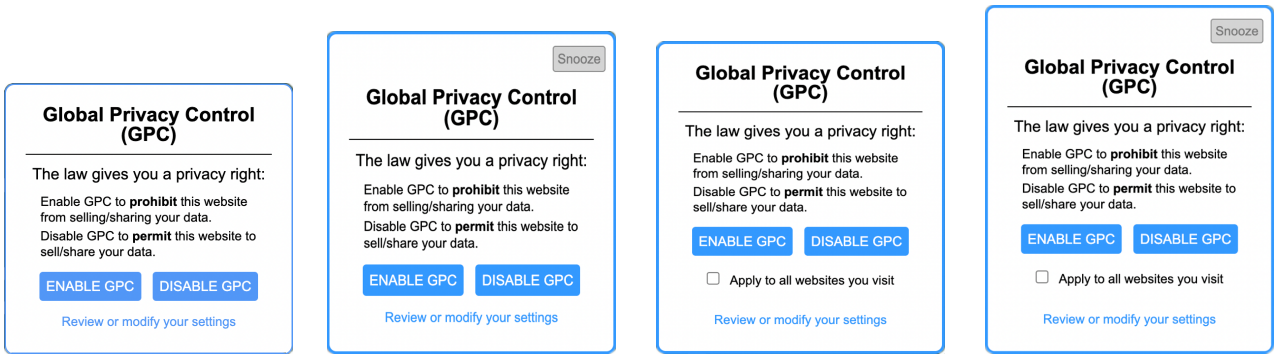

Figure 13: On the banner schemes, i.e., *SB-Base*, *S0-Snooze*, *S1-Apply-all*, and *S2-Snooze+Apply-all* (from left to right), participants are prompted for a GPC privacy choice via a banner on each new site they visit. *S0-Snooze* and *S2-Snooze+Apply-all* include a snooze button that prevents the banner from popping up on new sites for 12 hours at a time. *S1-Apply-all* and *S2-Snooze+Apply-all* include an apply-all feature that will apply a participant's choice to all future sites if they so wish. Participants can select to which sites they want to send GPC signals via a domain list on the settings page.

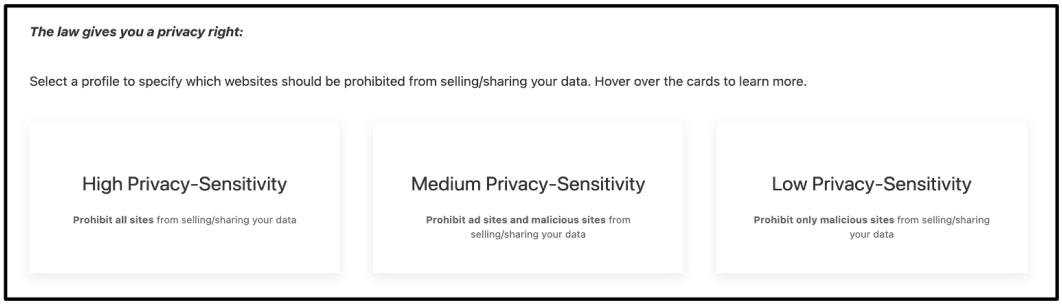

Figure 14: Upon installing our extension with the *S3-Profile* scheme, participants are prompted to choose a privacy profile. Their choice will then determine which sites will receive GPC signals. They may change this preference on the settings page. The settings page also contains a domain list.

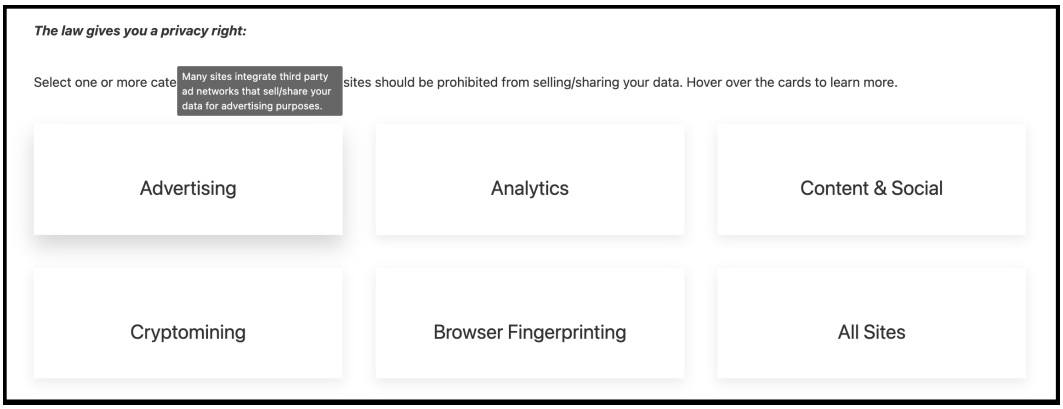

Figure 15: Upon installing our extension with the *S4-Website* scheme, participants are prompted to select the website categories that they would like to opt out from. Each category has a mouse-over tooltip with a more detailed description. Participants may change their categories on the settings page. The settings page also contains a domain list. We implemented the scheme with categories of third party sites from which people would be opted out. However, it would also be possible to implement it with first party site categories such as news, music, etc.

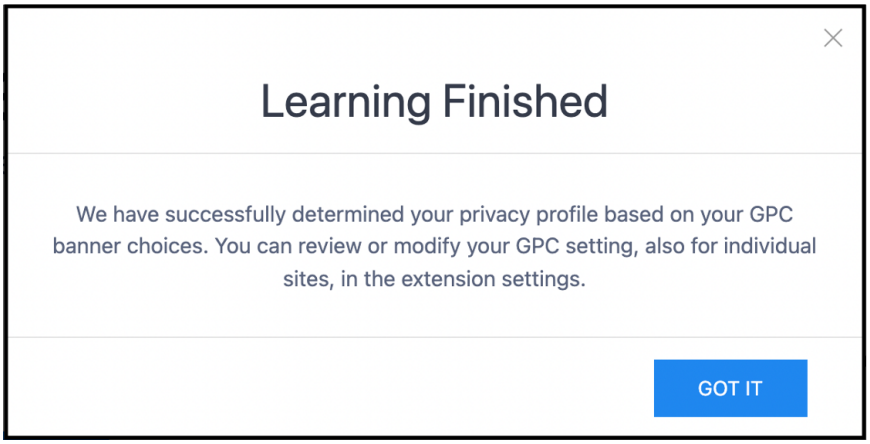

Figure 16: Upon installing our extension with the *S5-Learn* scheme, participants are prompted for their GPC choices via banners on the first 10 sites they visit. Their choices are then used to select a privacy profile that suits them best. The profiles are the same as for *S3-Profile*. After the learning period, participants are redirected to the settings page and shown which privacy profile they were assigned. They may change their privacy profile there. The settings page also contains a domain list.

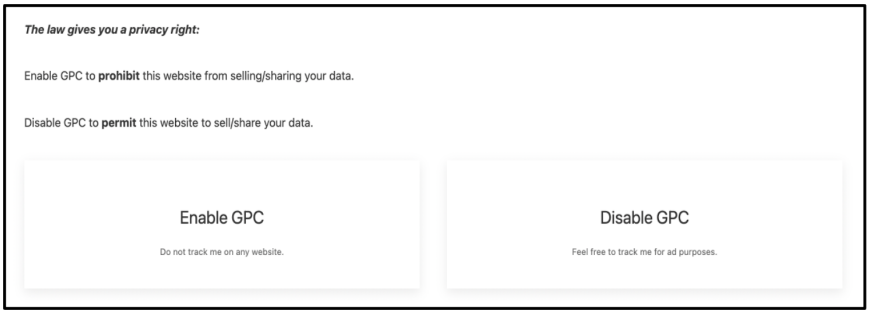

Figure 17: Upon installing our extension with the *S6-Universal* scheme, participants are prompted as to whether they would like to send GPC signals to all sites they visit or not. They may change this preference on the settings page. There is also a domain list on the settings page that participants may utilize if they so choose.

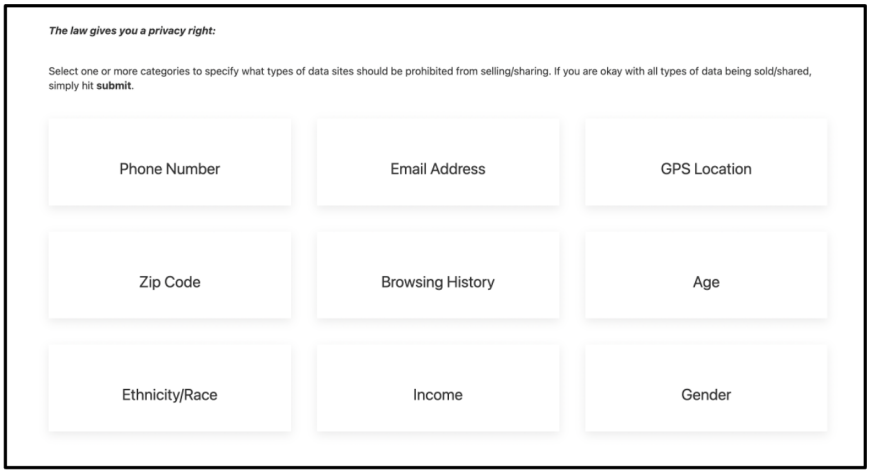

Figure 18: Upon installing our extension with the *S7-Data* scheme, participants are prompted to select the categories of data that they would not like to be shared with or sold to advertisers. Participants can adjust their settings on the settings page.

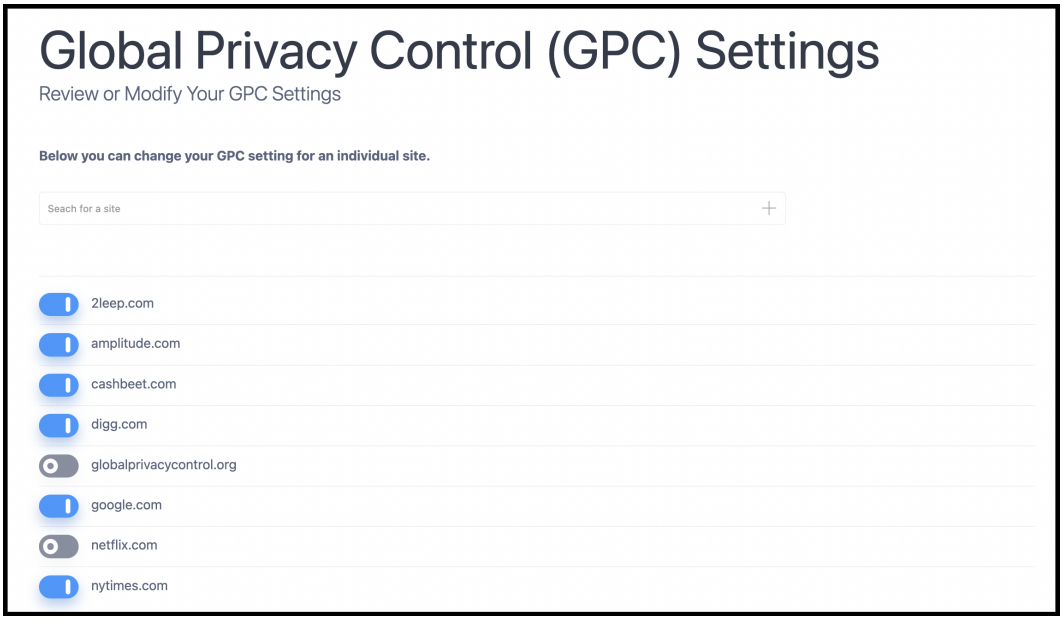

Figure 19: The domain list was present in all schemes except *S7-Data*. It allowed participants the option to make specific GPC privacy choices for each of their domains, i.e., the first party sites they intentionally visited while the extension was running.

A.4 Browsing Statistics

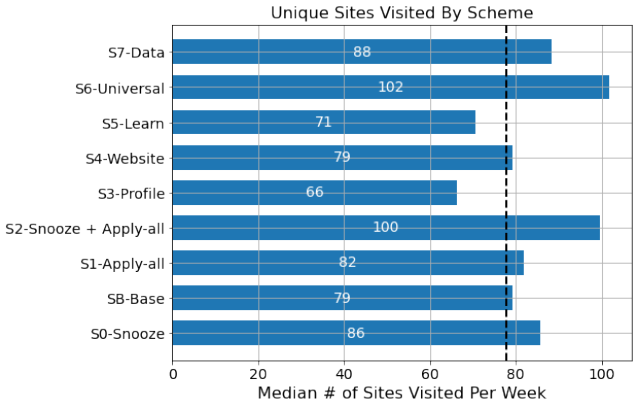

Figure 20: The median number of unique sites participants' visited per week by scheme. The dotted black line denotes the total median across schemes (78 sites).

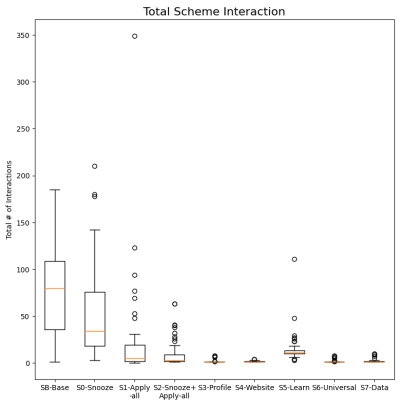

Figure 21: Participants' rates of interaction for their assigned schemes showing median, minimum, maximum, and the interquartile range for each. Interactions include site, privacy configuration, and snooze interactions (§4.1). Banner schemes naturally required far more interactions than category schemes.

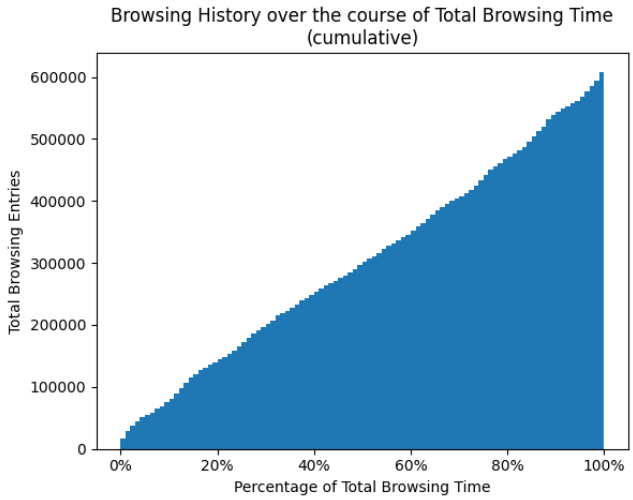

Figure 22: Distribution of participant browsing history entries across all schemes relative to total browsing time, cumulative. For example, the graph shows that after 20% of all participants' browsing time (around 1.4 days), around 140,000/600,000 (23%) total website visits had been recorded. The figure indicates a roughly linear relationship, with website visits accumulating at an approximately constant rate.

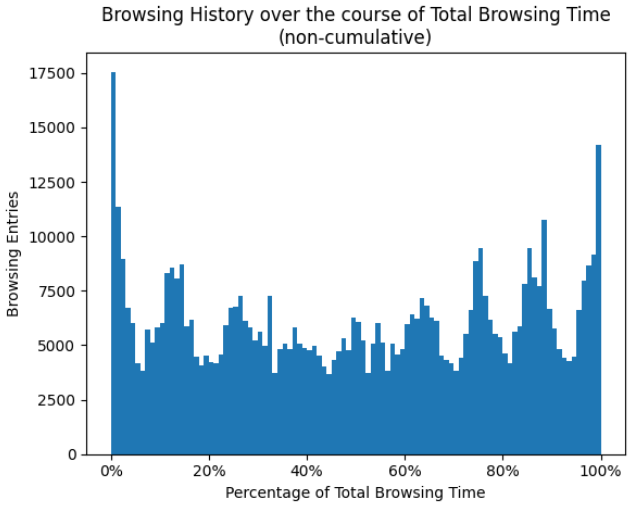

Figure 23: Distribution of participant browsing history entries across all schemes relative to total browsing time, non-cumulative. Of note are the spikes in activity around the start and end of the study periods, but, beyond that, activity is mostly constant save for various dips, presumably around nighttime.

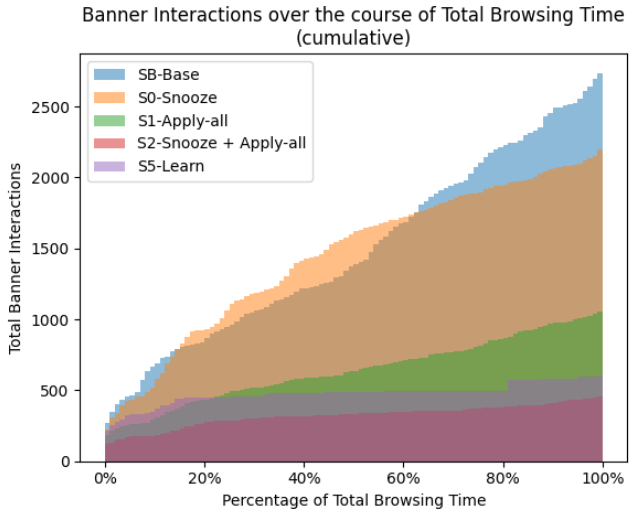

Figure 24: Distribution of participant banner interactions across banner schemes, including *S5-Learn* as a mixed banner-category scheme, relative to the total browsing time for each participant (around 7 days), cumulative. Curves that transition from steep to flat slopes indicate that most of participants' banner interactions for that scheme were concentrated in the start of the browsing period. Schemes with generalizability features, such as *S2-Snooze + Apply-all* and *S5-Learn* evidently required fewer banner interactions throughout the entire week compared to, for example, *SB-Base*.

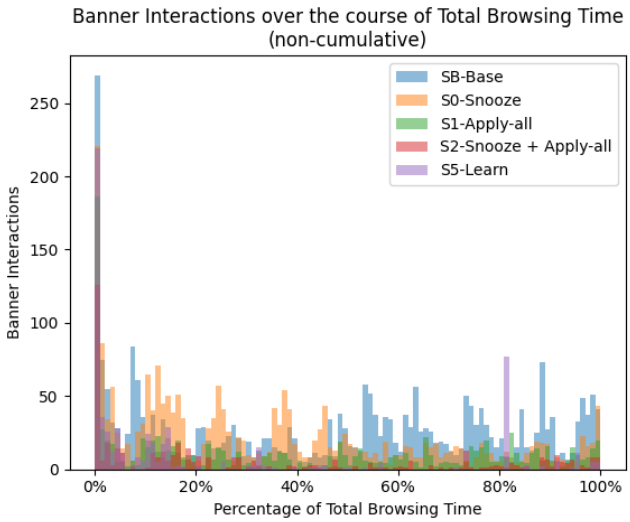

Figure 25: Distribution of participant banner interactions across banner schemes relative to the total browsing time for each participant, non-cumulative. More prominent spikes at the beginning of the browsing period and fewer afterwards indicate more interactions being made early on.

A.5 Selected Results for Significance Tests

Scheme	SB-Base		S0-Snooze		S1-Apply-all		S2-Snooze+Apply-all		S3-Profile		S4-Website		S5-Learn		S6-Universal		S7-Data
SB-Base																	
S0-Snooze	0.71	(0.63)															
S1-Apply-all	0.68	(0.53)	0.43	(0.24)													
S2-Snooze+Apply-all	**0.023**	(0.0038)	**0.0056**	(<0.001)	0.073	(0.018)											
S3-Profile	0.14	(0.047)	**0.043**	(0.0083)	0.32	(0.16)	0.50	(0.33)									
S4-Website	0.079	(0.024)	**0.023**	(0.0033)	0.19	(0.092)	0.63	(0.48)	0.83	(0.79)							
S5-Learn	0.074	(0.021)	**0.023**	(0.0031)	0.18	(0.077)	0.71	(0.61)	0.73	(0.67)	0.89	(0.86)					
S6-Universal	**0.023**	(0.0025)	**0.0051**	(<0.001)	0.061	(0.014)	0.97	(0.97)	0.50	(0.32)	0.63	(0.47)	0.71	(0.62)			
S7-Data	0.44	(0.26)	0.19	(0.091)	0.71	(0.60)	0.18	(0.075)	0.57	(0.40)	0.44	(0.27)	0.42	(0.22)	0.18	(0.065)	

Table 3: Results of pairwise Dunn test comparisons for Q17 ("Please select the extent to which you agree with the statement: 'I was able to make my opt out choices the way I wanted.' ") provided as p-values. P-values corrected via the Benjamini-Hochberg method are presented in black text, and those less than 0.050 are bolded (indicating rejection of the null hypothesis). Red parentheticals display uncorrected p-values. Comparisons of note for the paper's contents are further accentuated with grey backgrounds. Figures are rounded to two significant digits as applicable. Various significant comparisons, e.g., between *SB-Base* and *S3-Profile*, have been rendered insignificant by the correction. These comparisons are still noteworthy, but less so than the ones remaining significant.

Scheme	SB-Base		S0-Snooze		S1-Apply-all		S2-Snooze+Apply-all		S3-Profile		S4-Website		S5-Learn		S6-Universal		S7-Data
SB-Base																	
S0-Snooze	0.53	(0.46)															
S1-Apply-all	0.26	(0.19)	0.61	(0.54)													
S2-Snooze+Apply-all	**0.0019**	(<0.001)	**0.0095**	(0.0037)	**0.050***	(0.025)											
S3-Profile	**<0.001**	(<0.001)	**<0.001**	(<0.001)	**0.0040**	(0.0012)	0.38	(0.31)									
S4-Website	**<0.001**	(<0.001)	**<0.001**	(<0.001)	**<0.001**	(<0.001)	0.08	(0.043)	0.38	(0.32)							
S5-Learn	0.13	(0.079)	0.35	(0.27)	0.64	(0.62)	0.14	(0.090)	**0.017**	(0.0078)	**0.0012**	(<0.001)					
S6-Universal	**0.0045**	(0.0015)	**0.022**	(0.010)	0.11	(0.062)	0.63	(0.60)	0.16	(0.11)	**0.017**	(0.0071)	0.26	(0.20)			
S7-Data	**<0.001**	(<0.001)	**<0.001**	(<0.001)	**0.0019**	(<0.001)	0.25	(0.17)	0.69	(0.69)	0.63	(0.58)	**0.0089**	(0.0032)	0.09	(0.050)	

Table 4: Results of pairwise Dunn test comparisons for Q23 ("Please select the extent to which you agree with the statement: 'My normal use of the browser was disrupted by the opt out interface.' ") provided as p-values. P-values corrected via the Benjamini-Hochberg method are presented in black text, and those less than 0.050 are bolded (indicating rejection of the null hypothesis). Red parentheticals display uncorrected p-values. Comparisons of note for the paper's contents are further accentuated with grey backgrounds. Figures are rounded to two significant digits as applicable.

* **The more precise p-value for the comparison between *S1-Apply-all* and *S2-Snooze+Apply-all* is less than 0.050. It is shown as 0.050 due to rounding.**