

Data Sharing and Website Competition: The Role of Dark Patterns

Chiara Farronato, Andrey Fradkin, Tesary Lin

September 12, 2024

Preliminary and Incomplete

Regulations like the GDPR require firms to obtain consumer consent before using data. In response, some firms employ “dark patterns” — interface designs that nudge consumers to share data. We study the causal effects of these designs and how they vary across individuals and firms. To do so, we run a field experiment in which users download a browser extension that randomizes cookie consent interface designs as users browse the Internet. We find that consumers accept all cookies more than half of the time in the absence of dark patterns. Hiding consent options behind an additional click is the most effective dark pattern, while designs that only manipulate visual elements have smaller effects. Larger and better-known firms have moderately higher consent rates than other firms, giving them a slight competitive advantage. However, the effects of dark patterns do not vary systematically across site popularity. We find no evidence that frequent pop-ups increase choice fatigue.

All authors contributed equally to the paper. Chiara Farronato: Harvard Business School, NBER, and CEPR. Andrey Fradkin: Boston University and the MIT Initiative on the Digital Economy. Tesary Lin: Boston University. We thank Chloe Song for excellent research assistance. This work received generous support from the Internet Society Foundation and Boston University’s Digital Business Institute.

1. Introduction

Data is critical to the functioning of modern businesses, whether as an input into pricing or search ranking algorithms, or for targeted advertising. At the same time, the collection and use of data is a threat to consumer privacy. Over the past decade, regulatory and societal pressures have led companies to offer consumers choices about the collection and use of their data. At the same time, regulators across the continents have expressed concerns about how current data collection practices can negatively affect competition via data-enabled network effects.¹

To incentivize consumers to share data, businesses often structure these choices to make it hard for users to select options that limit data sharing, a phenomenon called *dark patterns* in public discourse and policy discussions.² Dark patterns are widespread. One example involves presenting users with only two options: “accept all” or “customize settings.” The option to reject non-essential cookies is hidden under the option to customize them. Firms frequently deploy this design, as it nudges users to share their data more readily than in settings where the reject option is easily accessible. Design patterns like this example are considered by regulators to be a serious problem that could pose consumer harm, and the European Union’s Digital Services Act explicitly bans dark patterns.³

In this project, we conduct a systematic evaluation of consent banners, and the impact of different types of dark patterns on consumer privacy choices. We also explore how dark patterns may exacerbate or ease data advantages of large companies. Lastly, we design our study to measure choice fatigue, meaning disengagement with consent banners by directly closing the banner without making a choice. We find that dark patterns effectively increase consent rates, but their effects do not systematically vary across websites. Although we do not find a causal effect of more frequent banners on consumer choices, we detect an increased propensity to disengage from cookie tracking choices over time.

We pursue our research questions with a field experiment conducted using Cookie Manager, a browser extension that changes the cookie consent interface seen by users as they browse the Internet. Cookie Manager is based on the Webmunk extension framework (Farronato, Fradkin, and Karr 2024). Cookie Manager presents six cookie consent variations from which we can identify the role of three types of dark patterns: removing choices, re-ordering choices to prioritize options with more data sharing, and empha-

¹See, for example, The Digital Markets Act that requires “gatekeeper” platforms to share data with smaller players when upon request: https://commission.europa.eu/strategy-and-policy/priorities-2019-2024/europe-fit-digital-age/digital-markets-act-ensuring-fair-and-open-digital-markets_en, and the FTC’s report to OECD on how data privacy and competition interacts: [https://one.oecd.org/document/DAF/COMP/WD\(2024\)29/en/pdf](https://one.oecd.org/document/DAF/COMP/WD(2024)29/en/pdf)

²<https://www.deceptive.design/>

³https://ec.europa.eu/commission/presscorner/detail/en/QANDA_20_2348.

sizing the choice preferred by the firm with a brighter color. Our experimental design presents interfaces at random across users and web domains, which allows us to identify the causal effects of these interfaces. Cookie manager attempts to enforce user choices to accept or reject whenever possible, so that user choices are realistic.

To implement our study, we recruit users from Prolific and ask them to install the web browser extension. After installation, the study consists of two parts. In the first part (*survey browsing*), we prompt users to visit a series of pre-specified websites in order to evaluate participants' preferences for privacy. Upon visiting a website for the first time, a cookie consent banner appears while the website content fades to gray in the background. The cookie consent form is randomly chosen among the interfaces we designed. Users must select an option or click 'x' to escape the form and continue to the website content. Cookie manager does nothing to cookie settings when users hit 'x', but 60% of users believe that clicking 'x' rejects all cookies.

In the second part (*organic browsing*), we observe organic browsing in the week following enrollment in the study. The combination of survey-based and organic website visits allows us to explore how different users make choices for the same website, and whether these choices are a good representation of the choices that they would make on websites that they organically choose to visit.

In the interface without dark patterns (*neutral interface*), the percentage of participants choosing "accept all" is 62% in the survey browsing and 55% during organic browsing. Rather than "always-accept" or "always reject," many participants make website-specific choices. When faced with the neutral interface, 53.5% of consumers vary their choices across sites, while the rest make the same choice regardless of website (and in that case, they mostly always accept cookies). Those who do vary their selections are more likely to share cookies to websites that are popular or familiar to them. Across our six treatments, 91% of consumers change their privacy choices across websites, showing that choice architecture results in additional variation in choices.

We find that deliberate obstruction has the strongest influence on privacy choices, while dark patterns that feature pure visual manipulations have weaker effects. In particular, hiding the "reject all" button from the main user interface decreases the probability of rejecting cookies by 17.1% in the survey visits and 9.4% in organic browsing. In comparison, re-ordering options so that the "accept all" is first increases acceptance by 2-3.3%, while graying out options other than "accept all" increases the acceptance probability by less than 2%. These findings suggest that dark patterns that increase frictions by adding clicks are more effective than designs that only involve changes in visual cues. These findings are consistent with websites' revealed preferences: as is shown by Utz et al. (2019), deliberate construction is present in 78.5% of the cookie banners they collected.

Perhaps surprisingly, the effects of dark patterns do not vary substantially with some important website characteristics. Consumers are influenced by dark patterns by the same degree, regardless of website popularity or their familiarity with websites. We do find some heterogeneous effects of dark patterns for websites that the participants state they normally visit. The lack of heterogeneity by domain rank alleviates concerns that dark patterns exacerbate entry barriers, or amplify data-enabled network effects (Hagiu and Wright 2023) that create incumbent advantages in the data economy. That said, we do find that regardless of the user interface, consumers are more likely to consent to data sharing on websites that they normally visit.

Lastly, we consider choice fatigue, which has become a concern as regulation has increased the frequency of consent banners across the Internet. In our setting, choice fatigue may be defined in two ways. First, we explore whether consumers decrease their probability of making an active choice as the frequency of choices increases. During the week of organic browsing, we randomize users into two groups: one receiving a consent banner every 10 minutes (more frequent pop-ups), and the other every 60 minutes. This experimental variation resulted in 76% more pop-ups in the 10-minute group. Our results suggest that there is no difference in users' choices across the two groups. Therefore, an exposure-based choice fatigue story does not hold. Second, we investigate how choices change simply as a function of time in the study. We find that users are more likely to close consent banners over time, suggesting diminishing attention over time, although unrelated to the frequency of pop-ups.

Our work relates to a broader literature on the consequences of behavioral consumers for competition (Huck and Zhou 2011; Spiegler 2014; Ho, Hogan, and Scott Morton 2017; Decarolis, Li, and Paternolillo 2023). This literature has explored how factors such as switching costs and obfuscation may hinder competition. We find that behavioral interventions are effective in allowing firms to collect more data. That said, the effectiveness of these interventions does not vary across websites, and thus policies that affect firms' abilities to use dark patterns will not have competitive effects through the volume of data collected. Nonetheless, since behavioral data may be most useful for targeted advertising by small firms, the removal of dark patterns may still have harmful downstream consequences on competition (Aridor et al. 2024). We will return to this point in the concluding section.

We contribute to the literature on the effects of dark patterns. Existing empirical work on dark patterns in privacy settings has mostly focused on describing their deployment (Mathur et al. 2019; Di Geronimo et al. 2020; Warberg et al. 2023). Efforts to measure how dark patterns affect privacy choices have mostly relied on lab or synthetic settings (Acquisti, John, and Loewenstein 2013; Utz et al. 2019; Luguri and Strahilevitz 2021; Habib et al. 2022; Lin and Strulov-Shlain 2023; Bielova et al. 2024; Baviskar et al. 2024), with

the exception of D'Assergio et al. (2022), who examines the effectiveness of persuasive language in re-permission emails in encouraging opt-in. Our approach has an advantage in realism over hypothetical or lab settings, since users in our study make consequential choices while going about their normal browsing activities. Unlike prior studies, we are able to study privacy choices across a wide selection of websites, which is crucial for understanding how dark patterns could potentially affect websites' competition and their access to consumer data.

Lastly, our work relates to the literature on the economics of privacy and on measuring privacy choices (Lin 2022; Collis et al. 2021; Tomaino, Wertenbroch, and Walters 2023; Tang 2023; Acquisti, John, and Loewenstein 2013). We note that choices in our setting do not simply reflect a user's value for privacy, but also the utility that some user may get from cookie tracking. That said, we are the first to measure privacy choices in the field over an extended period of time. This temporal dimension allows us to measure the sustained effects of choice architecture and whether they change as a function of learning or fatigue.

2. Institutional Background

The phrase "dark pattern" was coined by a computer scientist, Harry Brignull, to refer to design patterns that "deceive and manipulate users into taking actions they did not intend."⁴ Although the usage of dark patterns is not restricted to data exchange settings, companies routinely use them whenever consent is required for data collection and processing. For example, Utz et al. (2019) crawled major EU websites post-GDPR⁵ and found that 57.4% of these websites use dark patterns in their consent banners. Similarly, Nouwens et al. (2020) focused on the top 10,000 UK websites, and documented dark pattern deployment in over 80% of them.

Since these manipulative patterns could lead consumers to make choices that they would not have otherwise chosen, policymakers worry that the widespread deployment of dark patterns can cause consumer harm. Regulatory and legal intervention soon followed. For example, the Federal Trade Commission has fined large companies such as Epic Games and Amazon for their user interface designs that induce accidental purchases and obstruct cancellation to subscriptions.⁶ US States such as California, Colorado, and Connecticut have enacted privacy regulations that explicitly ban companies from using dark patterns

⁴<https://hallofshame.design/about/>

⁵GDPR, or General Data Protection Regulation, is a European Union regulation passed in 2018 that requires consumer explicit consent to data collection.

⁶<https://www.ftc.gov/news-events/news/press-releases/2023/03/ftc-finalizes-order-requiring-fortnite-maker-epic-games-pay-245-million-tricking-users-making>, <https://www.ftc.gov/news-events/news/press-releases/2023/03/ftc-finalizes-order-requiring-fortnite-maker-epic-games-pay-245-million-tricking-users-making>

to increase data collection.⁷ In the European Union, their landmark privacy regulation (General Data Protection Regulation, or GDPR) requires consent for data collection to be “freely given, specific, informed and unambiguous,” and Recital 32 specifically requires that consent should be granular to the purposes of data processing, and that default settings and inactions do not constitute consent.⁸ However, Bielova, Santos, and Gray (2024) argues that GDPR still leaves ample ambiguity on whether other dark patterns are allowed for encouraging consent. Regulators are also trying to prohibit the usage of dark patterns beyond the privacy realm, as is shown by the introduction of the bipartisan bill—Deceptive Experiences To Online Users Reduction (DETOUR) Act—in 2023.⁹

Prior research has documented a variety of dark patterns that companies deploy to advance data collection (Habib et al. 2022; Bielova, Santos, and Gray 2024). These different practices can be categorized into three main groups. The first group includes information or persuasion-based tactics. These strategies involve describing data sharing as more appealing than it may otherwise appear. Examples of this design include pre-prompts that apps can show users before Apple’s app tracking transparency (ATT) prompt,¹⁰ and wording that associates the non-sharing option with negative emotions, commonly known as the “confirm shaming” technique.¹¹ Recent work has failed to find evidence of the efficacy of these techniques. For example, Bielova, Santos, and Gray (2024) show that changes in consent banner texts do not significantly change privacy choices, presumably because consumers do not pay attention to these texts when interacting with the banners. Similarly, D’Assergio et al. (2022) show that adding persuasive language (other than giving incentives) in emails that request data collection opt-in do not improve consent rate.

The second type of pattern consists of obstruction tactics, or designs that increase frictions associated with consumer choices undesirable to the firm. Two most prominent examples involve setting defaults to “share all,” and designing what are known as “unequal path.” The latter strategy refers to designs that include “share all” on a main screen while the “reject all” option is hidden behind additional clicks (for example, under “settings”). The two designs are the two most popular dark patterns among websites. Indeed, Habib et al. (2022) show that unequal paths and defaults are present in 78.5% and 26% of consent banners, respectively.

releases/2023/06/ftc-takes-action-against-amazon-enrolling-consumers-amazon-prime-without-consent-sabotaging-their

⁷https://insightplus.bakermckenzie.com/bm/technology-media-telecommunications_1/united-states-consumer-protection-regulators-set-sights-on-dark-patterns.

⁸<https://gdpr-info.eu/issues/consent/>; <https://gdpr-info.eu/recitals/no-32/>.

⁹<https://www.warner.senate.gov/public/index.cfm/2023/7/warner-fischer-lead-bipartisan-reintroduction-of-legislation-to-ban-manipulative-dark-patterns>

¹⁰<https://www.appsflyer.com/blog/tips-strategy/apps-boost-att-opt-in/>

¹¹<https://www.deceptive.design/types/confirmshaming>

The third and final type of pattern consists of pure visual manipulations, or designs that influence choices via visual elements. One example of a visual manipulation is differential salience, in which designs gray out undesirable options or make the fonts smaller and harder to see. Another example is the reordering of the options to have the company’s preferred option on top. For instance, Apple’s ATT banner is accused by advertisers to nudge consumers away from sharing by ranking the “ask app not to track” on top and for using the phrase “tracking” which has a negative connotation. Although there is research comparing how Apple’s ATT prompt and its native app prompt affects sharing rates differently (Baviskar et al. 2024), no one has measured whether reordering options alone meaningfully influences data sharing choices.

We evaluate a broad range of dark patterns across many websites, making our results broadly applicable. In particular, we evaluate the effect of three different designs: deliberate obstruction (via hiding different options from the main screen), reordering options, and differential salience.¹² By combining browsing to websites of our choice and browsing that consumers organically engage in, we can validate the consistency of synthetic choices with real-world behavior.

3. Experimental Design

The goal of our experiment is to identify how people make privacy choices across many websites and choice architectures. To do this, we use Cookie Manager, a browser extension based on the Webmunk framework for browsing-based experiments (Farronato, Fradkin, and Karr 2024). Study participants install the extension on their Chrome browser. The extension manipulates the browsing experience by displaying pop-ups that prompt users to make consequential cookie tracking choices.

Figure 1 displays all the six interfaces we designed. Design C (“Sett-Acc-Rej”) is what we may consider as a neutral setting, where all options—customize, accept, reject—are displayed with similar colors. The other interfaces are manifestations of three types of dark patterns. For example, Design A (“Acc-Sett”) *hides* the reject option (*deliberate obstruction*); Design D (“Acc-Rej-Sett”) prioritizes accepting cookies by listing it as *the first option (re-ordering options)*; and Design F (“Acc-GreyRej-GreySett”) emphasizes the accept button with a *brighter color* than the other options (*differential salience*). Study participants can click on any of the options displayed, or avoid making an explicit choice by clicking on the X in the top right corner. If they click on cookie settings, they are presented with six different types of cookies to choose from, such as “information storage and access,”

¹²It is impossible to include all possible design patterns in a single study, as companies can always uncover new dark patterns via frequent testing and optimization.

“performance and analytics,” and “ad selection, delivery, and reporting” (see Appendix Table A1). Selecting all options is equivalent to accepting all cookies (and we will categorize it as such); selecting none of the options is equivalent to rejecting all cookies (and we will categorize it accordingly). To minimize choice friction, in this “customize setting” page we also allow consumers to either accept all cookies in one click, or reject all cookies with ease (as the default in this page is selecting none of the category-specific cookies).

In total, we have six different banner variations. In addition to the neutral interface, we have two designs with deliberate obstruction (one removes “Reject all cookies”, the other removes “Accept all cookies”), two re-ordered interfaces (one with “Accept all cookies” on top, the other with “Reject all cookies” on top), and one interface with differential salience (where “Accept all cookies” is at the top in blue, whereas the other options are below in grey).

Our banner can appear on any website. It replaces the organic cookie choice interface whenever one is present. Note that we experiment on US residents, so most of them rarely see such banners, given the lack of federal legislation on the topic. When an organic banner is detected, Cookie Manager attempts to enforce the accept or reject decision made through our own banner by clicking on the appropriate option in the organic banner. When users hit ‘x’ on our banner, cookie manager does not attempt to change any settings on the website.

The randomization of banners is at the *user by domain* level: the browser extension randomly selects a cookie interface to show to the participant when they first visit a domain since enrollment in the study, and tracks the corresponding user selection.¹³

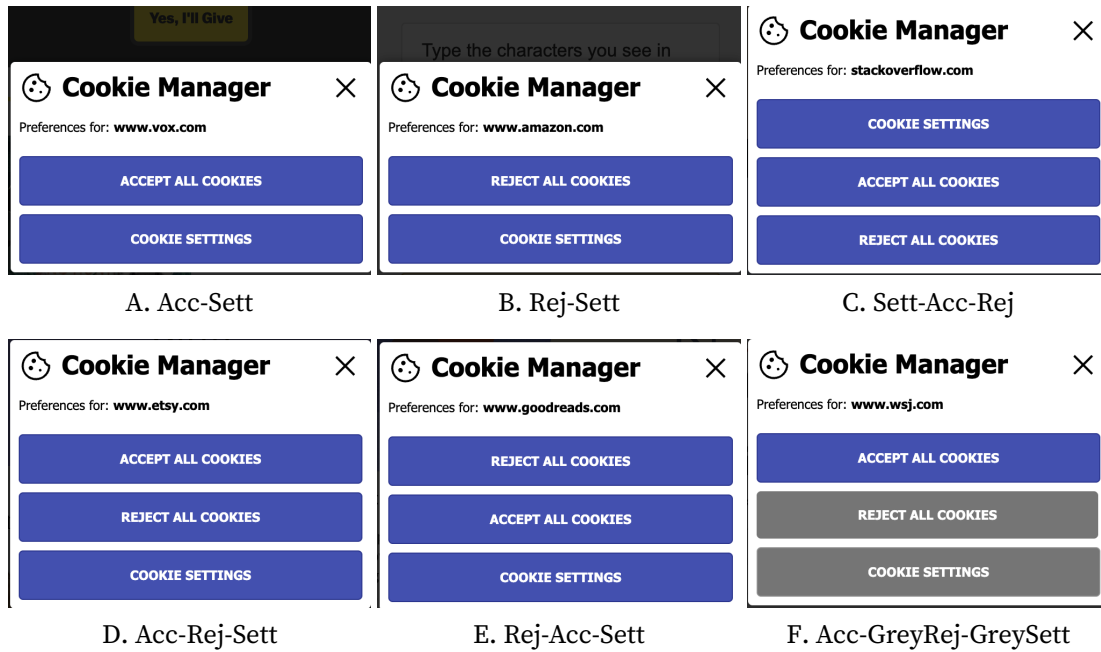
After eligibility and install of the browser extension, the study proceeds in two phases. The first phase is survey-based. Here, we ask participants to visit 20 pre-selected websites.¹⁴ Websites are randomly ordered when we ask participants to visit them. As they visit each of the websites, they make cookie sharing choices as a function of the assigned pop-up. We choose websites to cover a wide range of categories (from social media and shopping to search and news) and website popularity. The goal of this phase is to ensure we have choices across many individuals for each website, whether users normally visit them.

The second phase relies on participants’ organic browsing behavior. We ask participants to keep the extension installed for a week, during which we track their browsing

¹³In pilot studies, we tested for carryover effects of exposure to the initial choice architecture on all subsequent cookie choices, and found null effects. For this reason, we chose to randomize cookie interfaces at the user-web domain level, which allows more variation than randomizing at the user level only.

¹⁴The websites are the following: youtube.com, nytimes.com, appleinsider.com, yahoo.com, amazon.com, ebay.com, target.com, etsy.com, turo.com, stockx.com, espn.com, facebook.com, funnyordie.com, weather.com, duckduckgo.com, truewerk.com, thomannmusic.com, merrysky.com, seattletimes.com, semafor.com.

FIGURE 1. Consent Interface Design across Treatment Groups



Notes: The figure provides screenshots of the six cookie preference interfaces. Captions correspond to the labels used throughout the paper to refer to the treatment conditions. “Acc-Sett”: accept all-settings; “Rej-Sett”: reject all-settings; “Sett-Acc-Rej”: settings-accept all-reject all; “Acc-Rej-Sett”: accept all-reject all-settings; “Rej-Acc-Sett”: reject all-accept all-settings; “Acc-GreyRej-GreySett”: accept all-grey reject allgrey settings.

behavior. Instead of showing banners for every new web domain visited, we randomize users into one of two treatment conditions: in the *frequent pop-ups* condition, a pop-up appears every 10 minutes a user spends browsing the Internet; in the *infrequent pop-ups* condition, a pop-up appears every 60 minutes. At the end of the week, participants are asked to fill out a short outtake survey and uninstall the extension. We pay each participant \$7.50 at the conclusion of the study. The full set of survey questions is available in Appendix B.

4. Sample Description

We recruited participants on Prolific.¹⁵ We pre-registered recruiting 800 participants and expected 640 of them to complete the study.¹⁶ Appendix Table A1 confirms that our actual participants are close to the pre-registered numbers. A total of 1,227 Prolific users started answering questions to assess eligibility to participate in the study (line 1 in the table). A

¹⁵<https://www.prolific.com/>.

¹⁶<https://www.socialscisceregistry.org/trials/12862>.

respondent was deemed eligible to participate in the study if they resided in the US, were older than 18, spoke English as their primary language, and primarily used Chrome to browse the Internet. 75% of respondents were eligible (917 individuals). Among these, 877 of them consented to the study (line 3 in Table A1), but only 613 participants fully completed the study (line 9). For our final sample, we included everyone who completed the baseline survey and generated valid data points during the organic browsing phase, regardless of whether they proceed to the outtake survey stage. For our main analysis, we further excluded organic site visits before the end of survey phase and surveyed site visits after, as well as participants who did not get randomized into either the 10-minute or 60-minute treatment during the organic browsing phase (3% of the total).¹⁷ These restrictions allow us to maintain a consistent sample for the analysis of both organic and survey responses. As a result of these selection criteria, we have a total of 656 participants in our main analysis sample (line 8).

Table 1 presents descriptive statistics for the main sample. Starting with user demographics, we have a balanced sample between men (54%) and women (46%), and the average age is 38 years old. The median household income in our sample was \$50,000-\$74,999, with substantial variation, including 12% of households with an income of over \$150,000.

Next, we consider user browsing behavior. For the week preceding enrollment in the study, users visit an average of 51 unique domains. During the week of the study after the survey, we record 78 unique domains, suggesting that users do not avoid using the browser on which they are tracked. We also have high compliance during the survey period. On average, we record participants making choices on consent banners for 19 domains out of the 20 they were asked to visit.

Finally, during the week when we record their organic browsing behavior, our pop-up banners show up in 28 percent of the visited domains. This average masks heterogeneity induced by our experiment. For 47 percent of users, the pop-up appears every 10 minutes of browsing, which implies that 36 percent of the domains visited experience a pop-up. For the rest of users in our sample, the pop-up appears in 22 percent of the domains visited (see Appendix Table A2).

We verify pop-up design randomization in two ways. First, we run a proportion tests on pop-up distribution per website. The proportion test for the distribution of pop-ups across the survey websites has a p-value of 0.99, which fails to reject the null of balanced proportions across the 6 pop-up designs. Second, we perform covariate balance tests

¹⁷Due to an implementation challenge, not all users were randomized into the 10- or 60-minute treatment after they completed they survey. This affects only 3% of the users in the study who completed the initial survey.

TABLE 1. Summary Statistics

		Mean	Median	Std. Dev.
During Survey	Unique Domains in Prior Week	51.43	47.00	39.11
	Domains w. Banner	18.59	20.00	3.96
Post-Survey	Domains w. Banner	21.69	15.00	21.80
	Unique Domains Visited	77.79	60.00	62.13
	Unique URLs	1211.45	579.00	3537.73
Demographics	End Survey Completed	0.85	1.00	0.35
	Age	38.16	36.00	13.04
	Female	0.46	0.00	0.50
	Bachelor's or Above	0.18	0.00	0.39
Cookie Behavior	Accept-All Rate	0.53	0.62	0.37
	Close-Window Rate	0.28	0.15	0.33
	Reject-All Rate	0.15	0.00	0.28
Experimental Group	10-Minute Group	0.47	0.00	0.50
	60-Minute Group	0.53	1.00	0.50

Notes: The table shows descriptive statistics for the final study sample. Number of observations: 656.

by regressing user-level and domain-level covariates on treatment conditions (Appendix Table A3). We find no statistically significant differences across pop-up designs.

In the organic browsing phase, in addition to randomizing the pop-up design at the user-domain level, we also randomize the frequency of pop-up appearance at the user level. Appendix Table A2 provides descriptive statistics broken down by the two treatment groups. The two groups are balanced across all user characteristics. None of the differences are statistically significant, except for the number of banners shown in the two groups, which is induced by the experimental randomization. The proportion test for the distribution of pop-ups across the organic visits has a p-value of 0.00093, rejecting the null. Although significantly different, the proportions range from 0.158 to 0.175, which are close to the intended proportions (0.166).

5. Main Results

In this section, we describe the causal effects of dark patterns, additional determinants of consent decisions, heterogeneity across users and domains, and results on choice fatigue. We find that the majority of users accept all cookies in the neutral condition, and dark

patterns substantially affect consent decisions. There is little heterogeneity in choices across domain popularity but a lot of heterogeneity across users. Lastly, we find limited evidence of choice fatigue.

We observe consent decisions in two phases. In the first phase, which we call *survey*, users visit websites as part of the survey and respond to questions about site familiarity and usage. In the second context, which we call *organic*, users organically browse websites in the week after the first survey. We present most results separately for survey and organic phases, since users may behave differently during these two stages.

Figure 2 presents the choice distribution across treatment conditions, separately for survey choices (top panel) and organic choices (bottom panel). There are three main findings. First, participants choose to share their data with the websites they interact with for more than 50% of the time, even when the design is neutral. In the conditions for which the “Accept Cookies” options is immediately displayed, the accept rate is between 62 and 70 percent in the survey and 55 to 60 percent in organic browsing.¹⁸ The exception to this is for the interface where “Accept Cookies” is excluded from the initial options. In this condition, 21 percent in the survey and 17 percent in organic browsing achieve it by clicking all cookie options under “Settings.”

Second, the figures show a small percentage of granular choices (in light blue). Across all treatment conditions, participants make selective cookie choices only occasionally, ranging from 4% in the neutral condition to 6-9% in conditions that deliberately hide either the “accept” or the “reject” option.

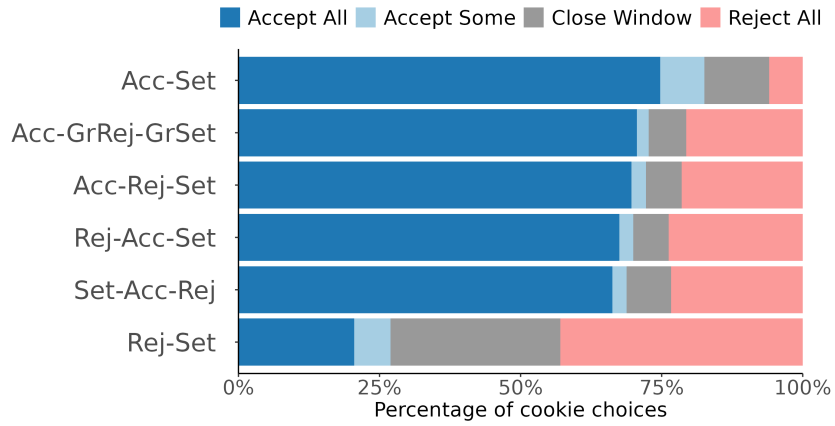
Third, participants react similarly to dark patterns when engaging in the survey and when browsing organically. This gives us confidence that choices during the survey, which may be considered more artificial, nonetheless represent dark pattern effects. The most notable difference between survey and organic behavior is that participants are more likely to close the window and less likely to explicitly reject cookies during the organic browsing phase. As a result, they are also less likely to make active “Accept Cookies” and “Reject Cookies” decisions in the organic phase.

Next, we measure the causal effects of dark patterns on consumer choices. We estimate the following regression specification:

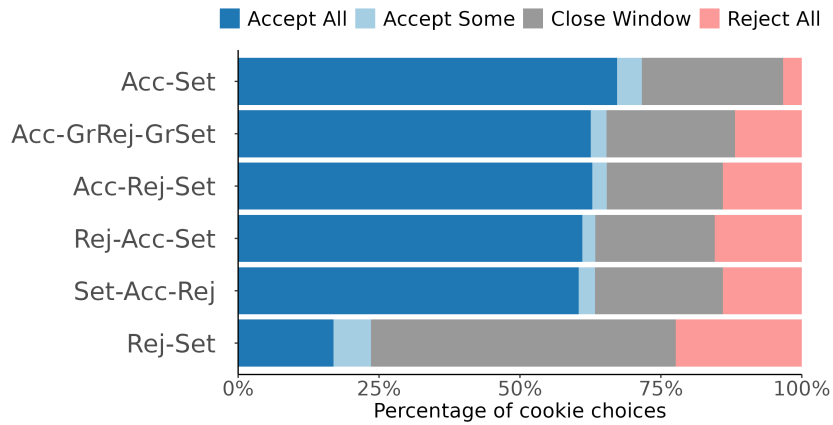
$$y_{is} = \beta_{acc-sett} + \beta_{acc-greyrej-greysett} + \beta_{acc-rej-sett} + \beta_{rej-acc-sett} + \beta_{rej-sett} + \mu_i + \nu_{c(s)} + \epsilon_{is}, \quad (1)$$

¹⁸Our acceptance rates are high, but if anything lower than prior evidence. For example, Bielova et al. (2024) show that 83% of participants in their artefactual study accept cookie tracking when the choice is offered with a “neutral” design.

FIGURE 2. Cookie Choices by Experimental Condition



A. Survey Choices



B. Organic Choices

Notes: This figure displays the proportions of cookie choices across banner design treatments. The possible choices are: Reject all cookies, Close window (i.e., the user clicks on the X of the pop-up window to close it), Accept some cookies (i.e., a user clicks on settings and select a subset of cookie types), and Accept all cookies. All choices indicate the final choice, e.g., “Accept All” includes instances where a participant clicks into the “settings” page and manually selects all cookies. Each row corresponds to a treatment condition. The mapping of the labels to each interface is presented in Figure 1.

where i denotes the participant, and s denotes the website. We include participant fixed effects μ_i and website category fixed effects $\nu_{c(s)}$.¹⁹ Each of the β coefficients measures

¹⁹We obtain website category information using WebShrinker (<https://webshrinker.com/>), a popular web categorization API which categorizes websites using labels from the Interactive Advertising Bureau (<https://www.iab.com/>).

the effect of a specific treatment condition relative to the neutral setting (Condition C in Figure 1).²⁰ We focus on three outcomes: whether the user accepts all cookie tracking, whether they refuse all tracking, and whether they close the window without making an active choice. Given its small share, the analysis on decision to select specific types of cookies is left to Appendix A5.

Table 2 displays our main results, with standard errors clustered at the participant level. Columns 1 through 3 focus on survey choices, whereas columns 4 through 6 focus on organic choices. Under the neutral setting, 62% of participants accept all cookie tracking, 23% reject all tracking, 11% close the window, and the rest select to be tracked by a subset of the cookies. Removing the reject button increases tracking the most, by 7.9 percentage points (or a 13% increase). Next, the design with accept at the top and the other options greyed out leads to a 3.6 percentage point increase (6%) in the share of participants choosing tracking. Putting the accept option at the top without differential salience marginally increases acceptance rates, by 2 percentage points. Having the reject option at the top is statistically equal to the neutral setting with small point estimates. Finally, removing the accept button has a large negative effect on the tracking share, which decreases by 46 percentage points (a 75% reduction).

Moving to effect on rejection rates in column 2, even the most-privacy preserving design (i.e., removing the accept all option) fails to move the reject rate above 50%. Specifically, this design increases the share of users who reject cookies from 23% to 43%. In contrast, hiding the reject button reduces the proportion rejecting by 17 percentage points (a 74% decline). Other conditions have smaller effects.

Lastly, we consider rates of closing out of pop-up window (Column 3). When the accept button is hidden, the proportion of participants closing the pop-up increases by over 200%. This suggests that a large share of users prefer to accept cookies when that option is easily available, but avoid explicitly rejecting cookies even when that is possible with the click of a button. The other conditions have smaller effects.

Similarly to the survey choices, dark patterns have important effects in organic websites as well (Columns 4 to 6). There is one main difference between the survey and organic results. The baseline proportion of participants closing the pop-up without making an active choice is 28% in organic websites (against 11% in survey choices). This result may reflect the fact that users have less patience or time for making active cookie tracking choices when browsing the Internet organically. However, as we show below, users' beliefs about the implications of closing the window translate to very similar proportions of users who accept and reject cookies across survey and organic choices.

²⁰Relative to the pre-registered specification, we have changed the baseline design to be the neutral design to be consistent with the existing literature in computer science (see Bielova et al. (2024) for example).

TABLE 2. Cookie Choices by Experimental Condition

	Survey			Organic		
	Accept All (1)	Reject All (2)	Close Window (3)	Accept All (4)	Reject All (5)	Close Window (6)
Acc-Set	0.079*** (0.012)	-0.172*** (0.014)	0.039*** (0.009)	0.055*** (0.011)	-0.094*** (0.012)	0.016 (0.012)
Rej-Set	-0.464*** (0.019)	0.199*** (0.016)	0.228*** (0.017)	-0.429*** (0.022)	0.094*** (0.012)	0.299*** (0.022)
Acc-Rej-Set	0.020* (0.010)	-0.007 (0.008)	-0.010 (0.006)	0.035** (0.011)	0.002 (0.007)	-0.041*** (0.009)
Rej-Acc-Set	0.005 (0.009)	0.012 (0.009)	-0.013* (0.006)	0.006 (0.010)	0.022** (0.008)	-0.026** (0.009)
Acc-GrRej-GrSet	0.037*** (0.009)	-0.019* (0.009)	-0.012* (0.006)	0.033** (0.010)	-0.017* (0.008)	-0.019* (0.009)
Benchmark group mean	0.62	0.23	0.11	0.56	0.13	0.28
R ²	0.653	0.582	0.573	0.580	0.522	0.510
Observations	12,142	12,142	12,142	14,163	14,163	14,163
Participant fixed effects	✓	✓	✓	✓	✓	✓
Domain Cat. fixed effects	✓	✓	✓	✓	✓	✓

Notes: Regression results of Equation 1 for three outcomes: accept all cookies, reject all cookies, and close window without making a choice. The results are presented separately for two different sets of choices: survey choices (columns 1 through 3) and organic choices (columns 4 through 6). Appendix Table A5 presents similar results for the decision to accept a subset of cookie types. * $p < 0.05$, ** $p < 0.01$; *** $p < 0.001$.

Table 3 repeats the analysis by estimating the effect of the three dark patterns we identified, rather than the effect of each treatment condition. This means estimating regressions of the following type:

$$Y_{is} = \gamma_{reject\ hidden} + \gamma_{accept\ hidden} + \gamma_{accept\ on\ top} + \gamma_{reject\ on\ top} + \gamma_{highlight\ accept} + \mu_i + \nu_{c(s)} + \epsilon_{is} \quad (2)$$

Compared to Equation 2, the dummies for treatment conditions are replaced with dummies describing which choice option is on top (accept or reject), which option is hidden (accept or reject), and which option is highlighted (accept).

Table 3 displays the results. Deliberate obstruction is the most effective pattern. Hiding the “reject” button from the main screen increases the probability of accepting all cookies by 7.9 percentage points during the survey phase and 5.3 percentage points during the organic browsing phase. The effect of hiding the “accept” button is even more drastic, decreasing the probability of accepting cookies by 47 percentage points during survey phase and 43.4 percentage points during organic browsing. Another impact of deliberate obstruction is increasing disengagement. In particular, hiding the “accept” button increases the probability of closing the pop-up window by 24.2 percentage points

(during the survey) to 32.6 percentage points (during organic browsing). The fact that hiding the “accept” button has a higher impact on cookie sharing choices than hiding the “reject” button likely reflects the fact that the baseline sharing probability is high, and thus hiding this option affects choices more.

In comparison, designs that rely on changing the visual presentation alone, mainly reordering and differential emphasis, are less effective in changing choices, reflected by both the small coefficients associated with them and the statistical insignificance of most estimates. We note that their effects increase during the organic phase, though still with small magnitudes. For example, putting the accept button at the top increases the proportion of accepting cookies by 3.3 percentage points while decreasing the probability of closing window by 4.0 percentage points. Similarly, ranking the reject button at the top increases the rejection rate by 2.2 percentage points while decreasing the close-window probability by 2.5 percentage points. On the other hand, graying out both the reject and the settings option leads participants to substitute from explicitly rejecting cookies to closing the window for 2% of the time, without affecting their cookie acceptance rate.

TABLE 3. Cookie Choices by Dark Pattern

	Survey			Organic		
	Accept All (1)	Reject All (2)	Close Window (3)	Accept All (4)	Reject All (5)	Close Window (6)
Reject Hidden	0.079*** (0.012)	-0.172*** (0.014)	0.039*** (0.009)	0.055*** (0.011)	-0.094*** (0.012)	0.016 (0.012)
Accept Hidden	-0.469*** (0.019)	0.187*** (0.016)	0.240*** (0.017)	-0.435*** (0.022)	0.071*** (0.011)	0.325*** (0.022)
Accept Top	0.020* (0.010)	-0.007 (0.008)	-0.010 (0.006)	0.035** (0.011)	0.002 (0.007)	-0.041*** (0.009)
Reject Top	0.005 (0.009)	0.012 (0.009)	-0.013* (0.006)	0.006 (0.010)	0.022** (0.008)	-0.026** (0.009)
Highlight Accept	0.017 (0.010)	-0.012 (0.009)	-0.002 (0.006)	-0.002 (0.010)	-0.019** (0.007)	0.022* (0.009)
Benchmark group mean	0.62	0.23	0.11	0.56	0.13	0.28
R ²	0.653	0.582	0.573	0.580	0.522	0.510
Observations	12,142	12,142	12,142	14,163	14,163	14,163
Participant fixed effects	✓	✓	✓	✓	✓	✓
Domain Cat. fixed effects	✓	✓	✓	✓	✓	✓

Notes: Regression results of Equation 2 for three outcomes: accept all cookies, reject all cookies, and close window without making a choice. The results are presented separately for two different sets of choices: survey choices (columns 1 through 3) and organic choices (columns 4 through 6). Appendix Table A6 presents similar results for the decision to accept a subset of cookie types. * $p < 0.05$, ** $p < 0.01$; *** $p < 0.001$.

These results point to three main conclusions: users often accept cookie tracking absent dark patterns while browsing the web; dark patterns that increase choice frictions

are effective in changing people's choices;²¹ closing the window without making an active choice is a frequent selection, even more true in the wild than in a synthetic survey-based setting, which shows the importance of website default tracking options in the absence of a choice.

Our results show that closing the window without making an active choice is a frequently chosen option. It is thus critical to identify users' beliefs when they make this choice. Among our participants, 60% believe that closing the window is akin to rejecting cookies. Another 26.2% believe the opposite, that closing the window means accepting cookie tracking. The rest of the participants believe neither is true, and instead conjecture that the website will fall back to its default settings, ask for consent during the next visit, or simply express uncertainty about what the website will do. Users who believe that closing the window is akin to rejecting cookies close the window at the same rate as those who think closing the window implicitly accepts cookies. However, both groups are 50% more likely to close the window than participants with ambivalent beliefs.

Incorporating user beliefs when closing the window into our results confirms that survey and organic choices are much more similar than it may appear at first sight. Indeed, adding these beliefs to the active choices from Figure 2 implies that 68 percent ($0.66 + 0.262 * 0.078$) of users in the survey and 66 percent ($0.6 + 0.262 * 0.23$) in organic browsing effectively accept cookies. Similarly, the results imply that 28 percent ($0.232 + 0.6 * 0.78$ in the survey and $0.139 + 0.6 * 0.23$ in organic browsing) of users effectively reject cookies.

Table A6 indicates that consumers tend not to make granular cookie choices, and would rather opt out of making choices altogether by closing the consent window. In the neutral design group, only 4% of participants accept a subset of cookie types; deliberately hiding either "accept all" or "reject all" options from the main screen encourages participants to check out the settings menu, but these designs still only increase the probability of granular choices by 2-5 percentage points. Among those who make granular selections, 83% choose to accept cookies for *preferences and functionality*, while only 7% accept cookies for *ad selection, delivery, and reporting* (see Appendix Table A7). This result suggests that targeted advertising is the least preferred use of consumer data, at least among the few users who make selective choices.

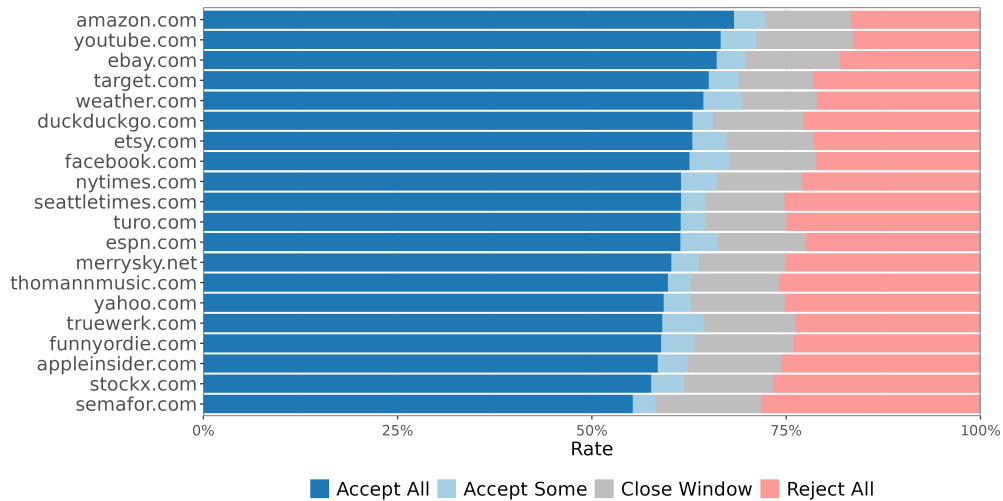
²¹These causal effects are broadly in line with existing findings in artefactual or survey experiments. For instance, Habib et al. (2022) compare a design where the reject option is hidden with a design where rejecting is the default, and found a sizable difference in choices among the two groups. Both Utz et al. (2019) and Vásquez Duque (2024) examine the effect of differential salience designs, and found small to no effect on choices.

5.1. Heterogeneity of Cookie Choices by Domain Characteristics and User Familiarity

A critical concern of our study is whether cookie consent forms and dark patterns affect competition. To investigate competitive effects, we consider heterogeneity of consent rates across domains and website categories. We find that users are more likely to consent for more familiar websites, but familiar and popular websites do not enjoy a differential advantage of using dark patterns.

We begin by looking at our survey domains, for which we have a balanced panel of user-by-domain interactions. The most known domains in our survey have the highest accept rates. Figure 3 displays cookie choices across individuals and pop-up designs for each website separately, ranked according to the propensity to accept all cookies. The websites at the top, such as amazon.com, youtube.com, and ebay.com, are well-known and frequently used. 66-68% of users accept cookie tracking on these websites. At the other extreme are lesser-known websites such as truewerk.com, merrysky.net, semafor.com, which only 2-3% of participants indicate they normally visit. Still, even for these sites, the probability to accept cookie tracking is at least 55%.

FIGURE 3. Cookie Choices by Survey Website



Notes: The figure shows the breakdown of user choices across the 20 websites in the survey. The websites are ordered from highest to lowest cookie acceptance rate.

To understand how cookie sharing choices vary across websites and user experience, we consider measures of user familiarity with the website. The first two measures come from our survey phase, where for each website, we ask consumers whether they have *heard of* and whether they *normally visit* a website. These variables are available for the 20 websites featured in the survey phase. We additionally construct two other metrics for all websites present in our data to characterize participants' familiarity with the site and their

popularity. To measure site familiarity, we compare each site visited during our study with the participant’s browsing history in the previous two weeks, and flag *pre-exp visit* as one when the website is listed in the browsing history. We also use log *domain rank* as a proxy for website popularity. The domain rank data comes from Tranco, which provides a stable ranking of websites based on an aggregation of several ranking approaches.²²

Table 4 shows how participants’ sharing decisions vary with these variables. Each of the three panels represents a set of regressions of Equation 1, to which we add additional controls: *normally visit* and *heard of* in Panel *a*, *pre-exp visit* in Panel *b*, and *domain rank* (*log 10*) in Panel *c*.

TABLE 4. Heterogeneity in Cookie Choices across Websites

Model:	Survey			Organic		
	Accept All (1)	Reject All (2)	Close Window (3)	Accept All (4)	Reject All (5)	Close Window (6)
<i>Panel a: Familiarity based on survey answers</i>						
Normally Visit	0.062*** (0.009)	-0.065*** (0.010)	-0.013* (0.006)			
Heard Of	0.012 (0.010)	-0.024* (0.010)	0.004 (0.008)			
<i>Panel b: Familiarity based on browsing history</i>						
Pre-Exp Visit	0.024* (0.011)	-0.027* (0.010)	-0.002 (0.007)	0.039*** (0.008)	0.003 (0.005)	-0.043*** (0.007)
<i>Panel c: Website popularity</i>						
Domain Rank (Log 10)	-0.017*** (0.003)	0.014*** (0.003)	0.002 (0.002)	0.001 (0.002)	-0.002 (0.002)	0.000 (0.002)

Notes: Regression results of Equation 1 to which we add explanatory variables to explore heterogeneity in cookie tracking choices. In Panel *a*, we add two dummies for whether the study participant has heard of the website, and for whether they study participant normally visits the website. In Panel *b*, we add a dummy for whether the study participant visited the website in the two weeks preceding the study (we obtain this information by collecting their Chrome browsing history). In Panel *c*, we add the website popularity rank (in logs). * $p < 0.05$, ** $p < 0.01$; *** $p < 0.001$.

In the survey websites, all three proxies of familiarity are correlated with cookie tracking choices in the expected direction. If a user normally visits a website,²³ has visited the website in the preceding two weeks, or if the website is more popular (i.e., lower rank), the user is more likely to accept cookies (columns 1 and 2). In particular, having a user interacting with a site that they normally visit is associated with a 6 percentage point

²²<https://tranco-list.eu/methodology>

²³The question about whether a user normally visits a websites happens after the treatment for a particular domain, so there could be some post-treatment bias in this question.

increase in the acceptance rate (panel a). On the other hand, having visited a website in the prior two weeks (panel b), and increasing site popularity 10-fold (panel c) are both associated with a 2 percentage point increase in the acceptance rate. These effects seem relatively small, especially in light of work on GDPR showing that it tilted the playing field in favor of larger and better-known sites (Goldberg, Johnson, and Shriver 2024).

In the organic choices, we do not have access to the first proxy for familiarity. However, if a user visited a website in the preceding two weeks (which in this case, we take as a proxy for the frequency with which a user visits) the user is 3.9 percentage points more likely to accept cookies (column 4). In this case however, the substitution is not coming from rejection, but rather from closing the window (column 6). In addition, the correlation between domain rank and acceptance decision is absent in the organic browsing data. This presumably happens because participants are already browsing their site of choice, thus a site being popular has less additional effect in changing their cookie tracking decision. In comparison, surveyed sites are pre-specified, in which case rankings and credibility of a site could matter more.

Next, we examine whether there was heterogeneity in treatment effects across websites. Figures A2 and A3 show that there is some heterogeneity in dark pattern effects across different surveyed site domains. Interestingly, whenever dark patterns have different effects across sites, they seem to alleviate participants' inclination to share data with popular and familiar sites (see Appendix Figures A8, A9, and A10).

On the other hand, we do not find evidence of substantial treatment effect heterogeneity during organic visits for domain ranking (A11). Neither do we find any evidence of treatment effect heterogeneity in terms of whether we observe a user having visited the website before (A12). Taken together, the ability to use dark patterns does not differently help popular and incumbent websites over others, at least in terms of the types of data collected.

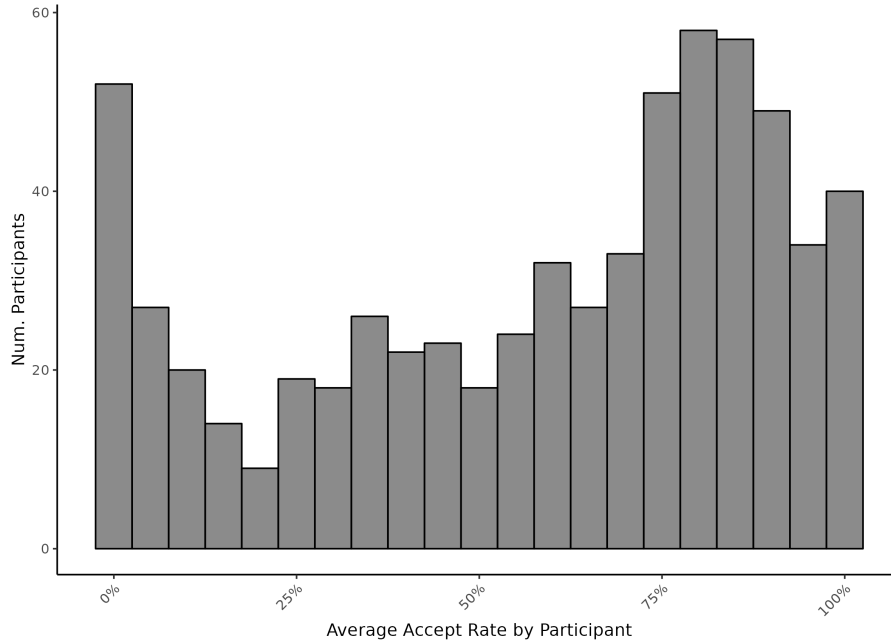
5.2. Heterogeneity across Individuals

Next, we consider how different individuals choose to accept cookies, and how this choice varies within individual across websites. We observe substantial heterogeneity in cookie choices across websites, consistent with the previous literature that privacy choices vary heavily with the economic context (Nissenbaum 2004; Lin 2022).

Figure 4 displays the distribution of accept rates across individuals, combining data from both the survey and organic data. We see a bimodal distribution, with some users never accepting and other users accepting most of the time. That said, an overwhelming

majority have at least some variation in choices — 91.8% of participants change their cookie sharing decisions across sites at least once. For illustrative purpose, we display the distribution of choices among 50 randomly selected participants (Figure 5) in the neutral condition.²⁴ Interestingly, we do not find systematic differences in demographics across those who accept more or less than the median amount (Table A13).

FIGURE 4. Distribution of Accept-All Rates by Participant (All Sites)



Notes: The figure shows the distribution of accept-all rates across users across survey and non-survey domains.

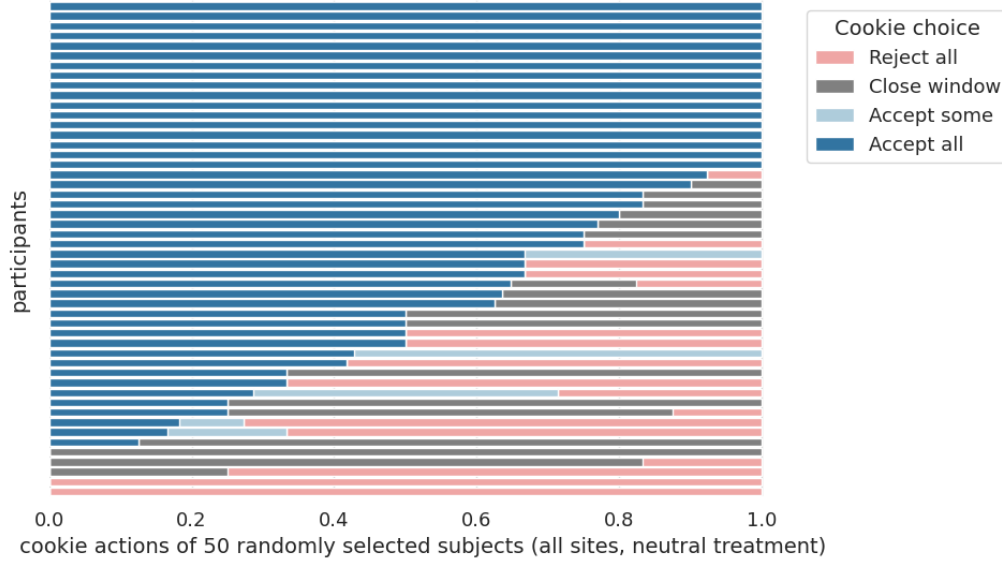
To isolate the choice variation induced by dark patterns from the variation explained by website differences, we examine the extent to which each user makes different cookie choices across websites when faced with the neutral design (*settings-accept-reject*). To allow for variation, we exclude participants who get exposed to the neutral design only once during organic browsing.²⁵ We estimate the following regression specification:

Holding fixed the banner design, 53.5% of participants change their data sharing decisions across sites; 36.1% choose to always accept cookies; and another 5.8% choose to always reject cookie tracking. We conclude that a substantial share of users are paying attention to the specific site making the consent request when making data sharing decisions.

²⁴The variation within users is due to a combination of the effects of users' differing preferences across websites, and simple clicking mistakes.

²⁵5% of study participants see the neutral design only once during the organic phase of the study.

FIGURE 5. Cookie Choice Variation within Individual (Neutral Treatment Group)



Notes: The figure shows the breakdown of user choices within the neutral design treatment 50 randomly selected participants with more than one choice, including both survey and organic site visits.

6. The Time Cost of Consent and Choice Fatigue

One criticism of GDPR and similar regulations that mandate consent is that consent pop-ups degrade the user experience and waste time. In this section, we consider the costs of asking users for consent and examine whether users become fatigued by these choices. To study these questions, we use a second source of randomization present in our design to estimate the causal effects of more frequent exposures to pop-ups.

We begin by calculating how much time people spend interacting with our consent banners and how this varies by banner type. To measure time spent on the pop-up, we calculate the time elapsed between the cookie pop-up and the recorded *final* action (accept all, accept some, reject all, close window). This time spent measure includes the time spent on intermediate clicks, such as on “settings” and when someone goes back and forth before closing the banner. We censor the time spent at 1 minute, since this is well above the 99th percentile of time spent on the banner (13 seconds during organic browsing).

The mean time spent on the banner is 6 seconds in the survey phase and 7.6 seconds in the organic phase. Table 5 displays the results of regressions of time spent on the treatment conditions in columns (1) and (3). The only banner for which we are able to detect effects is the Reject-Settings banner. Users spend more time on the banner in this circumstance, perhaps because they are looking to find the accept all option. In Columns

(2) and (4) we measure which of the four actions takes the most time, conditional on user, domain, and condition. Unsurprisingly, we find that “Accept Selected” takes the most time, since it requires users to click on the settings button and then make additional sub-selections.

TABLE 5. Effects of Dark Patterns on Time Spent

	Survey		Organic	
	Time Spent (Seconds)			
	(1)	(2)	(3)	(4)
Acc-Set	0.506 (0.288)		-0.474 (0.374)	
Rej-Set	2.172*** (0.276)		0.966* (0.389)	
Acc-Rej-Set	-0.364 (0.236)		-0.101 (0.403)	
Rej-Acc-Set	-0.206 (0.245)		-0.317 (0.379)	
Acc-GrRej-GrSet	-0.061 (0.252)		-0.321 (0.384)	
Accept Selected		7.892*** (0.643)		5.266*** (0.893)
Close Window		-0.158 (0.456)		1.267** (0.447)
Reject All		-0.657* (0.301)		0.715 (0.570)
Overall group mean	5.72	5.72	7.55	7.55
R ²	0.210	0.231	0.114	0.117
Observations	12,247	12,247	15,576	15,576
Participant fixed effects	✓	✓	✓	✓
Domain Cat. fixed effects	✓	✓	✓	✓
Condition fixed effects		✓		✓

Notes: Regression of the time spent between a cookie pop-up and final action, where each observation is a user by domain. Columns (1) and (3) contain controls for treatment arm, with the neutral design treatment as the omitted category; columns (2) and (4) control for the final action chosen, with “accept all” as the control group. * $p < 0.05$, ** $p < 0.01$; *** $p < 0.001$.

We can estimate the time costs in interacting with the consent banners. On average, a participant spent 7.6 seconds per domain to make their cookie consent decision. If the banner is shown on every unique domain they visit throughout the week (77.5/week in our

study), then their time spent on consent would be 9.8 minutes/week. A back-of-envelope calculation using the average US hourly wage published by the Bureau of Labor Statistics show that the spent amounts to \$5.73/week per consumer.²⁶ Looking at the effect of dark patterns on time spent, hiding the “accept” button increases the time costs to \$6.42/week. If consumers were to make selective cookie choices for each banner, the cost of time goes up to 12.82 seconds/domain, which amounts to \$9.73/week. In this sense, deliberate obstruction decreases consumer welfare not only because it distorts choices, but also because it increases the time cost of consent decisions.

Consider an alternative design in the policy discussions, which allows users to configure cookie choices at the browser level (aka. global privacy choice). Our calculation suggests that site-specific choices is better for consumers *if* the value they get from site-level customization is greater or equal to \$5.7/week.

6.1. Choice Fatigue

Next, we examine whether the attention users pay to choices changes as they receive more pop-ups. We compare the differences in choices between our 10-minute and 60-minute treatments to show that there isn’t choice fatigue due to the frequency of pop-ups.

The 10-minute treatment experiences 75% more pop-ups compared to the 60-minute treatment. Given this difference, we can see whether the frequency of choice types varies between these two conditions. We estimate the effects of this treatment in the following regression specification.

$$y_{is} = \beta_{10\text{ minutes}} + \gamma * \text{time in study}_{is} + \nu_{c(s)} + \epsilon_{is}. \quad (3)$$

The baseline is the condition where a user sees the pop-up every 60 minutes, while the alternative condition displays a pop-up every 10 minutes. We also control for the time a user has been in the study (post-survey), since this may be correlated with their overall engagement with the study.²⁷

Table 6 displays the results. We highlight two findings. First, we do not find a differential impact of pop-up frequency on data sharing choices, whether it is the acceptance rate or the inclination to close banners. Users make similar choices, whether they see a pop-up every 10 or 60 minutes. Furthermore, these null effects are precisely estimated, the 95% confidence interval excludes effects greater than 7%.

²⁶According to the Bureau of Labor Statistics, the average US hourly wage in July 2024 is \$35.07. See <https://www.bls.gov/news.release/empst.t19.htm>.

²⁷Adding this covariate does not affect whether we detect any treatment effects.

Second, the time in the study has effects on choice. Each additional day in the study increases the share of people closing the pop-up by two percentage points. Since study participants remain in the study for 7 days, which implies that they are 14 percentage points more likely to close the window at the end of the study compared to the first day. The substitution towards closing the window comes from both acceptance and rejection of cookie tracking.

TABLE 6. Fatigue in Cookie Choices During Organic Browsing

	Accept All (1)	Reject All (2)	Close Window (3)
10 Min Pop-up	0.003 (0.035)	0.008 (0.023)	-0.007 (0.030)
Time in Study (Days)	-0.009* (0.004)	-0.005* (0.002)	0.016*** (0.004)
R ²	0.008	0.004	0.011
Observations	14,163	14,163	14,163
Domain Cat. fixed effects	✓	✓	✓

Notes: This table shows estimates of Equation 3, where ‘10 Min Pop-up’ is an indicator for whether the user was in the treatment where pop-ups occurred at a frequency of once every 10 minutes. * $p < 0.05$, ** $p < 0.01$; *** $p < 0.001$.

It is tempting to interpret the time in study as another measure of choice fatigue, but time in time study is not randomly allocated. In Appendix Table A14 we add individual and hour of the day fixed effects, as well as control for the order of the domain visit. Even with these covariates, we see that time in the study reduces acceptance and increases close out. The most likely explanation for this effect is that participants reduce their engagement with the study over time, which is not a function of the frequency of the pop-ups.

7. Conclusions

We examined the impact of dark patterns on consumer privacy choices through a field experiment involving a browser extension that randomizes cookie consent interface designs. We found that consumers often consent to cookie tracking, and dark patterns can influence their choices. Specifically, hiding consent options behind an additional click is the most effective dark pattern, while visual manipulations like re-ordering or highlighting options have smaller effects. There is substantial heterogeneity in consumer choices depending on the website. This finding implies that global privacy control, a

solution that ask consumers to make a single choice applied to all websites, are unlikely to reflect the heterogeneity of their individual preferences across sites.

Larger and better-known firms achieve slightly higher consent rates, giving them a small edge in the data economy. However, the impact of dark patterns does not differentially benefit popular websites or those that participants are familiar with. If anything, dark patterns seem to alleviate consumers' inclination to share more data with popular and familiar sites. As a result, bans on dark patterns are unlikely to negatively affect the ability of smaller firms to compete with larger ones. In fact, if smaller firms also gain more benefit from the same amount of data, dark patterns could potentially intensify competition.

Any benefit of displaying frequent pop-ups needs to be balanced against the time cost of interacting with banners. We show that the average time cost of interacting with consent banners amounts to \$5.7 per week in policy regimes that mandate site-specific consent, and that dark patterns that hide certain options increase this time cost by 12%. In this sense, dark patterns using deliberate obstruction decrease consumer welfare not only by distorting choices but also by increasing time costs. We fail to find evidence of increasing choice fatigue due to pop-up frequency, which suggests that more frequent pop-ups may not affect the quality of choices by participants. However, consumers increase the frequency of directly closing windows over the course of our study, suggesting decreasing attention to our study or some fatigue over time.

Lastly, we've shown that consent rates are high even absent any dark patterns, which underscores the need to better understand why consumers seem willing to share their data in this context. One hypothesis is that they are conscious of the benefits from personalized experiences and targeted advertising. Another hypothesis is that users may be concerned that rejecting cookies could lead to worse experiences (such as having to log in every time they visit a website or being blocked from accessing essential site content). In our view, this is a phenomenon to explore in future research.

Our study has several limitations. First, the use of a browser extension with consent banners designed by us may not perfectly replicate real-world browsing behavior, as participants were aware of being observed. Second, while we included a variety of websites, our sample may not fully represent all types of online environments. Third, the study focused on short-term effects of dark patterns, and further research is needed to understand long-term implications and potential changes in consumer behavior over time. Finally, our findings are primarily relevant to the context of cookie consent banners and may not generalize to other types of privacy decisions or dark patterns used in different contexts.

References

- Acemoglu, Daron, Ali Makhdoumi, Azarakhsh Malekian, and Asu Ozdaglar. 2022. "Too much data: Prices and inefficiencies in data markets." *American Economic Journal: Microeconomics* 14 (4): 218–256.
- Acquisti, Alessandro, Leslie K John, and George Loewenstein. 2013. "What is privacy worth?" *The Journal of Legal Studies* 42 (2): 249–274.
- Aridor, Guy et al. 2024. "Evaluating The Impact of Privacy Regulation on E-Commerce Firms: Evidence from Apple's App Tracking Transparency." 4698374, SSRN Working Paper.
- Baviskar, Sagar, Iffat Chowdhury, Daniel Deisenroth, Beibei Li, and D Daniel Sokol. 2024. "ATT vs. Personalized Ads: User's Data Sharing Choices Under Apple's Divergent Consent Strategies." *Personalized Ads: User's Data Sharing Choices Under Apple's Divergent Consent Strategies* (June 28, 2024).
- Bergemann, Dirk, and Alessandro Bonatti. 2024. "Data, Competition, and Digital Platforms." *American Economic Review* 114 (8): 2553–95. <https://www.aeaweb.org/articles?id=10.1257/aer.20230478>.
- Bergemann, Dirk, Alessandro Bonatti, and Tan Gan. 2022. "The economics of social data." *The RAND Journal of Economics* 53 (2): 263–296.
- Bielova, Nataliia, Laura Litvine, Anysia Nguyen, Mariam Chammat, Vincent Toubiana, and Estelle Harry. 2024. "The Effect of Design Patterns on (Present and Future) Cookie Consent Decisions." In *USENIX Security Symposium*. USENIX Association. *Accepted for publication*.
- Bielova, Nataliia, Cristiana Santos, and Colin M Gray. 2024. "Two worlds apart! Closing the gap between regulating EU consent and user studies." *Harvard Journal of Law & Technology* 37 (3).
- Choi, Jay Pil, Doh-Shin Jeon, and Byung-Cheol Kim. 2019. "Privacy and personal data collection with information externalities." *Journal of Public Economics* 173: 113–124.
- Collis, Avinash, Alex Moehring, Ananya Sen, and Alessandro Acquisti. 2021. "Information frictions and heterogeneity in valuations of personal data." *Available at SSRN 3974826*.
- D'Assergio, Caterina, Puneet Manchanda, Elisa Montaguti, and Sara Valentini. 2022. "The race for data: Gaming or being gamed by the system?" *Available at SSRN 4250389*.
- Decarolis, Francesco, Muxin Li, and Filippo Paternollo. 2023. "Competition and Defaults in Online Search." *Available at SSRN 4660406*.
- Di Geronimo, Linda, Larissa Braz, Enrico Fregnan, Fabio Palomba, and Alberto Bacchelli. 2020. "UI dark patterns and where to find them: a study on mobile applications and user perception." In *Proceedings of the 2020 CHI conference on human factors in computing systems*,: 1–14.
- Farronato, Chiara, Andrey Fradkin, and Chris J. Karr. 2024. "Webmunk: A New Tool for Studying Online Consumer Behavior." White paper.
- Goldberg, Samuel G, Garrett A Johnson, and Scott K Shriver. 2024. "Regulating privacy online: An economic evaluation of the GDPR." *American Economic Journal: Economic Policy* 16 (1): 325–358.
- Habib, Hana, Megan Li, Ellie Young, and Lorrie Cranor. 2022. "'Okay, whatever': An evaluation of cookie consent interfaces." In *Proceedings of the 2022 CHI conference on human factors in computing systems*,: 1–27.
- Hagiu, Andrei, and Julian Wright. 2023. "Data-enabled learning, network effects, and competitive advantage." *The RAND Journal of Economics* 54 (4): 638–667.

- Ho, Kate, Joseph Hogan, and Fiona Scott Morton. 2017. "The impact of consumer inattention on insurer pricing in the Medicare Part D program." *The RAND Journal of Economics* 48 (4): 877–905.
- Huck, Steffen, and Jidong Zhou. 2011. "Consumer behavioural biases in competition: A survey." Ichihashi, Shota. 2021. "The economics of data externalities." *Journal of Economic Theory* 196: 105316.
- Lin, Tesary. 2022. "Valuing intrinsic and instrumental preferences for privacy." *Marketing Science* 41 (4): 663–681.
- Lin, Tesary, and Avner Strulov-Shlain. 2023. "Choice architecture, privacy valuations, and selection bias in consumer data." *arXiv preprint arXiv:2308.13496*.
- Luguri, Jamie, and Lior Jacob Strahilevitz. 2021. "Shining a light on dark patterns." *Journal of Legal Analysis* 13 (1): 43–109.
- Mathur, Arunesh, 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." *Proceedings of the ACM on Human-Computer Interaction* 3 (CSCW): 1–32.
- Nissenbaum, Helen. 2004. "Privacy as contextual integrity." *Wash. L. Rev.* 79: 119.
- Nouwens, Midas, 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*,: 1–13.
- Spiegler, Ran. 2014. "Competitive framing." *American Economic Journal: Microeconomics* 6 (3): 35–58.
- Tang, Huan. 2023. "The value of privacy: Evidence from online borrowers." *Available at SSRN*.
- Tomaino, Geoff, Klaus Wertenbroch, and Daniel J Walters. 2023. "Intransitivity of consumer preferences for privacy." *Journal of Marketing Research* 60 (3): 489–507.
- Utz, Christine, Martin Degeling, Sascha Fahl, Florian Schaub, and Thorsten Holz. 2019. "(Un) informed consent: Studying GDPR consent notices in the field." In *Proceedings of the 2019 acm sigsac conference on computer and communications security*,: 973–990.
- Vásquez Duque, Omar. 2024. "Beyond the Banner: Exploring User Knowledge of Cookies and Attitudes Toward Targeted Advertising." *Available at SSRN 4815717*.
- Warberg, Logan, Vincent Lefrere, Cristobal Cheyre, and Alessandro Acquisti. 2023. "Trends in Privacy Dialog Design after the GDPR: The Impact of Industry and Government Actions." In *Proceedings of the 22nd Workshop on Privacy in the Electronic Society*,: 107–121.

Appendix A. Additional Tables and Figures

FIGURE A1. Cookie Manager's User Interface

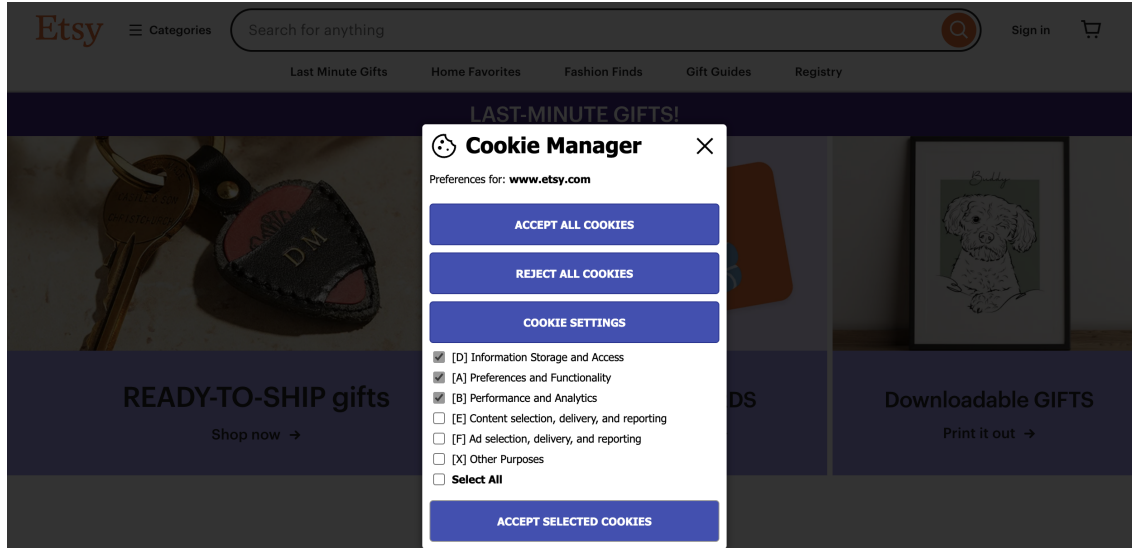


TABLE A1. Number of Participants across the Experimental Funnel

Stage	N	Percent	10 min	Percentage	60 min	Percentage
1) Start Survey	1,227	100.00				
2) Eligible for Study	917	74.74				
3) Study Consent	877	71.48				
4) Finished Survey	807	65.77				
5) Clicked All Links	804	65.53	359	100.00	418	100.00
6) Have Browser Cookie Data	789	64.30	357	99.44	415	99.28
7) Have Cookie Choice Data	787	64.14	356	99.16	414	99.04
8) Main Analysis Sample*	656	53.46	306	85.24	350	83.73
9) Finished Endline Survey	613	49.96	273	76.04	325	77.75

Notes: This table presents the number of study participants at every step of the study. After completing the initial survey, participants are randomly allocated to two treatment conditions: 10 minutes (where cookie pop-ups appear every 10 minutes of browsing), and 60 minutes (where cookie pop-ups appear every 60 minutes). Due to an implementation glitch, not all users are randomized into either the 10- or 60- minute treatment; 3% of participants kept seeing a banner for every new domain visited.

*: The main analysis sample in the second-to-last line restricts attention to users who have treatment assignment to either the 10-Minute or 60-Minute group, and for whom we observe at least one cookie selection both during and after the survey.

TABLE A2. Summary Statistics by Treatment

		10 Minutes			60 Minutes		
		Mean	Median	Std. Dev.	Mean	Median	Std. Dev.
During Survey	Unique Domains in Prior Week	52.28	47.00	40.26	50.68	46.00	38.12
	Domains w. Banner	18.67	20.00	3.71	18.53	20.00	4.16
Post-Survey	Domains w. Banner	28.23	18.00	27.21	15.95	12.00	13.19
	Unique Domains Visited	79.62	59.00	65.21	76.18	61.00	59.35
	Unique URLs	1454.68	579.00	4987.66	998.28	581.00	1274.56
Demographics	End Survey Completed	0.84	1.00	0.36	0.86	1.00	0.34
	Age	38.95	36.00	13.12	37.47	35.00	12.95
	Female	0.48	0.00	0.50	0.44	0.00	0.50
	Bachelor's or Above	0.18	0.00	0.39	0.18	0.00	0.39
Cookie Behavior	Accept-All Rate	0.52	0.62	0.36	0.54	0.66	0.37
	Close-Window Rate	0.27	0.14	0.32	0.29	0.15	0.33
	Reject-All Rate	0.17	0.01	0.29	0.14	0.00	0.26

Number of observations: 653

Notes: Summary statistics for two groups of users, those who experience a pop-up every 10 minutes of organic browsing, and those who experience a pop-up every 60 minutes. "Unique domains in prior week" refers to the number of unique domains visited 7 days before the experiment; we keep the length of the historical visit the same as the post-survey visit for comparability. Number of observations: 656

TABLE A3. Covariate Balance Check for Dark Pattern Randomization

	Age (1)	Female (2)	Bachelor's or Above (3)	Domain Rank (Log 10) (4)
Constant	38.720*** (0.196)	0.438*** (0.007)	0.182*** (0.006)	3.579*** (0.023)
Acc-GrRej-GrSet	0.061 (0.281)	-0.004 (0.011)	-0.009 (0.008)	-0.034 (0.033)
Acc-Rej-Set	-0.040 (0.281)	0.013 (0.011)	0.005 (0.008)	-0.004 (0.033)
Acc-Set	-0.023 (0.280)	-0.001 (0.010)	-0.007 (0.008)	-0.023 (0.033)
Rej-Acc-Set	0.343 (0.284)	0.011 (0.011)	-0.008 (0.008)	0.013 (0.033)
Rej-Set	0.312 (0.286)	-0.008 (0.011)	-0.011 (0.008)	-0.045 (0.033)
R ²	0.000	0.000	0.000	0.000
Observations	26,278	26,278	26,773	26,773

Notes: * $p < 0.05$, ** $p < 0.01$; *** $p < 0.001$.

TABLE A4. Covariate Balance Check for Banner Frequency Randomization

	# Survey Domains (1)	Age (2)	Female (3)	Bachelor's or Above (4)
Constant	18.537*** (0.222)	37.460*** (0.703)	0.443*** (0.027)	0.563*** (0.027)
10 Min Pop-up	0.123 (0.307)	1.560 (1.037)	0.039 (0.040)	-0.076 (0.039)
R ²	0.000	0.004	0.001	0.006
Observations	656	638	638	656

Notes: * $p < 0.05$, ** $p < 0.01$; *** $p < 0.001$.

TABLE A5. Choices of “Accept Some” Cookies by Experimental Condition

	Survey (1)	Organic (2)
Acc-Set	0.054*** (0.010)	0.023** (0.008)
Rej-Set	0.037*** (0.007)	0.036*** (0.007)
Acc-Rej-Set	-0.003 (0.004)	0.004 (0.003)
Rej-Acc-Set	-0.005 (0.005)	-0.002 (0.004)
Acc-GrRej-GrSet	-0.006 (0.004)	0.004 (0.003)
Benchmark group mean:	0.04	0.04
R ²	0.408	0.491
Observations	12,142	14,163
Participant fixed effects	✓	✓
Domain Cat. fixed effects	✓	✓

Notes: The table regressions of Equation 1, where the outcome is whether the user selects a subset of cookies. Otherwise the table is identical to Table 2. * $p < 0.05$, ** $p < 0.01$; *** $p < 0.001$.

TABLE A6. Choices of “Accept Some” Cookies by Dark Pattern

	Survey	Organic
	Accept Some	
	(1)	(2)
Reject Hidden	0.054*** (0.010)	0.023** (0.008)
Accept Hidden	0.041*** (0.007)	0.039*** (0.007)
Accept Top	-0.003 (0.004)	0.004 (0.003)
Reject Top	-0.005 (0.005)	-0.002 (0.004)
Highlight Accept	-0.003 (0.004)	-0.001 (0.003)
Benchmark group mean:	0.04	0.04
R ²	0.408	0.491
Observations	12,142	14,163
Participant fixed effects	✓	✓
Domain Cat. fixed effects	✓	✓

Notes: The table regressions of Equation 2, where the outcome is whether the user selects a subset of cookies. Otherwise the table is identical to Table 3. * $p < 0.05$, ** $p < 0.01$; *** $p < 0.001$.

TABLE A7. Types of Cookies Accepted among People Making Granular Choices

Cookie Type	Percentage Selected
Preferences and functionality	0.831
Information storage and access	0.630
Performance and analytics	0.600
Content selection, delivery, and reporting	0.388
Ad selection, delivery, and reporting	0.068
Other purposes	0.046

Notes: Percentage of different types of cookies selected among those who selectively accept some cookies but not all.

TABLE A8. Heterogeneity of Dark Pattern Effect by Domain Ranking (Survey Sites)

	Accept All (1)	Reject All (2)	Close Window (3)
Domain Rank (Log 10)	-0.020*** (0.005)	0.015** (0.005)	0.004 (0.003)
Domain Rank (Log 10) × Reject Hidden	0.007 (0.006)	-0.013* (0.006)	0.004 (0.004)
Domain Rank (Log 10) × Accept Hidden	0.024*** (0.007)	-0.018* (0.007)	0.002 (0.006)
Domain Rank (Log 10) × Accept Top	-0.001 (0.006)	0.006 (0.006)	-0.006 (0.004)
Domain Rank (Log 10) × Reject Top	-0.004 (0.006)	0.006 (0.006)	-0.004 (0.003)
Domain Rank (Log 10) × Highlight Accept	-0.004 (0.006)	-0.002 (0.006)	0.006 (0.004)
Reject Hidden	0.059* (0.023)	-0.131*** (0.022)	0.027 (0.015)
Accept Hidden	-0.544*** (0.030)	0.243*** (0.030)	0.235*** (0.025)
Accept Top	0.022 (0.021)	-0.026 (0.019)	0.008 (0.014)
Reject Top	0.019 (0.021)	-0.007 (0.020)	0.001 (0.012)
Highlight Accept	0.031 (0.021)	-0.006 (0.018)	-0.022 (0.014)
R ²	0.655	0.583	0.573
Observations	12,142	12,142	12,142
Participant fixed effects	✓	✓	✓
Domain Cat. fixed effects	✓	✓	✓

Notes: * $p < 0.05$, ** $p < 0.01$; *** $p < 0.001$.

TABLE A9. Heterogeneity of Dark Pattern Effect by Visit Frequency (Survey Sites)

	Accept All (1)	Reject All (2)	Close Window (3)
Normally Visit	0.103*** (0.015)	-0.091*** (0.015)	-0.015 (0.009)
Normally Visit × Reject Hidden	-0.046* (0.020)	0.070*** (0.019)	-0.014 (0.014)
Normally Visit × Accept Hidden	-0.081** (0.025)	0.010 (0.024)	0.020 (0.021)
Normally Visit × Accept Top	-0.015 (0.018)	-0.001 (0.016)	0.006 (0.012)
Normally Visit × Reject Top	-0.006 (0.018)	0.006 (0.017)	-0.003 (0.011)
Normally Visit × Highlight Accept	0.007 (0.020)	0.006 (0.018)	-0.006 (0.011)
Reject Hidden	0.095*** (0.014)	-0.195*** (0.016)	0.043*** (0.011)
Accept Hidden	-0.441*** (0.021)	0.184*** (0.017)	0.234*** (0.018)
Accept Top	0.024* (0.012)	-0.006 (0.010)	-0.012 (0.007)
Reject Top	0.007 (0.011)	0.011 (0.011)	-0.011 (0.007)
Highlight Accept	0.015 (0.012)	-0.015 (0.011)	0.000 (0.006)
R ²	0.658	0.587	0.573
Observations	12,142	12,142	12,142
Participant fixed effects	✓	✓	✓
Domain Cat. fixed effects	✓	✓	✓

Notes: “Normally visit” is a binary variable constructed from our question in the survey phase: “Do you normally visit website [X]?” and is available only for the 20 surveyed sites. * $p < 0.05$, ** $p < 0.01$; *** $p < 0.001$.

TABLE A10. Heterogeneity of Dark Pattern Effect by Prior Site Visits (Survey Sites)

	Accept All (1)	Reject All (2)	Close Window (3)
Has Prior Visit	0.043* (0.020)	-0.032 (0.018)	-0.014 (0.013)
Has Prior Visit × Reject Hidden	-0.029 (0.028)	0.052* (0.026)	-0.014 (0.017)
Has Prior Visit × Accept Hidden	-0.077* (0.032)	-0.011 (0.031)	0.049 (0.029)
Has Prior Visit × Accept Top	0.012 (0.028)	-0.015 (0.025)	-0.005 (0.021)
Has Prior Visit × Reject Top	0.021 (0.027)	-0.022 (0.025)	0.015 (0.018)
Has Prior Visit × Highlight Accept	0.011 (0.028)	-0.005 (0.025)	0.004 (0.020)
Reject Hidden	0.082*** (0.012)	-0.178*** (0.014)	0.041*** (0.009)
Accept Hidden	-0.458*** (0.019)	0.189*** (0.016)	0.234*** (0.017)
Accept Top	0.018 (0.010)	-0.005 (0.009)	-0.009 (0.006)
Reject Top	0.002 (0.010)	0.015 (0.010)	-0.014* (0.006)
Highlight Accept	0.016 (0.010)	-0.012 (0.010)	-0.003 (0.006)
R ²	0.654	0.583	0.574
Observations	12,142	12,142	12,142
Participant fixed effects	✓	✓	✓
Domain Cat. fixed effects	✓	✓	✓

Notes: “Has Prior Visit” is a binary variable indicating whether a domain has been visited by the participant two weeks prior to our experiment. * $p < 0.05$, ** $p < 0.01$; *** $p < 0.001$.

TABLE A11. Heterogeneity of Dark Pattern Effect by Domain Ranking (Organic Sites)

	Accept All (1)	Reject All (2)	Close Window (3)
Domain Rank (Log 10)	0.002 (0.005)	0.002 (0.003)	-0.004 (0.004)
Domain Rank (Log 10) × Reject Hidden	-0.004 (0.007)	0.000 (0.005)	0.004 (0.006)
Domain Rank (Log 10) × Accept Hidden	-0.009 (0.008)	0.004 (0.006)	0.009 (0.008)
Domain Rank (Log 10) × Accept Top	0.004 (0.006)	-0.005 (0.004)	0.000 (0.006)
Domain Rank (Log 10) × Reject Top	0.001 (0.006)	-0.004 (0.005)	0.001 (0.006)
Domain Rank (Log 10) × Highlight Accept	0.004 (0.007)	-0.006 (0.005)	0.002 (0.006)
Reject Hidden	0.073* (0.029)	-0.093*** (0.022)	0.002 (0.028)
Accept Hidden	-0.399*** (0.035)	0.056* (0.027)	0.290*** (0.035)
Accept Top	0.019 (0.027)	0.020 (0.019)	-0.042 (0.024)
Reject Top	0.001 (0.026)	0.039 (0.021)	-0.031 (0.026)
Highlight Accept	-0.019 (0.029)	0.004 (0.020)	0.014 (0.026)
R ²	0.580	0.522	0.510
Observations	14,163	14,163	14,163
Participant fixed effects	✓	✓	✓
Domain Cat. fixed effects	✓	✓	✓

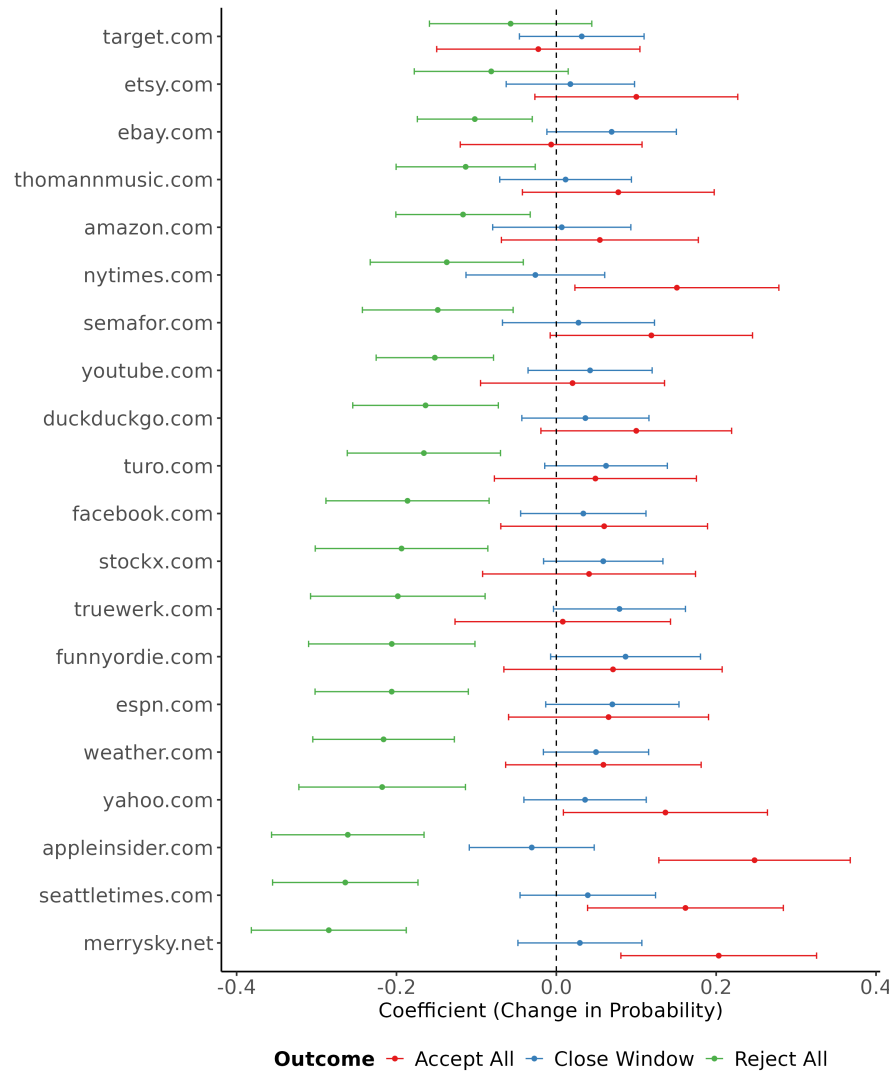
Notes: * $p < 0.05$, ** $p < 0.01$; *** $p < 0.001$.

TABLE A12. Heterogeneity of Dark Pattern Effect by Prior Site Visits (Organic Sites)

	Accept All (1)	Reject All (2)	Close Window (3)
Has Prior Visit	0.023 (0.017)	0.002 (0.012)	-0.029 (0.015)
Has Prior Visit × Reject Hidden	0.018 (0.021)	-0.001 (0.014)	-0.016 (0.019)
Has Prior Visit × Accept Hidden	-0.007 (0.026)	0.008 (0.018)	-0.025 (0.025)
Has Prior Visit × Accept Top	0.017 (0.019)	-0.002 (0.013)	-0.005 (0.018)
Has Prior Visit × Reject Top	0.023 (0.022)	-0.005 (0.015)	-0.004 (0.020)
Has Prior Visit × Highlight Accept	-0.010 (0.018)	0.002 (0.013)	0.006 (0.018)
Reject Hidden	0.049*** (0.013)	-0.094*** (0.014)	0.021 (0.014)
Accept Hidden	-0.434*** (0.025)	0.069*** (0.014)	0.335*** (0.024)
Accept Top	0.029* (0.012)	0.003 (0.008)	-0.040*** (0.011)
Reject Top	-0.002 (0.014)	0.024* (0.010)	-0.025* (0.012)
Highlight Accept	0.001 (0.012)	-0.020* (0.009)	0.020 (0.012)
R ²	0.581	0.522	0.512
Observations	14,163	14,163	14,163
Participant fixed effects	✓	✓	✓
Domain Cat. fixed effects	✓	✓	✓

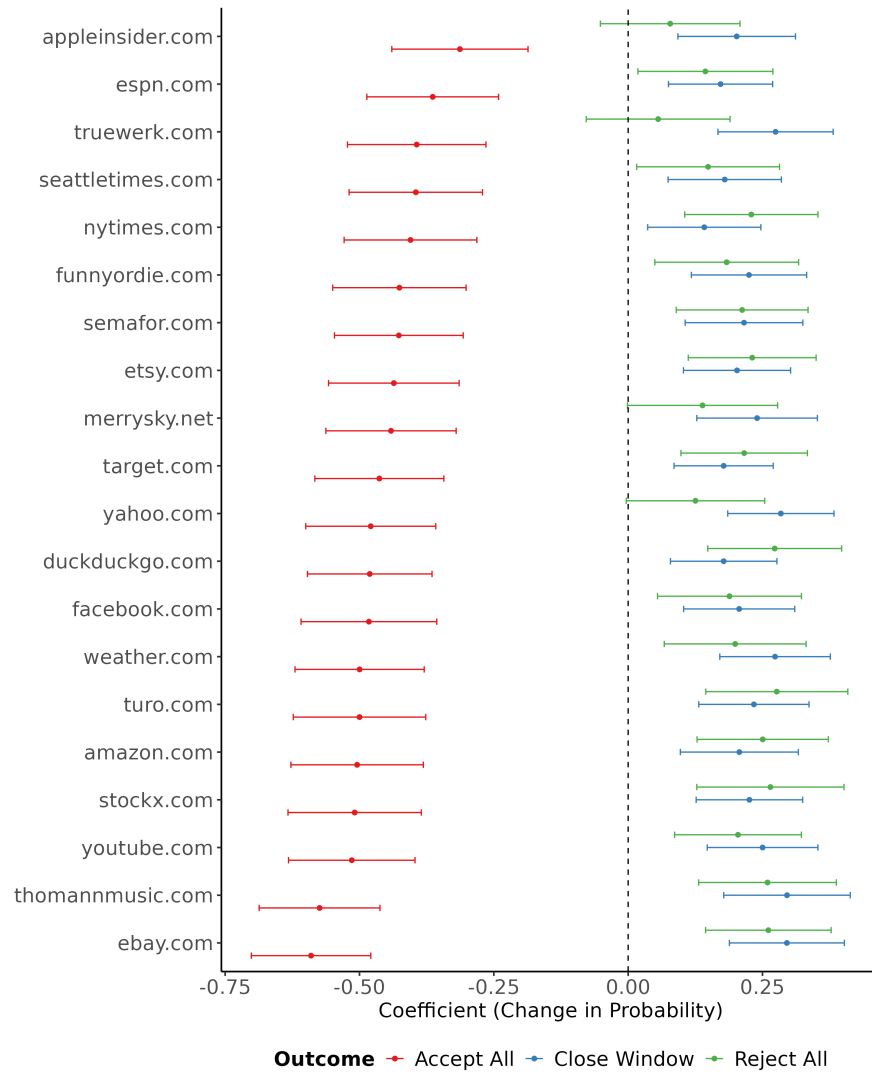
Notes: “Has Prior Visit” is a binary variable indicating whether a domain has been visited by the participant two weeks prior to our experiment. * $p < 0.05$, ** $p < 0.01$; *** $p < 0.001$.

FIGURE A2. Treatment Effects by Survey Domain (Acc-Set vs Neutral)



Notes: The figure shows the treatment effects (point estimates and 95% confidence intervals) of 'Acc-Set' vs the no dark pattern interface for each domain featured in the survey visit.

FIGURE A3. Treatment Effects by Survey Domain (Rej-Set vs Neutral)



Notes: The figure shows the treatment effects (point estimates and 95% confidence intervals) of 'Rej-Set' vs the no dark pattern interface for each domain featured in the survey visit.

TABLE A13. Demographic Differences Between Participants with High vs. Low Acceptance Rates

Variable	Mean: Low Accept	Mean: High Accept	Mean Diff.	p-value
Age	37.62	38.66	1.041	0.316
Female	0.5	0.43	-0.071	0.073
Bachelor's or Above	0.19	0.17	-0.015	0.613
Income > \$75,000	0.43	0.44	0.013	0.741

Notes: "Low Accept" and "High Accept" indicate the participants with their probability of choosing "accept all" lower and higher than the median, respectively.

TABLE A14. Fatigue in Cookie Choices During Organic Browsing (Additional Fixed Effects)

	Accept All (1)	Reject All (2)	Close Window (3)
Visit Order / 10	-0.017* (0.008)	-0.002 (0.003)	0.020* (0.008)
Time in Study (Days)	-0.001 (0.005)	-0.004 (0.002)	0.005 (0.005)
R ²	0.476	0.499	0.454
Observations	14,163	14,163	14,163
Domain Cat. fixed effects	✓	✓	✓
Participant fixed effects	✓	✓	✓
Hour fixed effects	✓	✓	✓

Notes: This table estimates a variant of Equation 3, which removes the banner frequency treatment and adds the order of which a domain is visited ("Visit Order") and additional fixed effects. * $p < 0.05$, ** $p < 0.01$; *** $p < 0.001$.

Appendix B. Survey Questions

This appendix presents the Qualtrics surveys used in the study:

- Intake.
- Outtake.

Device Transfer

The rest of the survey needs to be done on a Chrome browser. Please copy the link of the survey and reopen it in a Chrome browser to continue.

First Page

Would you like to help us understand online behavior and privacy choices? We are a team of Harvard and Boston University researchers who study the internet and how it affects society.

The study will take 30 minutes of your time over the course of the next day. We will ask you to fill out two surveys, clear the cookie data stored in your browser, install a browser extension vetted by Harvard and Boston University and keep it installed for seven days, and share information about your online behavior. Click below if you want to know more and discover if you qualify!

Eligibility Questions

Do you live in the United States?

No

Yes

Are you over 18 years old?

Yes

No

What is 12 minus 4? Regardless of the correct answer, you should always select the option with the value "seven". This is an attention check question.

6

8

7

5

What is the language you primarily speak?

Spanish

English

Other (please specify)

Which browser do you primarily use?

Others

Internet Explorer

Chrome

Microsoft Edge

Safari

Firefox

What was your total household income before taxes during the past 12 months?

Less than \$25,000

\$25,000-\$49,999

\$50,000-\$74,999

\$75,000-\$99,999

\$100,000-\$149,999

\$150,000 or more

Prefer not to say

What is the highest level of education you have completed?

Some high school or less

High school diploma or GED

Some college, but no degree

Associates or technical degree

Bachelor's degree

Graduate or professional degree (MA, MS, MBA, PhD, JD, MD, DDS etc.)

Prefer not to say

Not Eligible

Thank you for your answers! Unfortunately, you do not qualify to participate in our study.

Can you please return your submission on Prolific?

Consent

Congratulations! You are qualified to participate in our study.

Study Overview

The following is a summary with key information to help you decide whether you want to participate.

Why am I being invited to take part in a research study?

We invite you to take part in this research study because you are an English-speaking resident of the United States who uses Chrome to browse the web.

What should I know about a research study?

Research studies are conducted to better understand the choices we make. Whether or not you take part is completely up to you. Your decision will not be held against you. You can ask all the questions you want before you decide. You can even agree to take part and later change your mind.

Why is this research being done?

We want to better understand the online experience of people like you, how companies obtain user consent for the collection and use of their data, and how this affects user browsing experience. We hope that the results of this research will help inform data privacy policy.

How long will the research last and what will I need to do?

The study will last several days, but we will only ask you for 30 minutes of your time. Everything we ask you to do to participate in this research can be done from the comfort of your home. If you choose to participate, we'll ask you to:

- Complete two surveys:
 - The first survey will ask you some questions about yourself and your online browsing behavior. It will also ask you to visit some websites and make privacy choices.
 - The second survey will ask you about your experience during the study.
- Install the Cookie Manager browser extension, which is an application we developed for this study. We'll have instructions for you. The Cookie Manager extension will record your behavior and may tweak the interfaces through which you make cookie selections.
- Keep the extension installed for seven days, until the extension prompts you to uninstall it.

Will I be compensated for participating in this research?

Yes. You will be paid \$7.50 after completing the two surveys and keeping the Cookie Manager extension installed for several days.

Is there any way being in this study could be bad for me?

Since we may collect personal information, there is a risk of breach of confidentiality. We have worked hard to minimize this risk. For example, we will encrypt any data before storing it. Before accessing the data for analysis, we will also permanently delete all personal information that we may intentionally or unintentionally collect.

Will being in this study help me in any way?

We cannot promise any benefits to you or others from your taking part in this research. It

is possible, however, that our tweaks to your online browsing lead to a better (or worse) online experience.

Detailed Information

Withdrawing from the Study.

You can leave the research at any time; your decision will not be held against you. We may use the data you have shared with us prior to withdrawing as part of the study. We will provide simple instructions for how you can withdraw. Researchers can remove you from the research study without your approval. Possible reasons for removal include not complying with instructions to install the browser extension or intentionally avoiding data tracking through the extension.

Privacy.

Data security and privacy are important to us. During the course of the study we may collect personal information. The personal information that we know we are collecting will be deleted immediately. Other personal information that we inadvertently collect will be stored but removed after we finish collecting data.

We cannot promise complete secrecy, although efforts will be made to limit the use and disclosure of your personal information. Data will be encrypted and stored on secure servers and cannot be accessed by anyone outside the research team. At no time will study information be available over any public or private network in an unencrypted state.

In the future, when we publish our research, we will post anonymized data from this study in a data repository so that other researchers can reproduce our results. By then, no information that can identify you personally will be available, to us or others. We will not sell data from the study or share data for any commercial or marketing purposes.

Who can I talk to?

If you have questions, concerns, or complaints, or think the research has hurt you, do not hesitate to reach the research team on Prolific or cookie.manager.study@gmail.com.

Please indicate below whether you agree to participate in the study. Agreeing to participate means you are willing to install Cookie Manager (our web browser extension) for seven days, and complete the two surveys.

I agree to participate

I do not agree to participate

Not consent

Thank you for letting us know you do not want to participate. **Can you please return your submission on Prolific?**

Email

Thank you for your willingness to participate in our study!

Next, we will ask you to install *Cookie Manager*, a browser extension we developed to identify website tracking and to enable simplified privacy consent dialogs.

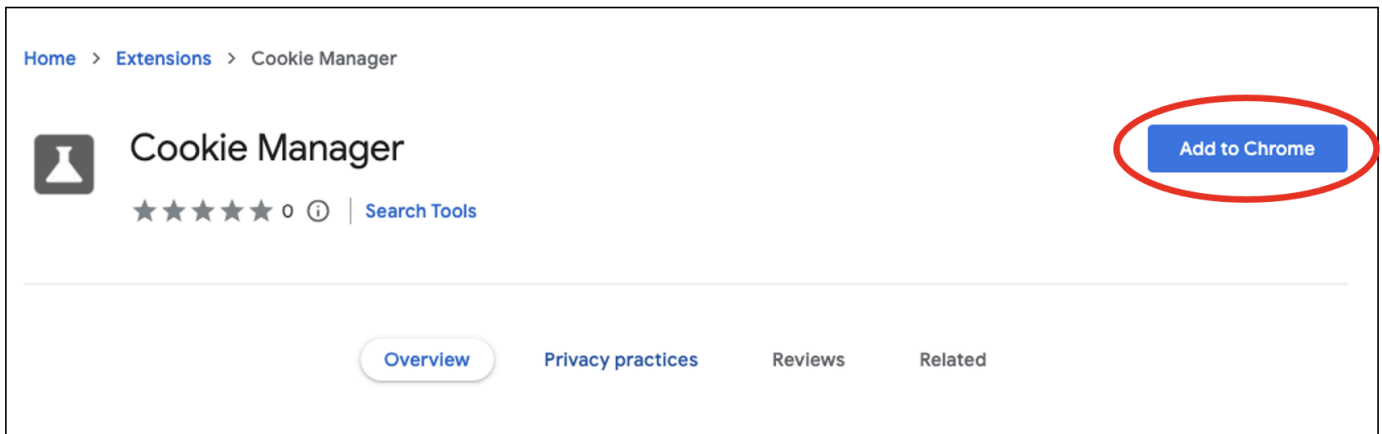
After installing the extension, you will see a consent-request popup window whenever you visit a website for the first time. If you make a choice, the extension will try to pass on your choices to the website. In most cases, if the website has already been collecting consent from users, it will recognize your choice and decide whether to continue tracking you based on your choice.

App Installation

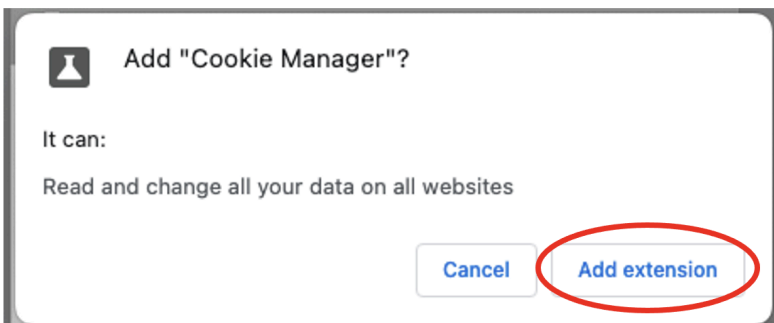
Cookie Manager Installation Instructions.

To install Cookie Manager, please **use Chrome** on the computer that you are using for online shopping:

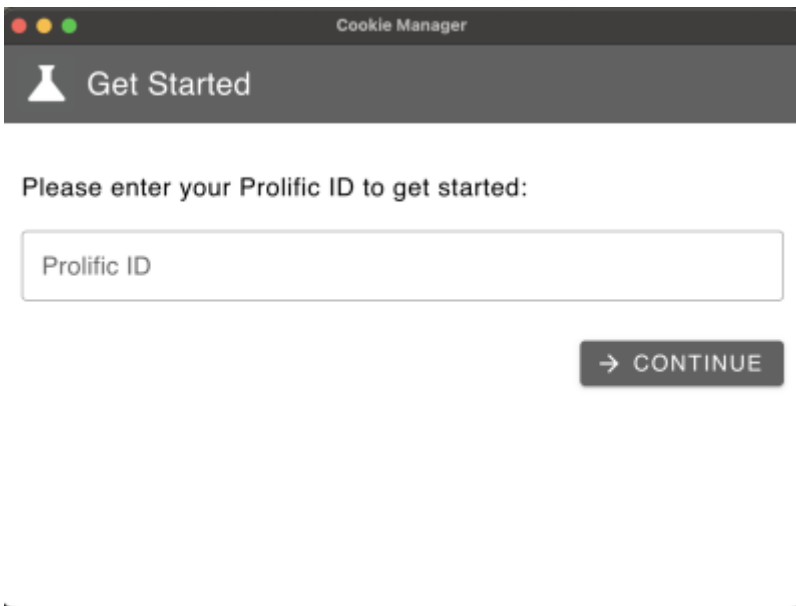
- Click [here](#).
- Click “Add to Chrome.”



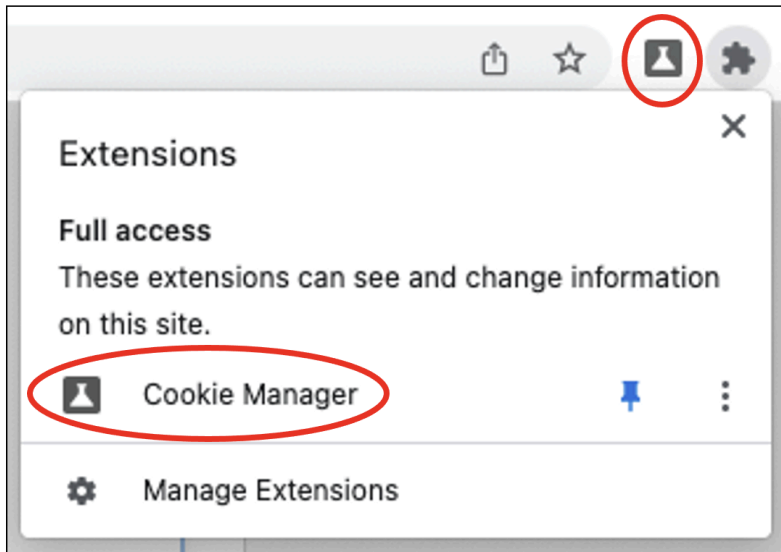
- When prompted, click “Add Extension.”



- You will be prompted to add your prolific id.



- You should now see the Cookie Manager icon on the top right corner of your browser. If you don't see it, it may be hidden under the puzzle icon, which is visible in the upper right corner of the screenshot below.



- You are all set.

If you have trouble installing Cookie Manager, email us at cookie.manager.study@gmail.com and we will help you with additional instructions.

Were you able to successfully install the extension?

Yes

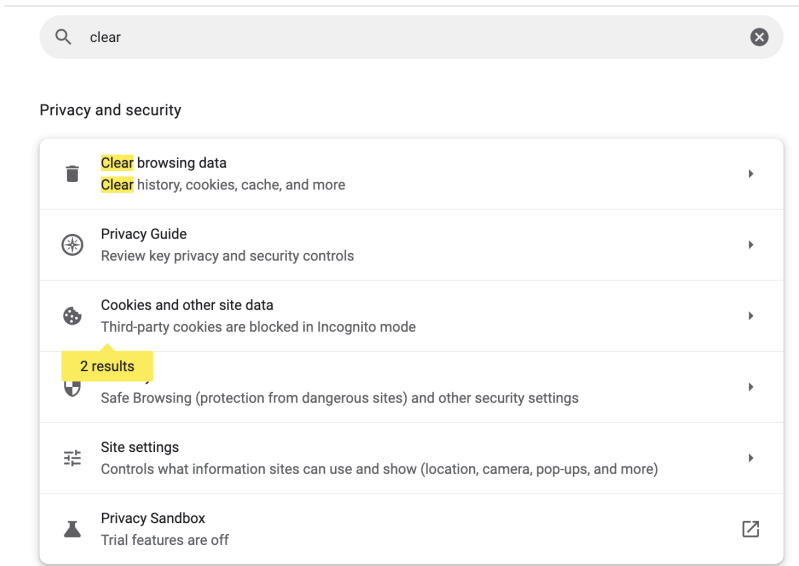
No

What difficulties have you encountered when installing the extension?

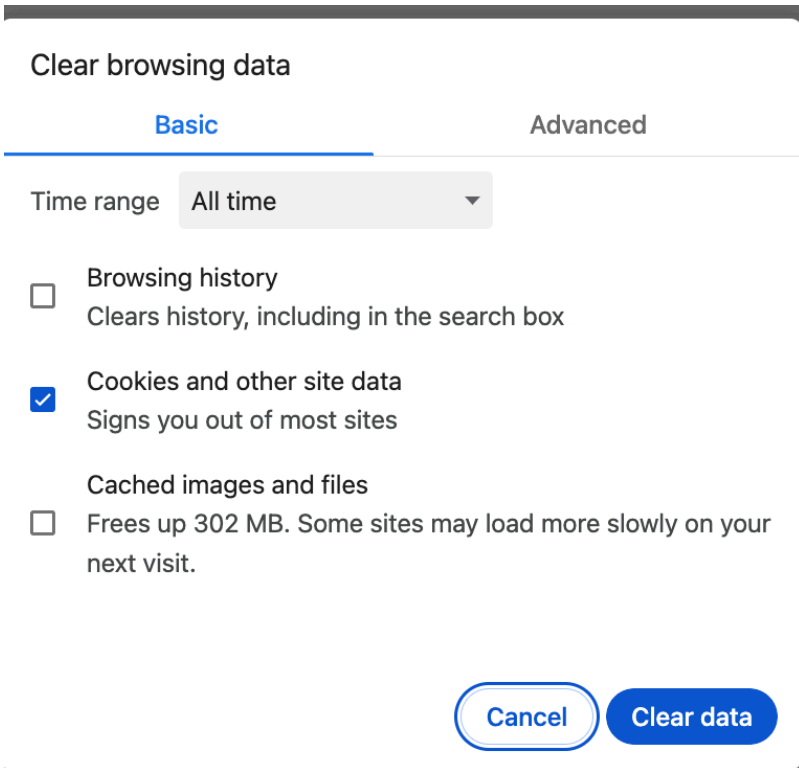
Clear Browsing History

Next, we will ask you to clear your cookie data. Please navigate to <chrome://settings/privacy?search=clear> (copy and paste the address directly on your

search bar), and click on "Clear Browsing Data". Then select **only 'cookies and other site data'**, and click clear data.



Select the time range to be "**All Time**" and select the cookies and other data check box, as seen below. Then click "Clear Data."



Were you able to clear your cookie data?

Yes

No

Intro to website navigation

Now that you have *Cookie Manager* installed, we will ask you to visit a list of 20 websites. Please wait until a banner shows up for each site and interact with the banner as you normally would. We will ask you to answer a few questions after each visit. After you finish the survey task, the frequency of pop-ups will drastically decrease.

Note: for your browsing action to be correctly registered in our database, please directly left-click on the link on the survey page to navigate to the website. If instead you right-click on the link and select "open on a new tab", a warning will continue showing up, meaning that our database has not recognized your click action.

YouTube

Please use Chrome to navigate to [youtube.com](https://www.youtube.com). Please wait until a banner shows up. Search for a video of your choice.

You haven't clicked on the link

Do you normally visit Youtube?

Yes

No

Have you ever heard of Youtube?

Yes

No

How often do you normally visit Youtube?

At least once a day

At least once a week

Less than once a week

Never

New York Times

Please use Chrome to navigate to nytimes.com. Please wait until a banner shows up. Click on an article of your choice.

You haven't clicked on the link

Do you normally visit NYTimes?

Yes

No

Have you ever heard of New York Times?

Yes

No

How often do you normally visit New York Times?

At least once a day

At least once a week

Less than once a week

Never

Apple Insider

Please use Chrome to navigate to appleinsider.com. Please wait until a banner shows up. Click on an article of your choice.

You haven't clicked on the link

Do you normally visit Apple Insider?

Yes

No

Have you ever heard of Apple Insider?

Yes

No

How often do you normally visit Apple Insider?

At least once a day

At least once a week

Less than once a week

Never

Yahoo

Please use Chrome to navigate to [yahoo.com](https://www.yahoo.com). Please wait until a banner shows up. Click on an article of your choice.

You haven't clicked on the link

Do you normally visit Yahoo?

Yes

No

Have you ever heard of Yahoo?

Yes

No

How often do you normally visit Yahoo?

At least once a day

At least once a week

Less than once a week

Never

Amazon

Please use Chrome to navigate to [amazon.com](https://www.amazon.com). Please wait until a banner shows up. Search for a product of your choice.

You haven't clicked on the link

Do you normally visit Amazon?

Yes

No

Have you ever heard of Amazon?

Yes

No

How often do you normally visit Amazon?

At least once a day

At least once a week

Less than once a week

Never

eBay

Please use Chrome to navigate to [ebay.com](https://www.ebay.com). Please wait until a banner shows up.
Search for a product of your choice.

You haven't clicked on the link

Do you normally visit eBay?

Yes

No

Have you ever heard of eBay?

Yes

No

How often do you normally visit eBay?

At least once a day

At least once a week

Less than once a week

Never

What is 6 divided by 2? Regardless of the correct answer, you should always select the option with the value "one". This is an attention check question.

2

3

1

Target

Please use Chrome to navigate to [target.com](https://www.target.com). Please wait until a banner shows up.
Search for a product of your choice.

You haven't clicked on the link

Do you normally visit Target?

Yes

No

Have you ever heard of Target?

Yes

No

How often do you normally visit Target?

At least once a day

At least once a week

Less than once a week

Never

Etsy

Please use Chrome to navigate to [etsy.com](https://www.etsy.com). Please wait until a banner shows up.
Search for a product of your choice.

You haven't clicked on the link

Do you normally visit Etsy?

Yes

No

Have you ever heard of Etsy?

Yes

No

How often do you normally visit Etsy?

At least once a day

At least once a week

Less than once a week

Never

Turo

Please use Chrome to navigate to turo.com. Please wait until a banner shows up. Click on a car of your choice.

You haven't clicked on the link

Do you normally visit Turo?

Yes

No

Have you ever heard of Turo?

Yes

No

How often do you normally visit Turo?

At least once a day

At least once a week

Less than once a week

Never

StockX

Please use Chrome to navigate to stockx.com. Please wait until a banner shows up. Search for a product of your choice.

You haven't clicked on the link

Do you normally visit StockX?

Yes

No

Have you ever heard of StockX?

Yes

No

How often do you normally visit StockX?

At least once a day

At least once a week

Less than once a week

Never

ESPN

Please use Chrome to navigate to espn.com. Please wait until a banner shows up. Click on an article of your choice.

You haven't clicked on the link

Do you normally visit ESPN?

Yes

No

Have you ever heard of ESPN?

Yes

No

How often do you normally visit ESPN?

At least once a day

At least once a week

Less than once a week

Never

Facebook

Please use Chrome to navigate to [facebook.com](https://www.facebook.com). Please wait until a banner shows up. Scroll down.

You haven't clicked on the link

Do you normally visit Facebook?

Yes

No

Have you ever heard of Facebook?

Yes

No

How often do you normally visit Facebook?

At least once a day

At least once a week

Less than once a week

Never

Funny Or Die

Please use Chrome to navigate to funnyordie.com. Please wait until a banner shows up. Click on an article of your choice.

You haven't clicked on the link

Do you normally visit Funny Or Die?

Yes

No

Have you ever heard of Funny Or Die?

Yes

No

How often do you normally visit Funny Or Die?

At least once a day

At least once a week

Less than once a week

Never

Weather

Please use Chrome to navigate to weather.com. Please wait until a banner shows up. Search for a location.

You haven't clicked on the link

Do you normally visit Weather.com?

Yes

No

Have you ever heard of Weather.com?

Yes

No

How often do you normally visit Weather.com?

At least once a day

At least once a week

Less than once a week

Never

DuckDuckGo

Please use Chrome to navigate to duckduckgo.com. Please wait until a banner shows up. Search for a product of your choice.

You haven't clicked on the link

Do you normally visit DuckDuckGo?

Yes

No

Have you ever heard of DuckDuckGo?

Yes

No

How often do you normally visit DuckDuckGo?

At least once a day

At least once a week

Less than once a week

Never

Truewerk

Please use Chrome to navigate to truewerk.com. Please wait until a banner shows up and navigate to an item.

You haven't clicked on the link

Do you normally visit Truewerk?

Yes

No

Have you ever heard of Truewerk?

Yes

No

How often do you normally visit Truewerk?

At least once a day

At least once a week

Less than once a week

Never

Thomann

Please use Chrome to navigate to thomannmusic.com. Please wait until a banner shows up and navigate to an item.

You haven't clicked on the link

Do you normally visit Thomann Music?

Yes

No

Have you ever heard of Thomann Music?

Yes

No

How often do you normally visit Thomann Music?

At least once a day

At least once a week

Less than once a week

Never

MerrySky

Please use Chrome to navigate to merrysky.com. Please wait until a banner shows up and search for a location.

You haven't clicked on the link

Do you normally visit Merry Sky?

Yes

No

Have you ever heard of Merry Sky?

Yes

No

How often do you normally visit Merry Sky?

At least once a day

At least once a week

Less than once a week

Never

Seattle Times

Please use Chrome to navigate to seattletimes.com. Please wait until a banner shows up and then click on an article.

You haven't clicked on the link

Do you normally visit Seattle Times?

Yes

No

Have you ever heard of Seattle Times?

Yes

No

How often do you normally visit Seattle Times?

At least once a day

At least once a week

Less than once a week

Never

Semafor

Please use Chrome to navigate to semafor.com. Please wait until a banner shows up and then click on an article.

You haven't clicked on the link

Do you normally visit Semafor?

Yes

No

Have you ever heard of Semafor?

Yes

No

How often do you normally visit Semafor?

At least once a day

At least once a week

Less than once a week

Never

Favorite website

Navigate to your favorite e-commerce website. Please wait until a banner shows up. Search for a product of your choice.

Paste the URL of the product below:

Did you see a cookie consent banner?

Yes

No

Did you make a choice on whether to allow for cookie sharing?

Yes, I allowed my preferred cookies and blocked unwanted cookies

Yes, I chose the default cookie sharing

No, I closed the cookie consent banner

No, I left the website without interacting with the consent banner

Questionnaire

Think about your browsing experiences on a typical day. Overall, how frequently do you encounter cookie consent banners?

Too frequently

A bit more frequently than ideal

Just right

A bit less frequently than ideal

Too infrequently

Overall, how would you rate the ease of navigation of the cookie consent interfaces on the websites you visit?

Very easy to navigate

Moderately easy to navigate

Neither easy nor hard to navigate

Moderately hard to navigate

Very hard to navigate

Overall, how would you rate the ease of making your preferred choices regarding cookie sharing on the websites you visit?

Very easy

Moderately easy

Neither easy nor hard

Moderately hard

Very hard

Which of the following best describes your behavior when deciding whether to share cookies online?

I reject most cookies

I consider both the website that is asking and the types of cookies involved before deciding whether to share them

I accept most cookies

I decide whether to share cookies based on what type of cookies they are

I decide whether to share cookies based on which website is asking

Part1-conclude

Thank you! To finish the rest of the study, we ask you to keep Cookie Manager installed for another seven days. You can continue your browsing activities as usual during this time. The frequency of pop-ups will drastically decrease over time. After the seven days have passed, the extension will prompt you with a survey and the instructions on how to uninstall the extension.

There is no completion code, since our system will detect completion automatically. Please make sure to click the next button below so that we register your response.

Powered by Qualtrics

Intro Page

Thank you for finishing our web browsing task! Now we will walk you through the uninstallation process of the browser extension. To complete the study, we just need to ask you a few more questions about the web browsing and cookie-sharing experiences while using our extension and in general.

Block 1

Think back about your browsing experiences after completing our 20-website visit task while Cookie Manager is installed. Overall, what do you think of the frequency with which cookie consent banners appear during that time?

Too frequent

A bit more frequent than ideal

Just right

A bit less frequent than ideal

Too infrequent

Block 2

Overall, how will you rate the ease of navigation of the cookie consent interface created by our browser extension?

Very easy to navigate

Moderately easy to navigate

Neither easy nor hard to navigate

Moderately hard to navigate

Very hard to navigate

Block 3

Overall, how will you rate the ease of making your preferred cookie sharing choices created by our browser extension?

Very easy

Moderately easy

Neither easy nor hard

Moderately hard

Very hard

Block 4

In the past week, which of the following statement best describes your behavior when deciding whether to share cookies online?

I accepted most cookies

I rejected most cookies

I chose whether to share cookies based on which website is asking

I chose whether to share cookies based on what types of cookie it is

I chose whether to share cookies based on what website is asking and what types of cookie it is

Why choice

Think back to a case when you accepted all cookies during the course of the study. Why did you do so?

Think back to a case when you chose **not** to accept all cookies during the course of the study. Why did you do so?

Block 5

Overall, how do you think the Cookie Manager extension changes your web browsing experience?

It improves my browsing experience by a lot

It improves my browsing experience slightly

It neither improves nor degrades my browsing experience

It degrades my browsing experience slightly

It degrades my browsing experience a lot

Block 12

Consider the cookie consent form below.



One option is to hit the 'x' button in the upper right. If you were to click this 'x', what do you think will happen?

All cookies are accepted.

None of the cookies are accepted.

Other, please explain:

Block 8

During the study period, did you take any actions to change how you browse the internet?

No

Yes, I used a different browser or device.

Yes, I browsed the internet less.

Yes, I did something else. Please specify.

Block 9

As you browse the internet, which information do you think advertisers have about you?
Check all that apply.

Your demographic information

Your prior website visits

Your interests

Your prior purchases

Your social media posts

Your address

Your credit score

Block 10

Thinking about privacy polices you might come across online or on your smartphone.
Which of the following comes closer to your view, even if neither is exactly right?

Just something I have to get past in order to use a product or service.

A meaningful part of my decision to use a product or service.

Privacy means different things to different people today. In thinking about all of your online browsing, please state how important it is for you to be in control of who can get info about you.

Not all imporant	Not very imporant	Somewhat Important	Very Imporant
1	2	3	4
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

Block 11

What do you think are the benefits of sharing the data listed above with the advertisers?

What do you think are the costs of sharing the data listed above with the advertisers?

Block 6

Do you have any suggestions to help us improve the design of the Cookie Manager extension or the design of our study in general?

Block 14

Consider a tool that that allows you to specify how you would like to answer cookie consent questions online. This tool will then automatically hide all cookie pop-ups and answer them in they way you specified. For example, if you stated that you wanted to accept cookies for all websites, the tool would do so.

Please select how much better or worse the tool is than manually answering the cookie consent form for each website.

Much worse than manual choice	Worse	Somewhat worse	Similar	Somewhat better	Better	Much better than manual choice
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

How much would you be willing to pay for the tool?

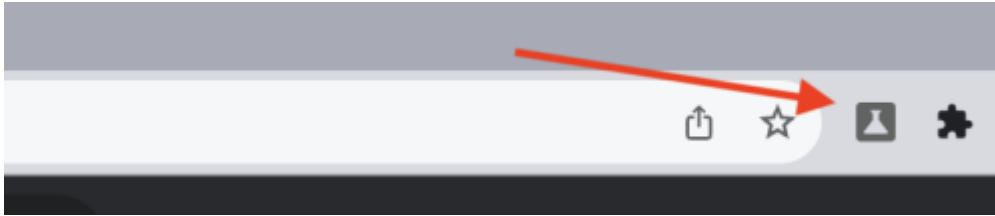
Please enter the price in the text box below.

Instructions for how to download and configure the tool, called Consent-O-Matic, are available [here](#).

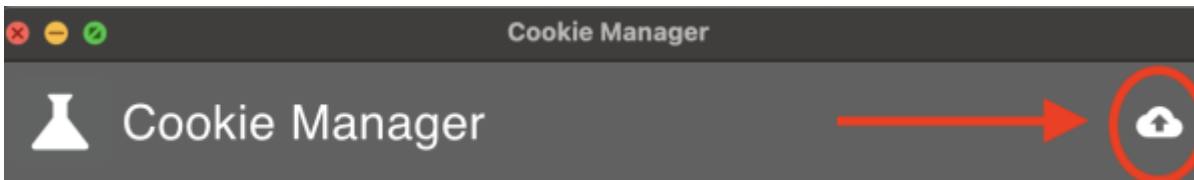
Please click the arrow below to continue the survey.

Block 7

Thank you! We will ask you to upload your data to us prior to uninstalling the extension. Please click on the Cookie Manager extension icon in your Chrome browser.



You should see a pop-up. Please click on the cloud button with an arrow. Completing this step ensures that your participation in our study and the associated data are properly recorded.

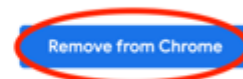


Now that you've clicked the cloud button, you can now proceed to uninstall the extension. Completing this step ensures that we stop collecting your browsing data going forward.

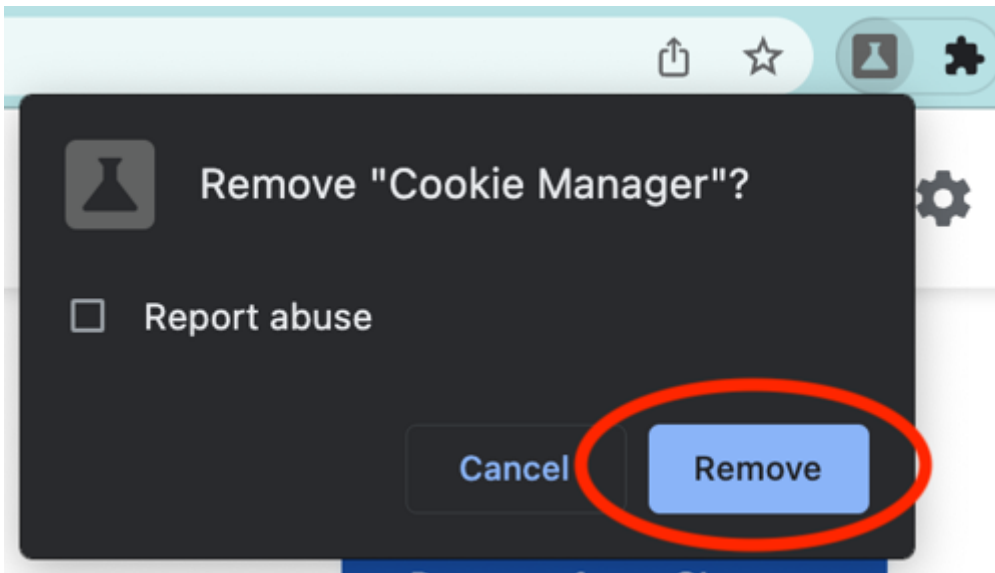
Here's how you can remove the cookie manager extension:

- Click [here](#).
- Click "Remove from Chrome."

[Home](#) > [Extensions](#) > [Cookie Manager](#)



- Confirm by clicking "Remove" on the pop-up window appearing on the top right corner of your browser.



- You're all set.

If you have trouble uninstalling Cookie Manager, email us at cookie.manager.study@gmail.com and we will help you with additional instructions.

Please click the arrow below to finish the survey.

Powered by Qualtrics