

The Economics of Data Sharing: Consumer Preferences and the Role of “Dark Patterns”

Chiara Farronato, Andrey Fradkin, Tesary Lin

June 4, 2025

Preliminary and Incomplete

Abstract

Regulations mandating consumer consent for data collection and use have led firms to employ “dark patterns”—interface designs that encourage data sharing. We study the causal effects of these designs on consumer consent choices and explore their implications for consumer welfare. We conduct a field experiment where a browser extension randomizes consent interfaces while study participants browse the internet. We find that consumers accept all cookies over half of the time absent dark patterns, with substantial user heterogeneity. Hiding options behind an extra click significantly sways choices toward visible options, while purely visual manipulations have smaller impacts. Users also frequently close banners without active selection, highlighting the importance of default settings. While better-known firms achieve higher baseline consent rates, dark patterns do not amplify this advantage. We use a structural model to show that the consumer surplus-maximizing banner removes dark patterns and defaults to accepting cookies upon consumer inaction. A browser-level global consent choice further improves consumer welfare compared to site-by-site consent interactions, primarily by reducing time costs.

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 and Hayden Schrauff for excellent research assistance, and Audacious Software for software development. This work received generous support from the Internet Society Foundation and Boston University’s Digital Business Institute.

1. Introduction

Consumer data is a key input for pricing, algorithmic recommendations, and targeted advertising. Yet, the collection and use of this data have sparked privacy concerns, spurring regulatory and societal pressure for firms to grant consumers more control. In response, firms now regularly solicit consent for data use, most commonly online through cookie consent banners.

Data use consent practices are controversial for at least three reasons. First, many firms design consent banners to favor data sharing, potentially preventing consumers from making their preferred choice. These choice architectures, often labeled *dark patterns* in public discourse,¹ include deliberate obstruction (e.g., hiding rejection options behind extra clicks), reordering options to favor sharing, or using differential visual salience. Second, requiring each firm to obtain user consent may lead to frequent consent banners, disrupting user experience and increasing the cognitive burden of repeated choices. Lastly, consent requirements risk benefiting large and prominent firms due to consumers’ higher propensity to share data with familiar brands.²

In this paper, we study the role of choice architecture in online cookie consent decisions and its effect on consumer welfare. We conduct a field experiment that randomly assigns one of six consent interfaces across users and web domains. We use this experimental variation to measure the causal effects of choice architecture. We also estimate a structural model of consumer choices to share data, allowing us to compare welfare across a variety of policies, including those that ban dark patterns and allow for browser-level, rather than site by site, data sharing choices.

We find that deliberate obstruction is the most effective design in swaying choices towards readily visible options, whereas purely visual manipulations like reordering or highlighting with different colors have minimal impact. Prominent websites generally achieve higher consent rates, but the effect of dark patterns does not systematically differ across websites, which does not substantially exacerbate incumbent advantages in data collection. Users exhibit considerable heterogeneity in their baseline data-sharing preferences, and a significant fraction frequently close consent banners without making an active selection, underscoring the critical role of website default settings.

¹<https://www.deceptive.design/>

²For example, the Digital Markets Act requires “gatekeeper” platforms to share data with smaller players 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. The FTC’s report to OECD also discusses how data privacy and competition interact: [https://one.oecd.org/document/DAF/COMP/WD\(2024\)29/en/pdf](https://one.oecd.org/document/DAF/COMP/WD(2024)29/en/pdf)

Our welfare analysis indicates that the consumer surplus-maximizing consent interface is one without any dark patterns, combined with an “accept all cookies” default for users who close the banner without choosing. This optimal banner design substantially increases welfare compared to the most common design in the US, which hides rejection options while defaulting users to data sharing. It also dominates the EU design, where the default tends to be no sharing. Finally, we find that a browser-level global consent mechanism can yield even larger welfare gains by eliminating repeated site-specific interactions and their associated time costs.

Our study is of particular relevance to regulators, who are increasingly concerned about choice architecture. For example, in the EU, both the Digital Services Act and the Artificial Intelligence Act regulate dark patterns.³ In the U.S., several states and the Federal Trade Commission have taken action against deceptive interface designs,⁴ though there is no federal standard. Prior audits (Habib et al. 2022; Bielova, Santos, and Gray 2024) show that dark patterns—particularly those involving hidden options or visual manipulation—are present on the vast majority of websites requiring consent. These patterns raise regulatory concerns because they can undermine meaningful user choice.

Our experiment is enabled by Cookie Manager, a customized browser extension we developed in order to assign and randomize cookie consent interfaces for users as they browse the internet. Cookie Manager is based on the Webmunk extension framework (Farronato, Fradkin, and Karr 2024), and enforces users’ cookie consent choices whenever easily implementable, making user choices incentive compatible.

In the absence of dark patterns, consent banners should display three clear options: accept all cookie tracking, reject all (except for essential cookies), and customize by use. Our browser extension displays one of six different consent interfaces, which vary in their use of three types of dark patterns: deliberate obstruction (i.e., removing options such as ‘reject all’ from the main banner), reordering options to prioritize those with more data sharing, and highlighting the option to share all data with a different color. We randomize these dark patterns across users and web domains.

³[https://www.europarl.europa.eu/RegData/etudes/ATAG/2025/767191/EPRS_ATA\(2025\)767191_EN.pdf](https://www.europarl.europa.eu/RegData/etudes/ATAG/2025/767191/EPRS_ATA(2025)767191_EN.pdf)

⁴For example, the Federal Trade Commission has fined Epic Games and Amazon for user interface designs that induce accidental purchases and obstruct cancellation of subscriptions (<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> and <https://www.ftc.gov/news-events/news/press-releases/2023/06/ftc-takes-action-against-amazon-enrolling-consumers-amazon-prime-without-consent-sabotaging-their>). US States such as California, Colorado, and Connecticut have enacted privacy regulations that explicitly ban companies from using dark patterns to increase data collection. (https://insightplus.bakermckenzie.com/bm/technology-media-telecommunications_1/united-states-consumer-protection-regulators-set-sights-on-dark-patterns).)

⁴The extension looks for common HTML, CSS, and textual patterns indicating buttons to accept or reject cookies. It then selects those buttons if detected.

We use Prolific to recruit US consumers who consent to install Cookie Manager. The study consists of two phases. In the first phase (*survey browsing*), we prompt participants to visit specific websites. This structured browsing allows us to evaluate their privacy preferences across the same set of websites, regardless of whether they would organically visit them. In the second phase (*organic browsing*), we observe participants' natural browsing behavior for a week following the survey phase. As before, we randomize the design of the consent pop-ups at the website and participant level.⁵ Together, these two phases allow us to identify data sharing preferences and dark pattern effects across websites, and characterize how survey-based results map to field-based choices.

When participants encounter a consent pop-up without any dark patterns (a *neutral interface*), 65% choose to “accept all cookies” during the survey phase, and 60% do so during the organic phase. While most participants consistently accept or reject cookies in the absence of dark patterns, 31% exhibit heterogeneous preferences across websites, showing a greater tendency to accept cookies on popular or familiar sites. In contrast, dark patterns significantly increase variability in consumer choices: 91% of participants change their privacy preferences across websites when dark patterns are present, compared to just 30% with a neutral interface.

We find that deliberate obstruction has the strongest influence on privacy choices, while dark patterns that feature pure visual manipulations have weaker effects. Hiding the “reject all cookies” button reduces the probability of rejecting cookies by 17.1% in survey visits and 9.4% in organic browsing. The sizable effect of deliberate obstruction is consistent with websites' strategic choices. As shown by Utz et al. (2019), deliberate construction is present in 78.5% of cookie banners, making it the most commonly used dark pattern. In comparison, presenting all options but reordering them so that the “accept all cookies” is displayed at the top only increases consent rates by 2-3.5%. Additionally, graying out options other than “accept all cookies” increases the acceptance probability by less than 2%.

Perhaps surprisingly, the effects of dark patterns do not vary substantially with website characteristics such as popularity or user familiarity. Absent dark patterns, consumers are more likely to consent to data sharing on popular or familiar websites during the survey phase. Dark patterns do not increase users' propensity to share data with popular or familiar websites during organic browsing, and if anything, they seem to alleviate such a tendency during the survey phase. These findings challenge the hypothesis that

⁵After the survey phase, participants are further randomized into two groups: one group experiences consent banners at most every 10 minutes (the “10-minute” frequency treatment), whereas another group experiences consent banners at most every 60 minutes (the “60-minute” treatment). This randomization, detailed in Appendix B, allows us to test for choice fatigue by examining whether the frequency of pop-ups influences user behavior. As discussed in the appendix, we find no significant effects of banner frequency on consent choices.

dark patterns increase entry barriers or amplify data-enabled network effects (Hagiu and Wright 2023), which would otherwise reinforce incumbent advantages in the data economy.

We also measure the time participants spend interacting with banners. On average during the organic browsing phase, a participant spends 7.4 seconds interacting with a neutral banner. Extrapolating this number to scenarios where the banners are present on every domain and assuming a value of time of \$36/hour (the average hourly wage in the US), we estimate the weekly cost of interacting with consent pop-ups to be between \$1.4 and \$2.7 per week per participant.⁶

We use the experimental variation in consent interfaces to estimate a structural model of consumer consent decisions. Our analysis relies primarily on revealed preferences, except in cases where users close the consent banner without making an active choice. In those cases, defaults matter, so we base choices on elicited consumer beliefs, but compute welfare according to the actual default.

Our results highlight that frictions in the consent interface, such as the effort required to open the settings menu or navigate within it, are substantial relative to the differences in how much users value the various data-sharing options. Even small obstacles can outweigh meaningful preferences. The cost of closing the banner without making a choice is also generally lower than the cost of engaging with settings, making inaction an attractive fallback for many users. Importantly, there is significant variation in preferences between user types. Some individuals have strong and consistent preferences for one option over the others (either always accepting or always rejecting cookies). Others assign more moderate differences between the data sharing options.

The estimates allow us to evaluate different consent policies. First, we consider banning dark patterns, akin to the recent EU Digital Services Act.⁷ In this regime, website default will matter, given consumers' increasing inclination to close consent windows without explicit choices. We find that a neutral design combined with a share-all default maximizes consumer surplus within our sample. This policy improves consumer surplus compared to the existing norm in the EU, where the default is no-sharing, and in the US, where dark patterns are predominant.

An alternative policy recommendation involves removing the site-specific consent banners altogether and instead allowing consumers to implement browser-level consent. We find that the browser-level choice outperforms consent-based policies by a wide margin, improving consumer welfare by at least 73% compared with policies based on

⁶Our participants interact with 54 unique web domains on any given week.

⁷https://commission.europa.eu/strategy-and-policy/priorities-2019-2024/europe-fit-digital-age/digital-services-act_en.

website-by-website consent. This recommendation is further supported by our survey data: 54% of participants indicate a preference for browser-level or global privacy controls, rather than interacting with cookie banners on each site.

Our work contributes to the growing literature on the economics of privacy, particularly studies that measure privacy preferences (Lin 2022; Collis et al. 2021; Ke and Sudhir 2023; Tomaino, Wertenbroch, and Walters 2023; Tang 2023; Acquisti, John, and Loewenstein 2013) and examine the consumer impact of data collection practices (Miller and Tucker 2018; Tang 2019; Zhao, Yildirim, and Chintagunta 2021; Bian et al. 2023). Importantly, we highlight that privacy choices in our setting are shaped not only by consumers' aversion to data sharing but also by their perceived benefits from cookie-based personalization and functionality. Our contribution lies in measuring these choices in a naturalistic, real-world setting and using them to estimate consumer surplus under alternative data collection regimes, thus offering a policy-relevant view of privacy trade-offs.

We also contribute to the broader literature on choice architecture, particularly in the context of privacy decisions. Most empirical work on dark patterns has focused on documenting their prevalence across online platforms (Mathur et al. 2019; Di Geronimo et al. 2020; Nouwens et al. 2020; Warberg et al. 2023). Attempts to quantify their effects on user behavior have typically relied on lab experiments or artificial 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). A few exceptions, such as D'Assergio et al. (2022) and Müller-Tribbensee, Miller, and Skiera (2024), examine persuasive design in real-world contexts. Our study extends this literature by capturing privacy choices made during everyday browsing, thereby offering greater real-world relevance. In addition, our analysis covers a wide range of websites, allowing us to assess how dark patterns interact with domain familiarity and popularity, key factors in understanding their implications for consumer behavior and competitive access to data.

Finally, our work relates to the existing literature on behavioral biases and their implication for firm competition (Huck and Zhou 2011; Spiegler 2014; Ho, Hogan, and Scott Morton 2017; Decarolis, Li, and Paternollo 2023). This literature examines how factors such as switching costs and obfuscation strategies can limit competition in product markets. Our findings extend this analysis to the domain of data collection strategies, showing that firms leverage manipulative interface designs to steer users toward sharing more data. The effectiveness of these interventions appears constant across websites of varying popularity. Thus, policies that target the use of dark patterns may not have pro-competitive effects. In fact, to the extent that the marginal value of additional user data is higher for smaller firms (for example, because consumer data is particularly useful

for targeted advertising and personalization (Aridor et al. 2024; Johnson et al. 2024)), a ban on dark patterns may even negatively impact competition.

The rest of the paper is structured as follows. Section 2 presents our experimental design and describes the study participants. We discuss our descriptive findings and experimental treatment effects in Section 3, and our model of user privacy preferences in Section 4. We discuss the time cost of consent in subsection 3.3 and conclude in Section 5.

2. Study Design

In this section, we present our study design. First, we introduce our consent interfaces, then we discuss the study phases and randomization, the sample, and summary statistics.

Consent Interfaces. The goal of our experiment is to identify how people make cookie tracking 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 consent interfaces that prompt users to make consequential cookie tracking choices.

The design of our consent interfaces is motivated by prior work documenting how companies use dark patterns to encourage data sharing (Habib et al. 2022; Bielova, Santos, and Gray 2024). These patterns fall into three main categories. The first, obstruction tactics, create friction around privacy-protecting choices—for example, by setting defaults to “accept all cookies” (“accept all” henceforth) or hiding “reject all cookies” (“reject all” henceforth) behind additional clicks. These are among the most common patterns observed online (Habib et al. 2022). The second type, visual manipulation, alters option salience through layout, color, or order—for instance, graying out “reject all” or placing “accept all” first. The third category involves persuasion tactics, such as framing data sharing in a more favorable light. We abstain from studying the effects of persuasion and framing, since text is high-dimensional, and a comprehensive study would require many more treatment arms than subjects that we can recruit.

Figure 1 displays the interfaces we designed. Design C (“Set-Acc-Rej”) is what we consider a neutral setting, where we remove all the dark patterns we set out to test. The other interfaces are manifestations of obstruction and visual manipulation. For example, Design A (“Acc-Set”) hides the reject option, making rejection only possible through “Cookie Settings” (*deliberate obstruction*); Design D (“Acc-Rej-Set”) prioritizes accepting cookies by listing it as *the first option (reordering options)*; and Design F (“Acc-

GreyRej-GreySet”) emphasizes the accept button with a *brighter color* than the other options (*differential salience*).

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,” the other removes “accept all”), two reordered interfaces (one with “accept all” on top, the other with “reject all” on top), and one interface with differential salience (where “accept all” is at the top in blue, whereas the other options are below in gray).

Participants can click on any of the options displayed, or avoid making an explicit choice by clicking the *X* in the top right corner. If they click “*X*”, the website will implement its default data-sharing setting, which is normally “accept all” for US websites. If they click on “Cookie Settings,” they are presented with six different types of cookies to choose from, such as information storage & access, performance & analytics, and ad selection delivery & reporting (see Appendix Table A1). Selecting all options is equivalent to accepting all cookies; selecting none of the options is equivalent to rejecting all cookies. To minimize choice friction under “Cookie Settings”, we allow consumers to either accept all cookies in one click, or reject all cookies with ease (as the default on this page is selecting none of the category-specific cookies).


Our banner can appear on any website. It replaces the organic cookie choice interface when one is present, and provides a consent interface when one is absent.⁸ To ensure participants’ data-sharing choices are incentive-compatible, we enforce their decision to the greatest extent possible by integrating our extension with a script that detects elements of a webpage related to cookie consent through a set of rules, some of which were taken from other open-source packages (e.g., DuckDuckGo’s AutoConsent) and others custom-made by Audacious software. Cookie Manager attempts to enforce the decision made through our own banner by selecting the appropriate option in the cookie consent form natively displayed on the website. When users close the consent window without an active choice, Cookie Manager does not attempt to change any settings on the website. We communicate the cookie choice enforcement to our participants.

Study Phases and Randomization. As further discussed below, we recruit participants through Prolific, screen for eligibility, and instruct qualified and consenting users to install Cookie Manager on their Chrome browser. The study proceeds in two phases: a structured survey phase followed by an organic browsing phase.

⁸We experiment on US residents, so most of them do not see consent banners as often, due to the lack of federal privacy regulations that require consent.

⁸We tell users “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.”

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-Set”: accept-settings; “Rej-Set”: reject-settings; “Set-Acc-Rej”: settings-accept-reject; “Acc-Rej-Set”: accept-reject-settings; “Rej-Acc-Set”: reject-accept-settings; “Acc-GreyRej-GreySet”: accept-grey reject-grey-settings.

In the survey phase, participants are asked to visit 20 pre-selected websites (see Appendix A4 for a full list), presented in randomized order. As participants browse each site, the extension displays a randomly assigned cookie consent banner from those displayed in Figure 1, and records their choice. The selected websites span a range of categories (including social media, e-commerce, and news) and popularity levels. This phase serves two purposes: first, it ensures sufficient observations per domain to estimate dark pattern effects using fixed-effect specifications; second, it allows us to measure participants’ familiarity with each site, which would be difficult to infer from organic behavior alone.

In the organic phase, participants continue their regular browsing for one week while the extension remains active. Consent banners are displayed upon visiting domains for which there has not previously been a consent choice during the study. To reduce disruption, banners appear no more than once every 10 minutes. As a result, we do not record consent choices for some domains participants visit during this phase.⁹ At the

⁹ In this second phase, we also introduce variation in banner exposure frequency. Specifically, we randomize users into two frequency treatments: in the frequent pop-ups group, banners appear at most every 10 minutes; in the infrequent pop-ups group, banners appear every 60 minutes. To implement this, a countdown starts after each banner interaction, and a new banner appears after the 10- or 60-minute

end of the week, participants complete an exit survey and uninstall the extension. Each participant receives \$7.50 upon study completion. Survey instruments are provided in Appendix D.

Consent interface randomization occurs at the user-by-domain level: the extension randomly selects a banner design when a participant visits a domain for the first time since enrolling in the study, and tracks the corresponding choice.¹⁰ This approach increases statistical power by leveraging within-user variation across sites.

Sample Description. We recruited participants on Prolific, and restricted our participants to adults residing in the US who primarily speak English and use Chrome as their main browser.¹¹ We pre-registered recruiting 800 participants and expected 640 of them to complete the study.¹² Our actual participants are close to the pre-registered numbers (see Appendix Table A5 for the conversion funnel). A total of 1,227 Prolific users started the study; 75% of respondents were eligible. Among these, 877 consented to the study.

Our final sample includes all participants who completed the baseline survey and generated valid data during the organic browsing phase, regardless of whether they completed the exit survey. For our main analysis, we exclude participants who, due to an implementation issue, continued to receive a consent interface on every new domain visit rather than at the intended less frequent interval.¹³ We also exclude participants who reported making choices based on their expectations about the study, rather than their genuine privacy preferences. These restrictions ensure a consistent sample across both survey and organic phases. After applying these criteria, our main analysis sample consists of 563 participants.

Table 1 presents descriptive statistics for the main sample. 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 is \$50,000-\$74,999, with substantial variation, including 12% of households with an income of over \$150,000. In the week prior to enrollment, users visited an average of 53 unique domains. During the study week (excluding the survey-assigned websites), participants visited an average of 54 unique domains, suggesting that

threshold is crossed and a new domain is visited, whichever occurs later. We do not find differences in consumer behavior across these two treatments.

¹⁰In pilot testing, we found no evidence of carryover effects from prior exposures to specific banner designs.

¹¹<https://www.prolific.com/>.

¹²<https://www.socialscienceregistry.org/trials/12862>.

¹³Approximately 3% of users were not assigned to the frequency condition and saw banners for every new domain throughout the organic phase.

most did not alter their typical browsing behavior in response to being tracked. During the organic phase, our consent interfaces appear in 41% of the visited domains.¹⁴

TABLE 1. Summary Statistics

		Mean	Median	Std. Dev.
During Survey	Unique Domains in Prior Week	53.27	49.00	39.69
	Domains w. Banner	19.67	20.00	0.82
Post-Survey	Domains w. Banner	22.40	15.00	21.94
	Unique Domains Visited	54.45	39.00	51.19
	Unique URLs	670.75	393.00	817.41
	End Survey Completed	0.86	1.00	0.35
Demographics	Age	37.97	36.00	12.82
	Female	0.45	0.00	0.50
	Bachelor’s or Above	0.17	0.00	0.38
	Income > \$75,000	0.44	0.00	0.50
Cookie Behavior	Accept-All Rate	0.54	0.64	0.36
	Close-Window Rate	0.26	0.14	0.32
	Reject-All Rate	0.16	0.00	0.28

Notes: The table shows user-level descriptive statistics for the final study sample. Number of observations: 563

We verify effective randomization of consent pop-ups in two ways. First, we run a proportion test on the distribution of pop-ups per website. The proportion test for the distribution of pop-ups across the survey websites has a p-value of 0.96, which fails to reject the null of balanced proportions across the six pop-up designs. Second, we perform covariate balance tests by regressing user- and domain-level covariates on treatment conditions (Appendix Table A6). We find no statistically significant differences across pop-up designs.¹⁵

3. Experimental Results

In this section, we present reduced-form evidence on the causal effects of dark patterns, explore heterogeneity in privacy choices across users and domains, and examine participants’ beliefs about cookie tracking. We find that the majority of users accept all cookies when presented with a neutral interface, and that dark patterns that increase choice

¹⁴This average masks heterogeneity induced by our experiment (see footnote 9). Participants in the 10-minute frequency treatment are exposed to the pop-up for 53% of the domains visited. In comparison, participants in the 60-minute frequency treatment see the pop-up in 30% of the domains visited.

¹⁵In 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 (footnote 9).

friction significantly shift consent behavior. While domain-level factors such as popularity have modest effects, user-level heterogeneity is substantial. Survey responses further indicate that most participants had at least a basic understanding of how cookies function and the potential consequences of data sharing, suggesting that consent decisions were meaningfully informed. We conclude by quantifying the time costs associated with making cookie preference decisions.

3.1. The Effect of Dark Patterns on Data Sharing Choices

Figure 2 presents the choice distribution across treatment conditions, separately for the survey (Panel A) and organic phases (Panel B), as user behavior may differ between structured tasks and natural browsing. Three main findings are worth highlighting. First, participants share their data with the websites more than 50% of the time, even when the banner design is neutral. In particular, absent deliberate nudging, the accept-all rate is 65 percent in the survey and 61 percent during organic browsing.¹⁶ The exception is the condition where we hide “accept all” from the main screen, where 19 percent of consumers choose to accept all cookies by going to the “cookie settings” page.

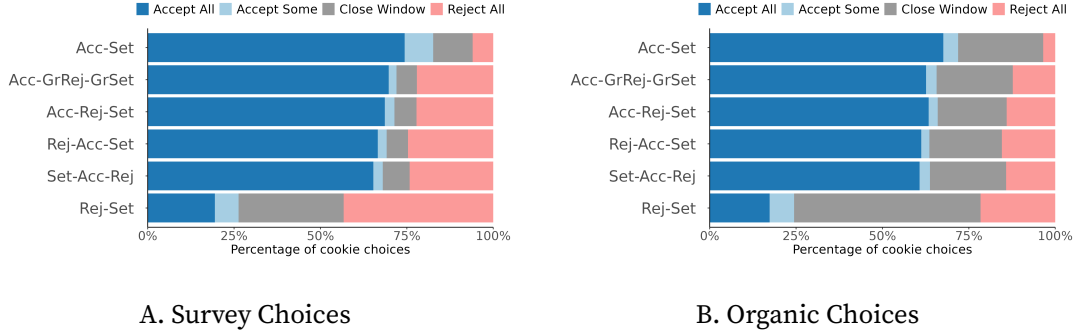
Second, granular choices, defined as accepting only a subset of cookies, are relatively rare. Across all treatment conditions, participants make selective choices infrequently, ranging from 3% in the neutral interface to 6–9% when either the “accept all” or “reject all” option is deliberately hidden.

Third, participants respond similarly to dark patterns across both the survey and organic browsing phases, with one key exception: they are substantially more likely to close the consent window during organic browsing. When focused on the task of making cookie selections in the survey, participants are far less likely to disengage. This difference in behavior highlights the role of attention and the importance of default settings when users are primarily engaged in non-privacy-related browsing tasks. However, incorporating users’ beliefs about what happens when they close the window (further discussed in Section 3.2) suggests that the underlying preferences across the two phases are more aligned than the raw data imply. After imputing passive choices using belief data, we estimate that 68 percent of users in the survey phase and 66 percent in the organic phase intend to accept cookies. Similarly, 28 percent in both phases appear to intend to reject them. Together, these findings underscore two points: first, that survey-based

¹⁶Our acceptance rates are high but lower than prior evidence in lab studies that also examine privacy choices absent dark patterns, which document the probability of choosing “accept all” to be 77% (Vázquez Duque 2024) to 83% (Bielova et al. 2024). Existing reports on ATT consent rates are often lower, at around 44% as of 2024; see: <https://www.appsflyer.com/company/newsroom/pr/att-data-findings/>. However, Apple is known to use non-neutral languages that discourage sharing (Baviskar et al. 2024), and we suspect that some mobile data (such as location) can be more sensitive compared to web behavior data.

choices, though potentially more artificial, still reliably capture the effects of dark patterns; and second, that default settings play a critical role when users shift their attention from privacy decisions to other browsing tasks.

FIGURE 2. Cookie Choices by Experimental Condition



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

To quantify the causal effects of our treatments on consumer choice we estimate the following type of regressions:¹⁷

$$y_{ij} = \beta_{acc-set_{ij}} + \beta_{acc-greyrej-greyset_{ij}} + \beta_{acc-rej-set_{ij}} + \beta_{rej-acc-set_{ij}} + \beta_{rej-set_{ij}} + \mu_i + \nu_{c(j)} + \epsilon_{ij}. \quad (1)$$

Here, i denotes the participant and j denotes the website. We include participant fixed effects μ_i and website category fixed effects $\nu_{c(j)}$, where categories are based on Ghostery’s taxonomy¹⁸ and extended to the full set of websites in our study using ChatGPT.¹⁹

¹⁷Our treatments often combine different types of manipulations into a single interface. For example, “acc-set” simultaneously prioritizes “accept all” and hides “reject all.” To see the separate effects of individual dark patterns (re-ordering, obstruction, and highlighting) on consumer choices, please refer to Appendix A.

¹⁸<https://www.ghostery.com/>.

¹⁹We use ChatGPT 4o to classify the websites. We use the following prompt: “Classify the website domains listed below into one of the following major categories (and only one of the following, do not include categories not in this list and try to limit how often other is selected): ‘Reference Website’, ‘Entertainment Website’, ‘Business Website’, ‘E Commerce Website’, ‘Adult Website’, ‘News and Portals Website’, ‘Recreation Website’, ‘Banking Website’, ‘Government Website’, ‘Political Website’, ‘Other’. Each domain is separated by ‘+ – – +’. Please always return a 10-element sequence of classifications separated by ‘– – –’. The list of domains is: [the list of domains follows].”

Each of the β coefficients measures the effect of a specific treatment condition relative to the neutral interface (Condition C in Figure 1).²⁰ We focus on three outcomes: accepting all cookies, rejecting all cookies, and closing the window without making an active choice. Given its small share, the analysis of the decision to select specific types of cookies is left to Appendix Table A8.

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	0.083*** (0.013)	-0.178*** (0.015)	0.039*** (0.009)	0.054*** (0.012)	-0.094*** (0.013)	0.019 (0.013)
Acc-GrRej-GrSet	0.035*** (0.010)	-0.017 (0.010)	-0.012 (0.006)	0.031** (0.011)	-0.016 (0.008)	-0.020* (0.010)
Acc-Rej-Set	0.020 (0.011)	-0.007 (0.009)	-0.010 (0.006)	0.038** (0.011)	-0.001 (0.007)	-0.041*** (0.010)
Rej-Acc-Set	0.003 (0.010)	0.014 (0.010)	-0.012* (0.006)	0.004 (0.011)	0.020* (0.009)	-0.023* (0.010)
Rej-Set	-0.464*** (0.020)	0.193*** (0.017)	0.233*** (0.018)	-0.427*** (0.024)	0.086*** (0.012)	0.302*** (0.023)
Benchmark group mean	0.65	0.24	0.08	0.61	0.14	0.22
R ²	0.646	0.579	0.562	0.571	0.522	0.494
Observations	11,075	11,075	11,075	12,610	12,610	12,610
Participant fixed effects	✓	✓	✓	✓	✓	✓
Domain Cat. fixed effects	✓	✓	✓	✓	✓	✓

Notes: Regression results of Equation 1 for three outcomes: accept all cookies, reject all cookies, and close the 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). Standard errors clustered at the participant level. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Table 2 displays our main results, with standard errors clustered at the participant level. Columns 1 through 3 focus on survey choices and jointly capture how each treatment affects substitution across consent choices, whereas columns 4 through 6 focus on organic choices.

Treatments involving deliberate obstruction are the most effective at steering consumers. Removing the “reject all” option (the *Acc-Set* treatment) decreases rejection rates by 18 percentage points (a 74% decline). The modal user substitutes towards acceptance, whose rates increase by 8 percentage points in the survey and 5 percentage points in the organic phase. Closing the window increases modestly in the survey and remains flat in the organic phase.

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

Removing the “accept all” option (the *Rej-Set* treatment) has even bigger effects, given that acceptance is the most frequent choice at baseline. Indeed, this treatment decreases acceptance rates by 46 percentage points in the survey phase (a 71% decrease) and by 43 percentage points in the organic phase (a 70% decrease). In the survey, this leads to large shifts toward both rejecting cookies (up 19 percentage points) and closing the window (up 23 percentage points). In the organic phase, the dominant response is again to close the window, which increases by 30 percentage points, more than doubling the baseline close rate.

Visual manipulations—such as reordering and highlighting—have modest effects, concentrated during the organic phase, perhaps because users are not as focused on cookie preferences while naturally browsing the internet. Placing “accept all” at the top (*Acc-Rej-Set*) increases acceptance by 3.8 percentage points in the organic phase (a 6% increase), while its effect in the survey is negligible. Placing “reject all” at the top (*Rej-Acc-Set*) slightly increases rejections in the organic phase (up 2 percentage points, a 14% increase), but has no meaningful effect in the survey. Highlighting “accept all” (*Acc-GrRej-GrSet*) modestly increases acceptance in the organic phase by 3.1 percentage points, but this effect is nearly identical to that of simply placing “accept all” at the top without additional visual cues. Most of the substitutions appear to come from users who would otherwise have closed the window.

Appendix Table A8 shows that consumers rarely make granular cookie choices unless prompted by design interventions. In the neutral condition, only 3% of users accept a subset of cookie categories. However, hiding either the “accept all” or “reject all” option from the main screen nudges users to explore the settings menu, increasing the likelihood of granular choices by 2-5 percentage points.

Among those who do make selective choices, 83% accept cookies for *preferences and functionality*, while only 7% opt in to *ad selection, delivery, and reporting* (see Appendix Table A9). This pattern suggests that targeted advertising is the least preferred use of consumer data, at least among users who selectively consent to cookie tracking.

These results point to three main conclusions. First, users often accept cookie tracking absent dark patterns while browsing the web. Second, dark patterns relying on obstruction are far more powerful than those based on visual design alone, although visual manipulations can be modestly effective when participants are not explicitly focused on the task of managing cookie preferences.²¹ This, closing the window without making an active

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 find 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 find small to no effect on choices.

choice is a frequent selection, even more true in the wild than in a synthetic survey-based setting, signifying the importance of website defaults.

3.2. Heterogeneity across Websites and Consumers

In addition to the effects of dark patterns, we examine how cookie choices vary across websites and individuals, to help assess the potential competitive implications of consent design. Prior work has documented that privacy decisions are highly context-dependent (Nissenbaum 2004; Lin 2022), and our results extend this by showing substantial heterogeneity at the user level even holding contexts fixed. Overall in our study, 92.2% of participants change their cookie-sharing decisions across sites at least once. This variation is a combination of the effects of dark patterns, which we have already showed, and systematic variation across websites and users, both of which are discussed next.

Website-Level Heterogeneity. We begin by investigating whether consent behavior differs across websites and whether dark patterns amplify or mitigate these differences. Specifically, we assess how cookie-sharing decisions vary with website category, users' prior experience, and overall domain popularity.


Figure 3 displays consent rates by site category, using e-commerce as the reference group. The estimates correspond to the category fixed effects from Equation 1, where the outcome is "accept all." Users are substantially less likely to accept cookies on adult, political, and government websites. This suggests that participants are more hesitant to share data with websites perceived as more sensitive, and more comfortable doing so on sites, like e-commerce, that provide recommendations for more mundane purposes.

Next, we examine how user familiarity with a website shapes consent decisions. We draw on two sources to measure participants' experience with each site. First, during the survey phase, participants reported whether they had heard of and regularly visited each of the 20 assigned websites. Second, we use browsing history data to determine whether a given domain had been visited in the two weeks prior to enrollment in the study. To capture overall site popularity, we use log domain ranks from Tranco,²² which aggregates rankings from multiple public sources.

Panels *a* and *b* of Table 3 report regression results examining how cookie choices vary with domain-level measures of user familiarity. Panel *a* focuses on the survey phase, using self-reported measures of website experience. Participants are 2.6 percentage points more likely to accept cookies on websites they report having heard of, and an additional

²²<https://tranco-list.eu/>.

FIGURE 3. Consent Rates by Site Category



Notes: Differences in cookie acceptance rates by website category. The plot shows the estimated category fixed effects from Equation 1, where the outcome is “accept all” and the baseline category is e-commerce websites. Bars denote 95% confidence intervals.

7.1 percentage points more likely on websites they say they normally visit. These effects are statistically significant and primarily offset a decline in the likelihood of rejecting cookies, suggesting substitution between acceptance and rejection rather than closing the window.

Panel *b* combines data from both the survey and organic phases, using observed proxies for experience. Participants are 1.8 percentage points more likely to accept cookies on websites they had visited prior to the experiment. Additionally, domain popularity (measured by the log of domain rank) is associated with greater willingness to accept cookies instead of closing the window, suggesting that users are more inclined to share data on better-known sites.

Together, these results support the interpretation that user familiarity, whether self-reported or inferred from usage data, is associated with a greater likelihood of accepting cookies. This is consistent with our open-ended survey responses, in which participants frequently cite trust in the website or brand as a motivation for acceptance.

We also test whether dark patterns are more effective on more or less familiar websites. Appendix Figure A2 and Appendix Tables A10 and A11 show limited heterogeneity in the effectiveness of dark patterns across sites. If anything, dark patterns appear to dampen

TABLE 3. Heterogeneity in Cookie Choices across Websites and Users

	Accept All (1)	Reject All (2)	Close Window (3)
<i>Panel a: Experience based on survey answers (survey data only)</i>			
Normally Visit	0.071*** (0.009)	-0.072*** (0.010)	-0.010 (0.006)
Heard Of	0.026** (0.009)	-0.026** (0.009)	-0.002 (0.006)
R ²	0.652	0.587	0.562
Observations	11,075	11,075	11,075
<i>Panel b: Experience based on browsing history and site popularity</i>			
Pre-Exp Visit	0.018* (0.008)	-0.031*** (0.006)	0.011 (0.006)
Domain Rank (Log 10)	-0.010*** (0.002)	0.002 (0.002)	0.007*** (0.002)
R ²	0.518	0.470	0.422
Observations	23,685	23,685	23,685
<i>Panel c: Demographics</i>			
Bachelor's or Above	-0.086 (0.044)	0.017 (0.034)	0.056 (0.039)
Income > \$75,000	-0.004 (0.031)	-0.001 (0.023)	0.001 (0.023)
Age	0.001 (0.001)	-0.001 (0.001)	-0.002** (0.001)
Female	-0.083** (0.029)	0.048* (0.023)	0.042 (0.023)
R ²	0.147	0.069	0.122
Observations	22,460	22,460	22,460
<i>Panel d: Beliefs about Privacy</i>			
Privacy: Meaningful Decision	0.025 (0.035)	-0.005 (0.028)	-0.030 (0.024)
High Value Privacy	-0.013 (0.049)	0.070* (0.028)	-0.078 (0.043)
Outtake Missing	-0.060 (0.057)	0.074 (0.041)	-0.047 (0.047)
R ²	0.135	0.063	0.113
Observations	23,685	23,685	23,685

Notes: Regression results of Equation 1, in which we add explanatory variables to explore heterogeneity in cookie tracking choices across websites. In Panel *a*, we add two dummies to indicate whether the study participant has heard of the website and whether the study participant normally visits the website (both questions are answered as part of the survey browsing, so we only include survey choices). In Panel *b*, we aggregate across browsing and survey data allowing for a level shift between the two. The dimensions of heterogeneity are proxied by a dummy to indicate whether the study participant visited the website in the two weeks preceding the study (we obtain this information by collecting their Chrome browsing history) and the website popularity rank (in logs) from Tranco. Panel *c* includes individual demographics (age, gender, income, and education). Panel *d* focuses on self-reported and observed privacy attitudes. We include indicators for participants' beliefs about the meaningfulness of privacy decisions, whether they highly value privacy, and whether they left the study prior to completing the final survey. Regressions in panel *c* and *d* require us to exclude individual fixed effects. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

participants' tendency to share more data with popular and familiar websites, suggesting they are unlikely to reinforce data-driven competitive advantages in practice.

Individual-Level Heterogeneity and User Beliefs. We next examine variation in cookie preferences across users. While earlier sections documented heterogeneity in behavior across websites, Panel c of Table 3 shifts the focus to demographic predictors of privacy choices. Specifically, we test whether demographic characteristics help explain variation in cookie acceptance, rejection, or banner closure. Overall, most demographic factors have limited predictive power, but some meaningful patterns emerge. Women and participants with at least a bachelor degree are modestly less likely to accept all cookies (by 6.5 and 8.8 percentage points, respectively), and older participants are slightly less likely to close the banner without making a choice.

While these patterns suggest some demographic tendencies in privacy behavior, they raise a further question: to what extent are users' choices informed by an understanding of what cookies do and how they affect data sharing? Indeed, a common concern in studies of online privacy behavior is that users may lack the information needed to make meaningful choices. To address this, our endline survey asked participants to explain their reasons for accepting or rejecting cookies during the study, and to describe their broader beliefs about the consequences of data sharing.

Among those who accepted all cookies, a majority cited reasons consistent with informed behavior. 36% reported doing so because they trusted the website or viewed the site as reputable (e.g., "It was a website I totally trusted"), while 20% cited practical motivations such as faster navigation or convenience (e.g., "It was just easier to do so. I didn't want to have to spend time looking through my options every time I opened a new site"). Another 17% accepted cookies to ensure proper website functionality (e.g., "To make the websites function well").

Similarly, among those who rejected cookies, the most common explanations were privacy-related: 38% cited concerns about unfamiliar or untrustworthy websites ("Did not trust site"), 9% explicitly referenced tracking or advertising concerns ("I didn't want to be tracked by advertisers on that website"), and 12% pointed to a general preference for privacy ("in some website i felt not to accept some cookies because i may not want the website to have my data").

We also asked participants what they believed would happen if they closed the banner using the "X" button. 61% believed this action meant rejecting cookies, while 25% believed it meant accepting cookies. The remaining respondents expressed uncertainty. Notably, consumers with uncertain beliefs close windows at a rate of 13%, while those with certain

beliefs close windows for 20% of the time. This variation in beliefs is accounted for in our structural model in the next section, but the overall distribution suggests that most users made choices with a specific outcome in mind, which we described earlier in Section 3.1.

Finally, when asked about the benefits and costs of sharing data (regardless of their preferred choice), participants demonstrated an intuitive grasp of the core trade-offs. Roughly 69% of respondents identified personalized ads and 15% identified better site functionality as potential benefits, while over 50% cited privacy loss (43%), security risks (16%), or increased advertising (14%) as key concerns. Only a small subset (less than 5%) reported being unsure or unaware of what cookies do.

To further test whether these attitudes align with actions, we examine whether stated privacy values are correlated with actual consent behavior. Participants who say they highly value privacy are significantly more likely to reject all cookies (by 7 percentage points). Still, the majority of coefficients are indistinguishable from zero, and the low R^2 values across models suggest that much of the variation in privacy behavior remains unexplained by observables, likely driven by individual preferences or traits not captured in the survey.

Appendix Tables A13 and A12 further validate that self-reported privacy attitudes align with observed behavior. Participants who report having accepted or rejected most cookies during the study indeed exhibit those respective behaviors in the data. Moreover, when we regress cookie choices on participants' stated motivations, we find patterns that are directionally consistent with expectations. Respondents who said they accepted cookies due to trust or site functionality are more likely to have accepted cookies, while those citing distrust, unfamiliarity, or privacy concerns are less likely to do so. These belief-based results are statistically noisy and measured post hoc, so we include them primarily for completeness in the appendix. Nonetheless, they provide additional support for the interpretation that stated privacy attitudes reflect meaningful choices.

Taken together, these results suggest that most participants made consent choices with at least some understanding of the role of cookies and the implications of data sharing.

Despite observable characteristics explaining some of the variation in cookie decisions, a large share of the variation remains unexplained. To disentangle user-driven preferences from domain-specific effects, we estimate a random effects model in which we regress the probability of accepting all cookies on treatment indicators, participant random effects, and domain random effects, both with and without the covariates used in the panels of Figure 3. Appendix Table A17 shows that the standard deviation of the participant random effect is more than five times greater than that of the domain effect,

even after adjusting for all covariates. This highlights that individual user differences contribute far more to cookie acceptance behavior than differences between websites.

3.3. The Time Cost of Consent

A common criticism of consent-based privacy regulations is that repeated pop-ups degrade user experience and impose time costs. Here, we quantify the direct costs of asking users for consent through banner interactions.

We begin by calculating how much time people spend interacting with our consent banners and how that time varies by banner design. To measure time spent on the pop-up, we calculate the time elapsed between the cookie pop-up and the recorded *final* action. This measure includes time spent on intermediate clicks, as well as any back-and-forth interactions before the banner is closed. We censor the time spent at 60 seconds, which is well above the 99th percentile of time spent on the banner (13 seconds during organic browsing) and likely reflects task switching rather than genuine time spent.

Like in Section 3.1, we estimate Equation 1 where the outcome is the time spent interacting with the banner. Table 4 displays the results. In the neutral design condition, consumers on average spend 5.25 seconds per banner in the survey phase (column 1) and 7.34 seconds in the organic phase (column 2). The only design that significantly affects time spent is the *rej-set* design, which hides “accept all,” the most commonly chosen option.

To translate this into monetary terms, and as a preview of how we will quantify welfare in the next section, we perform a back-of-the-envelope calculation using the U.S. average hourly wage of \$36 as a baseline.²³ With users spending 7.34 seconds per domain and visiting 54.5 unique domains per week on average, this implies a weekly time cost of approximately \$4 per user. Using the higher \$69/hour estimate from Greminger, Huang, and Morozov (2023) as an upper bound increases this estimate to \$7.67 per week.

Dark patterns that add friction further raise the time cost. In particular, when the “accept all” option is hidden, the time spent increases, resulting in a weekly cost of \$4.60 at the baseline rate (\$8.80 using the upper bound). These findings underscore that deliberate obstruction not only distorts user choices but also imposes significant and quantifiable time costs. In the next section, we use these reduced-form findings to inform a unified model of user behavior and evaluate the welfare implications of alternative consent policies.

²³The \$36/hour estimate is based on the average U.S. wage reported by the Bureau of Labor Statistics (<https://www.bls.gov/news.release/empsit.t19.htm>). In contrast, Greminger, Huang, and Morozov (2023) estimate a higher opportunity cost of time—\$69/hour—by accounting for both wage-based value and the disutility of searching for products online. Their figure can be interpreted as an upper bound.

TABLE 4. Time Spent by Experimental Condition

	Survey	Organic
	Time Spent (Seconds)	
	(1)	(2)
Acc-Set	0.497 (0.291)	-0.524 (0.382)
Acc-GrRej-GrSet	-0.144 (0.259)	-0.192 (0.393)
Acc-Rej-Set	-0.426 (0.234)	0.144 (0.421)
Rej-Acc-Set	-0.257 (0.242)	-0.363 (0.379)
Rej-Set	2.125*** (0.281)	1.110** (0.423)
Benchmark group mean	5.25	7.34
R ²	0.193	0.120
Observations	11,075	12,565
Participant fixed effects	✓	✓
Domain Cat. fixed effects	✓	✓

Notes: Regression results of Equation 1 where the outcome is the time spent interacting with the cookie consent banner. The results are presented separately for two different sets of choices: survey choices (column 1) and organic choices (column 2). Standard errors clustered at the participant level. * $p < 0.05$, ** $p < 0.01$; *** $p < 0.001$.

4. Structural Model and Estimation

To quantify the welfare implications of different consent choice architectures, we now introduce a model of consumer data sharing decisions that incorporates the influence of dark patterns. The model allows us to translate observed choices into welfare measures and quantify the costs imposed by choice frictions, in particular, obstruction and visual manipulations. More importantly, it enables us to evaluate policy interventions not directly tested in our experiments, such as policies that require alternative default settings or browser-level consent solutions.

We model consumer utility as a function of the selected option and the frictions introduced by the design of the consent interface. (For ease of comparison, Table 5 summarizes the utilities for all options in the model and how they vary across conditions.) When faced with a cookie consent banner, a consumer can choose among four different actions: {accept all, reject all, customize settings, close window}. For notational simplicity, we omit participant and website subscripts and let θ_k denote the utility of data-sharing

option k . Specifically, $k = acc$ corresponds to accepting all cookies, $k = rej$ to rejecting all cookies (with utility θ_{rej} normalized to zero), and $k = set$ to customizing settings.²⁴

The utility of closing the consent banner depends on a consumer’s belief about the website’s data collection default and the cost of clicking to close the pop-up, κ . The cost term κ captures many reasons why closing the pop-up is differentially attractive than making an explicit choice.²⁵

As discussed in Section 3.2, we use the endline survey to elicit users’ beliefs about the website’s default behavior when the consent banner is closed. Let ρ_{belief} denote participants’ subjective probability that closing the window results in full cookie acceptance. For respondents who believe that closing the window means accepting cookies (52%), we set $\rho_{acc} = 1$; for those who believe it means rejecting cookies (22%), we set $\rho_{rej} = 0$. For participants who are uncertain (12%), we estimate their belief as ρ_{unc} . Under the assumption of risk neutrality, ρ_{unc} can be interpreted as the probability that these participants assign to the event that closing the window implies full cookie acceptance. For the remaining participants (14%) who did not complete the endline survey, we estimate ρ_{miss} and interpret this parameter analogously, as the probability they implicitly assign to cookie acceptance upon closing the window.²⁶

Choice architecture influences consumer decisions by either introducing friction or altering the salience of available options. We model friction as click-based costs and visual salience as multiplicative adjustments to utility. Specifically, consumers incur a cost C_{set} for clicking on the settings button to access additional options, and an additional cost $C_{set-acc}$ for selecting “accept all” once inside the settings menu (since this option requires an extra click relative to rejecting all).

Visual manipulations—such as changing the order or appearance of options—affect perceived utility. We model these salience effects as multiplicative factors applied to the baseline utility of each option. For instance, placing the “accept all” option at the top of the list increases its utility from θ_k to $(1 + \delta_{top})\theta_k$, assuming θ_k is positive. Likewise, visually de-emphasizing an option (e.g., graying it out) reduces its utility to $(1 + \delta_{grey})\theta_k$, where $\delta_{grey} < 0$ reflects reduced salience. Given the normalization $\theta_{rej} = 0$, the utilities of

²⁴Our design does not separately identify the benefit of granular cookie selection from the cost of doing so. As a result, a user may assign high utility to sharing only selected cookies, but still avoid customization due to the associated cost.

²⁵For instance, a consumer may be less inclined to close the pop-up if the option is visually less prominent or if the consequences of doing so are ambiguous. Conversely, some users may prefer this option simply because it is consistently available and requires minimal effort.

²⁶We do not apply any discount for participants who explicitly believe that closing the window results in either full acceptance or full rejection. Our data (see Section 3.1) indicate that these users substitute between closing the window and their belief-consistent action at roughly a 1:1 ratio.

other alternatives may be estimated to be negative. To ensure proper estimation of the δ parameters while allowing for negative θ_k , we write $\theta_k + \delta_{top}|\theta_k|$ and $\theta_k + \delta_{grey}|\theta_k|$.

TABLE 5. Utilities Across Interface Conditions

Treatment	Accept All	Reject All	Customize Settings	Close Window
Acc-Set	$\theta_{acc} + \delta_{top} \theta_{acc} $	$-C_{set}$	$\theta_{set} - C_{set}$	$\rho_{belief}\theta_{acc} + \kappa$
Rej-Set	$\theta_{acc} - C_{set} - C_{set-acc}$	0	$\theta_{set} - C_{set}$	$\rho_{belief}\theta_{acc} + \kappa$
Set-Acc-Rej	θ_{acc}	0	$\theta_{set} + \delta_{top} \theta_{set} - C_{set}$	$\rho_{belief}\theta_{acc} + \kappa$
Acc-Rej-Set	$\theta_{acc} + \delta_{top} \theta_{acc} $	0	$\theta_{set} - C_{set}$	$\rho_{belief}\theta_{acc} + \kappa$
Rej-Acc-Set	θ_{acc}	0	$\theta_{set} - C_{set}$	$\rho_{belief}\theta_{acc} + \kappa$
Acc-GrRej-GrSet	$\theta_{acc} + \delta_{top} \theta_{acc} $	0	$\theta_{set} + \delta_{grey} \theta_{set} - C_{set}$	$\rho_{belief}\theta_{acc} + \kappa$

To allow for preference heterogeneity, we allow our set of parameters (θ_{acc} , θ_{set} , C_{set} , $C_{set-acc}$, δ_{top} , δ_{grey} , κ , ρ) to vary across three different groups of participants. We restrict attention to choices made when facing the neutral interface, and divide study participants into three groups: those who always accepted all cookie tracking (*acceptors*), those who always rejected them (*rejectors*), and those who varied their choices across websites when facing the neutral condition (*discerners*). Implicitly, we assume that the neutral condition reflects their true preferences, and estimate their parameters based on how choice architecture influences their choices under non-neutral interfaces. ρ_{belief} can take 4 values: $\rho_{acc} = 1$ for respondents who believe that by closing the window they accept all cookies, $\rho_{rej} = 0$ if they believe they reject all cookies, ρ_{unc} , and ρ_{miss} (to be estimated) for respondents who are unsure and for non-respondents.

A consumer i of type $t(i) \in \{\text{acceptors}, \text{rejectors}, \text{discerners}\}$ picks the option k that maximizes their utility given the treatment for website j :

$$\text{Max}_{k \in \{\text{acc}, \text{rej}, \text{set}, \text{close}\}} U_{t(i)k|treatment_{ij}} + \epsilon_{ijk}.$$

We assume that the errors ϵ_{ijk} are i.i.d. distributed according to a Type-1 extreme value distribution. Thus, we have the following choice probabilities:

$$Pr(i \text{ chooses option } k \text{ for website } j | treatment_{ij}) = \frac{\exp(U_{t(i)k|treatment_{ij}})}{\sum_{k \in \{\text{acc}, \text{rej}, \text{set}, \text{close}\}} \exp(U_{t(i)j|treatment_{ij}})}.$$

We estimate the model using maximum likelihood based on data from both the survey and organic phases. Identification of the model parameters relies on our experimental variation, which randomly assigns consent banners at the user-by-domain level. Preferences for specific data sharing choices are identified through their choices in the benchmark group without dark patterns. The costs associated with friction (e.g., accessing settings)

are identified by comparing choice probabilities under interfaces with and without “accept all” and “reject all” available on the main screen. Visual salience effects are identified by comparing choices between interfaces that reorder or gray out options while maintaining the same choice set. Beliefs about what happens when the banner is closed are separately elicited through the endline survey, which helps us infer the costs associated with closing the window. By estimating the model separately for user types based on their behavior in the neutral condition, we also capture preference heterogeneity without imposing strong parametric assumptions.

TABLE 6. Parameter Estimates Across Consumer Subsets

Parameter	Explanation	Pooled Estimate	Acceptors	Rejectors	Discerners
θ_{acc}	Utility of accepting all	0.917*** (0.004)	2.072*** (0.005)	-2.243*** (0.021)	0.460*** (0.040)
θ_{set}	Utility of granular choice	-0.143*** (0.010)	-0.397*** (0.008)	-0.948*** (0.016)	0.423*** (0.008)
C_{set}	Cost of clicking settings	1.618*** (0.010)	2.212*** (0.007)	2.317*** (0.016)	1.717*** (0.017)
$C_{set-acc}$	Cost of clicking accept after settings	0.103*** (0.011)	-0.131*** (0.015)	0.422*** (0.010)	0.123*** (0.009)
δ_{top}	Effect of ranking on top	0.375*** (0.009)	0.381*** (0.005)	0.097*** (0.023)	0.334*** (0.024)
δ_{grey}	Effect of being grayed out	-1.134*** (0.011)	-0.451*** (0.012)	-0.634*** (0.008)	-0.285*** (0.042)
κ	Utility shifter of closing window	-0.447*** (0.015)	-0.062*** (0.009)	-0.971*** (0.011)	-0.926*** (0.038)
ρ_{unc}	Pr(accept close) for uncertain	-0.563*** (0.022)	-0.132*** (0.062)	0.505*** (0.042)	-0.868*** (0.124)
ρ_{miss}	Pr(accept close) for missing	0.246*** (0.035)	0.337*** (0.037)	0.177*** (0.014)	-0.752*** (0.015)

Notes: Acceptors (55%), rejectors (10%), and discerners (31%) are defined using choices in the neutral design condition. The pooled estimates includes all three subsets and the rest of participants who always close window in the neutral condition. Standard errors are clustered at the participant level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 6 presents the parameter estimates from our structural model. The first column reports pooled estimates without allowing for consumer heterogeneity, while the next three columns show estimates separately for the three consumer types: acceptors, rejectors, and discerners. As expected, the utility of accepting all cookies (θ_{acc}) is strongly positive for acceptors, strongly negative for rejectors, and moderate for discerners. The utility of making a granular choice (θ_{set}) is negative for both acceptors and rejectors—consistent with their tendency to make binary choices—but positive for discerners, who appear more willing to engage with setting customization.

The friction parameters reveal that interface design imposes substantial behavioral costs. The cost of clicking on the “settings” button (C_{set}) ranges from 1.7 to 2.3 across consumer types, roughly matching the magnitude (in absolute terms) of the utility from

accepting cookies relative to rejecting them. The additional cost of selecting “accept all” from within the settings menu ($C_{set-acc}$) is lower but still meaningful, highlighting that even small added steps can discourage action. For always acceptors, this cost estimate implies they face lower friction when clicking on settings and choosing to accept, compared to the friction they would face when attempting to reject cookies.

Visual salience effects are more modest. Ranking an option at the top (δ_{top}) increases its utility, particularly for acceptors, while graying out options (δ_{grey}) generally reduces their utility, most strongly for rejectors.

Finally, the utility shifter for closing the banner without making a choice (κ) is lowest for rejectors and discerners, suggesting they perceive this option as especially unattractive. Still, its magnitude is smaller than C_{set} , implying that closing the window is less costly than navigating to the settings menu. The risk aversion parameter (ρ) is large and negative for discerners, consistent with ambiguity aversion among users who are uncertain about the consequences of closing the banner.

4.1. Consumer Surplus under Counterfactual Policies

The structural estimates allow us to evaluate the welfare implications of different consent policies. We find that choice frictions meaningfully deter users from selecting their preferred option, motivating policies that aim to reduce these barriers. One approach is to mandate neutral interfaces free of deliberate obstruction. In such cases, the website’s default behavior becomes critical, as many users close the banner without making a choice. We therefore compare neutral designs with both share-all and reject-all defaults. A second approach replaces site-by-site consent with a browser-level setting (“global privacy control”)²⁷, which reduces interaction costs but limits customization, making its welfare impact ambiguous.

Specifically, we consider the following proposals:

1. **US status quo:** An “acc-set” design, with websites defaulting to collect all cookies upon inaction (i.e., when the user closes the banner);
2. **EU norm:** A neutral interface, with websites defaulting to no cookie tracking upon inaction;
3. **Consumer welfare-maximizing banner:** A neutral interface, with websites defaulting to collect all cookies upon inaction;
4. **Global Privacy Control (Sharing):** A browser-level choice that automatically accepts all cookies across websites;

²⁷<https://globalprivacycontrol.org>.

5. **Global Privacy Control (Not Sharing):** A browser-level choice that automatically rejects all cookies across websites (e.g., in line with COPPA and similar privacy rules, see Johnson et al. (2024)).

The correct welfare calculation needs to account for consumers' incorrect beliefs when they close the consent banner. To this end, we allow choice probabilities to be guided by consumer beliefs, while computing the actual utility using websites' default action upon consumer inaction. Based on Train (2015), consumer surplus under potentially mistaken belief can be expressed as follows:

$$CS = \log \left(\sum_{t(i), k \in \{acc, rej, set, close\}} e^{U_{t(i)k|policy}} \right) + \sum_{t(i), k \in \{acc, rej, set, close\}} Pr_{t(i)k|policy} \cdot d_{t(i)k|policy},$$

where $U_{t(i)k|policy}$ is the perceived utility associated with different options under a given policy, $Pr_{t(i)k|policy}$ is the probability of consumers of type $t(i)$ choosing option k , and $d_{t(i)k|policy}$ is the difference between the actual and the perceived utility.²⁸ We then calculate the surplus for each of the policies above.

For choices under the US status quo, utilities are based on the banner that hides “reject all.” Since the default upon closing the window is sharing, we adjust the welfare calculation for consumers who believe the default is rejection or are unsure. Dropping the notation that all values are conditional on the “acc-set” banner, the second term in the consumer surplus formula simplifies to:

$$\sum_{t(i), k \in \{acc, rej, set, close\}} Pr_{t(i)k} \cdot d_{t(i)k} = \sum_{t(i)} Pr_{t(i)close} \cdot \left[U_{t(i)acc} - Pr(b_i = 1|t(i)close)U_{t(i)acc} - Pr(b_i = 0.5|t(i)close)(0.5 + \rho)U_{t(i)acc} \right].$$

For consumer surplus under the EU norm, utilities are based on the banner without dark patterns, with a no-sharing default. Therefore, we need to adjust the welfare calculations for the consumers who think the default is accepting or are unsure:

$$\sum_{t(i), k \in \{acc, rej, set, close\}} Pr_{t(i)k} \cdot d_{t(i)k} = \sum_{t(i)} Pr_{t(i)close} \cdot \left[-Pr(b_i = 1|t(i)close)U_{t(i)acc} - Pr(b_i = 0.5|t(i)close)(0.5 + \rho)U_{t(i)acc} \right].$$

²⁸It is important to note that we consider the error term associated with closing the window as a separate i.i.d. draw, even though the choice is perceived as resulting in the same outcome as one of the existing alternatives.

For the consumer welfare-maximizing banner, we use the banner without dark pattern, with a sharing default, which applies a correction to consumer surplus that is similar to the US status quo, except that utilities (and thus choices) are based on the banner without any dark patterns.

The surplus under global privacy controls reflects two components: the utility loss from not being able to customize cookie preferences, and the gain from avoiding repeated interactions with consent banners. To capture this, consumer surplus is set equal to $U_{t(i)k}$, where $k \in \{\text{accept, reject}\}$ corresponds to the browser-level choice in each of the two counterfactuals. We then add the time savings from bypassing banner interactions. Treating C_{set} as a proxy for the time cost of clicking into settings, the time-saving benefit is calculated as $C_{\text{set}} \times (\text{average time spent per banner} / \text{time to click "settings"})$.²⁹

Lastly, to express welfare estimates in dollar terms, we must make assumptions about how to map utility to monetary values, since our model does not include prices. This translation is inherently challenging: our data do not provide direct estimates of consumers' willingness to pay for privacy or for avoiding additional clicks. Instead, we draw on external sources to assign a dollar value of time. We then use these external benchmarks to scale our welfare estimates, leveraging the fact that the relative magnitudes of consumer surplus and friction parameters are identified within our model.

To estimate the monetary value of C_{set} , we perform a back-of-the-envelope calculation using observed time spent and the opportunity cost of time. Appendix Table A18 quantifies the additional time users take when they choose to "accept all" or "reject all" when their preferred option is hidden through deliberate obstruction. On average, acceptors take the longest time, 4.9 seconds followed by discerners at 2.9 seconds, and rejectors at 2.5 seconds. To estimate the opportunity cost of time, we use participants' self-reported annual household income,³⁰ and assign an implied hourly wage based on the midpoint of each income bracket. We use the number of domains visited per week to convert C_{set} to a dollar value. So, our welfare measures are per user per week.

To estimate the monetary value of C_{set} , we perform a back-of-the-envelope calculation using observed interaction times and the opportunity cost of time. Appendix Table A18 reports the additional time users spend when selecting "accept all" or "reject all" after their preferred option is hidden through deliberate obstruction. On average, acceptors take the longest (4.9 seconds), followed by discerners (2.9 seconds), and rejectors (2.5 seconds).

²⁹For time spent per banner, we use the average interaction time during the organic phase under the neutral design, which reflects the time cost of making a choice absent deliberate obstruction.

³⁰During the intake survey, participants report their household income over the past 12 months, selecting from six predefined brackets: \$0–25k, \$25–50k, \$50–75k, \$75–100k, \$100–150k, and \$150k or more.

To approximate the opportunity cost of time, we use participants’ self-reported annual household income,³¹ and assign an implied hourly wage based on the midpoint of each bracket. We then multiply these wage estimates by the average number of unique domains visited per week (54.45 domains) to convert C_{set} into a weekly dollar value. Accordingly, all welfare estimates are expressed on a per-user, per-week basis. For example, a participant in the \$75–100k income bracket (\$42 implied hourly wage) who is classified as a rejector would incur a weekly time cost C_{set} of \$1.60, calculated as $(2.5 \times 42 \times 54.45)/3,600$. This procedure yields 18 distinct values for the cost of C_{set} , one for each combination of user type (acceptors, rejectors, discerners) and income bracket (six levels). Consumer surplus estimates are then computed as weighted averages across these groups.

Table 7 reports our consumer surplus estimates in dollar terms. The first column shows pooled results, while the remaining columns allow for heterogeneity across the three consumer types. For ease of comparison, we also report the corresponding welfare values in utility units in Appendix Table A15. Although the per-choice effects are relatively small, our back-of-the-envelope calculations suggest that the cumulative impact of improved interface design on consumer welfare can be substantial.

TABLE 7. Consumer Surplus Under Counterfactual Policies (\$, Weekly)

Counterfactual	Pooled Estimate	Accepters	Rejectors	Discerners
US Status Quo	2.85	3.77	-0.57	1.53
EU Norm	2.92	3.52	0.56	1.76
CS Maximizing Banner	3.16	4.01	0.38	1.81
Global Accept All	5.97	6.82	3.76	5.24
Global Reject All	4.65	4.14	5.2	4.79

Notes: The table reports scaled consumer surplus dollar-value estimates under various counterfactual policies. The US status quo refers to an *acc-set* interface, combined with an “accept all” default when consumers close window; the CS maximizing banner refers to a neutral interface with an “accept all” default when consumers close window; the EU norm refers to a neutral interface with a “reject all” default when consumers close window. Global accept all and global reject all forces each individual to either always accept or always reject all cookies. The pooled estimate column refers to the estimate across all subjects, and the other columns correspond to estimates by user type.

Our welfare estimates show that, relative to the US status quo, the consumer surplus-maximizing policy, which adopts a neutral interface and defaults to accepting cookies upon inaction, increases average surplus by 13%, with much larger gains for rejectors. These welfare improvements stem from reducing choice frictions while aligning defaults with the preferences of the average user.

³¹During the intake survey, participants report their household income over the past 12 months, choosing from six predefined brackets: \$0–25k, \$25–50k, \$50–75k, \$75–100k, \$100–150k, and \$150k or more.

Defaulting users to rejecting cookies upon inaction has two opposing effects on consumer welfare. On one hand, it may improve outcomes for users who expect or prefer a reject-all default. On the other hand, it may reduce welfare for the majority of users who prefer to share data. Our counterfactual analysis suggests that the net effect is negative. Holding the interface design fixed, the EU Norm, defined as a neutral interface paired with a reject-all default, results in 8% lower consumer surplus compared to the welfare-maximizing policy. As expected, this policy benefits rejectors, who gain from having the default aligned with their preferences. This policy still outperforms the US status quo, increasing average surplus by 4%, due to its use of a neutral design that reduces choice friction.

Remarkably, both global privacy control policies yield higher consumer surplus compared to even the optimal consent-based policy. In particular, both acceptors and discerners benefit the most from a global accept policy, with surplus gains increasing by 70-190% compared to the optimal consent-based policy, while rejectors benefit the most from a global reject policy, enjoying a surplus of \$5.2 per week compared to \$0.62 under the EU norm, which is their preferred banner policy. We note that global choice benefit even the discerners, who would customize data sharing decisions by website when site-specific banners are given. This is partly because they spend relatively more time on each banner compared to other groups, thus the benefit of saving time outweighs their cost of not customizing cookie choices by site.

We also conduct additional sensitivity analysis for the counterfactual policy evaluation. First, we compute the time spent per banner for the optimal global control policy (global accept) to yield the same surplus as the optimal consent-based policy (CS-Max banner). This time amounts to 2.45s, which is 33% of their current average time spent for each banner. Second, we validate these policy evaluation results using participants' stated preferences for global privacy control tools in their endline responses. 62% of our participants say they believe having such a tool is better than site-by-site choice, while only 13% say it would be worse (Appendix Figure A3). The average valuation for such a tool is at \$4.00. This number lies within our estimated surplus gain moving from the US Status Quo to each segment's favorite global privacy control option: accepters would increase their surplus by \$3.03, rejectors by \$5.30, and discerners by \$3.48. While stated preferences may not exactly coincide with revealed preferences, the similarity in their magnitudes gives reassurance that our surplus estimates are sensible.

Overall, these results indicate substantial welfare gains from allowing users to pre-specify global privacy preferences that match their individual data-sharing attitudes.

5. Conclusions

In this paper, we examine the impact of dark patterns on consumer privacy choices through a field experiment that randomizes cookie consent interface designs while participants browse the internet. We find that the most effective form of choice architecture involves concealing options- i.e., placing “reject all” behind an additional click- which substantially decreases the selection of the hidden option and instead prompts consumers towards abandoning consent choice. In contrast, visual manipulations such as reordering or highlighting have more modest effects.

While better-known websites receive slightly higher consent rates on average, we do not find that dark patterns systematically amplify this advantage. If anything, these patterns reduce users’ baseline inclination to share more data with popular sites. This suggests that dark patterns are unlikely to harm the competitive position of smaller firms. In fact, if the marginal value of data is higher for newer or smaller players, dark patterns could even intensify competition.

We also show that consent-based privacy regimes impose real-time costs. Interacting with cookie banners across websites takes roughly 6.6 minutes per week, translating to a time cost between \$3.99 and \$7.64 depending on the benchmark value of time. Designs that obstruct consent can increase this cost by up to 70%. Thus, deliberate obstruction reduces welfare not only by distorting choices but also by increasing cognitive and time burdens.

To quantify these welfare losses, we estimate a structural model of privacy choices and evaluate alternative policy scenarios. We find that a neutral interface displaying all three choices upfront increases consumer surplus by 13% relative to the U.S. status quo, which hides rejection behind settings. A browser-level consent tool, where users specify their preferences once, further improves welfare by at least 73%, largely by eliminating repetitive, site-by-site interactions.

We observe relatively high consent rates even in the absence of dark patterns, raising the question of why consumers are generally willing to share their data. Possible explanations include a desire for personalized services, convenience, or concern about degraded website functionality when rejecting cookies. Exploring these motivations in greater depth remains an important direction for future work.

Our study has several limitations. First, we rely on a revealed preference approach, which assumes that observed choices reflect underlying utility. This assumption may not fully hold if users lack a clear understanding of the consequences of accepting or rejecting cookies. To partially address this, our welfare analysis treats closing the banner

as a belief-driven action, guided by the user’s stated expectations, but computes welfare based on the actual default implemented upon inaction. While this approach helps bridge the gap between perceived and actual outcomes, it does not fully eliminate uncertainty. Nonetheless, our endline survey indicates that most participants understand at a high level that cookies enhance website functionality and enable targeted advertising.

Second, our experimental setup, which uses a browser extension to recruit study participants online, may not perfectly replicate the natural browsing experience, especially given that participants knew they were being studied and selected into participating. Third, our analysis captures only short-term behavior. Longer-term dynamics, such as adaptation, reduced engagement, or shifting attention over time, remain unexplored. Finally, our findings are specific to cookie banners and may not generalize to other privacy contexts or types of dark patterns. We leave these important questions to future research.

References

- 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).
- Bian, Bo, Michaela Pagel, Huan Tang, and Devesh Raval. 2023. “Consumer surveillance and financial fraud.”, National Bureau of Economic Research.
- 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).
- 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 Paternolli. 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.
- Greminger, Rafael P, Yufeng Huang, and Ilya Morozov. 2023. "Make Every Second Count: Time Allocation in Online Shopping." *Available at SSRN* 4643819.
- 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."
- Johnson, Garrett, Tesary Lin, James C Cooper, and Liang Zhong. 2024. "COPPAocalypse? The Youtube Settlement's Impact on Kids Content."
- Ke, T Tony, and K Sudhir. 2023. "Privacy rights and data security: GDPR and personal data markets." *Management Science* 69 (8): 4389-4412.
- 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.
- Miller, Amalia R, and Catherine Tucker. 2018. "Privacy protection, personalized medicine, and genetic testing." *Management Science* 64 (10): 4648-4668.
- Müller-Tribbensee, Timo, Klaus M Miller, and Bernd Skiera. 2024. "Paying for Privacy: Pay-or-Tracking Walls." *arXiv preprint arXiv:2403.03610*.
- 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. 2019. "The value of privacy: Evidence from online borrowers." *Available at SSRN*.
- 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.
- Train, Kenneth. 2015. "Welfare calculations in discrete choice models when anticipated and experienced attributes differ: A guide with examples." *Journal of choice modelling* 16: 15-22.
- 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.
- Zhao, Yu, Pinar Yildirim, and Pradeep K Chintagunta. 2021. "Privacy regulations and online search friction: Evidence from GDPR." *Available at SSRN 3903599*.

Appendix A. Effects on Consumer Choice by Type of Dark Pattern

This appendix presents reduced-form results on the effects of specific types of dark patterns on consumer choices.

To quantify the effect of dark patterns on privacy choices, we identify the presence of reordering, highlighting, and obstructions in our treatments. This means estimating the following regression:

$$y_{ij} = \gamma_{reject\ hidden_{ij}} + \gamma_{accept\ hidden_{ij}} + \gamma_{accept\ on\ top_{ij}} + \gamma_{reject\ on\ top_{ij}} + \gamma_{highlight\ accept_{ij}} + \mu_i + \nu_{c(j)} + \epsilon_{ij}, \quad (A1)$$

where i indexes study participants and j indexes websites. We include participant fixed effects μ_i and website category fixed effects $\nu_{c(j)}$, where categories are based on Ghostery’s taxonomy³² and extended to the full set of websites in our study using ChatGPT.³³

The γ indicators capture the specific nudges present in the consent interface for user i on website j : which option is placed at the top (accept or reject), which option is hidden (accept or reject), and whether the “accept all” option is visually highlighted. To interpret these coefficients relative to our treatment conditions in Figure 1, $\gamma_{accept\ on\ top_{ij}}$ represents the effect of placing “accept all” first, relative to the neutral interface (Treatment C), while $\gamma_{reject\ on\ top_{ij}}$ captures the analogous effect for prioritizing “reject all.” The remaining γ coefficients reflect the incremental impact of additional nudges—hiding or highlighting—beyond the positioning of “accept all” or “reject all” at the top of the banner.

We focus on three outcomes: accepting all cookies, rejecting all cookies, and closing the window without making an active choice. Table A1 presents the results. Columns 1–3 jointly capture how each type of dark pattern affects substitution across consent choices in the survey phase; columns 4–6 provide the corresponding analysis for the organic phase.

Deliberate obstruction emerges as the most effective dark pattern. Hiding the “reject all” button from the main screen reduces rejection rates by 17 percentage points in the survey phase (a 71% decrease) and by 9 percentage points in the organic phase (a 66% decrease). In the survey phase, participants shift both to accepting cookies, up 6.3 percentage points (a 10% increase), and to closing the window, which rises by 4.9 percentage points

³²<https://www.ghostery.com/>.

³³We used ChatGPT 4o to classify the websites. We used the following prompt: *Classify the website domains listed below into one of the following major categories (and only one of the following, do not include categories not in this list and try to limit how often other is selected): 'Reference Website', 'Entertainment Website', 'Business Website', 'E Commerce Website', 'Adult Website', 'News and Portals Website', 'Recreation Website', 'Banking Website', 'Government Website', 'Political Website', 'Other'.*

TABLE A1. 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	0.063*** (0.013)	-0.171*** (0.015)	0.049*** (0.008)	0.017 (0.012)	-0.093*** (0.012)	0.060*** (0.012)
Accept Hidden	-0.467*** (0.020)	0.179*** (0.017)	0.246*** (0.018)	-0.431*** (0.023)	0.065*** (0.012)	0.325*** (0.023)
Accept Top	0.020 (0.011)	-0.007 (0.009)	-0.010 (0.006)	0.038** (0.011)	-0.001 (0.007)	-0.041*** (0.010)
Reject Top	0.003 (0.010)	0.014 (0.010)	-0.012* (0.006)	0.004 (0.011)	0.020* (0.009)	-0.023* (0.010)
Highlight Accept	0.015 (0.010)	-0.010 (0.010)	-0.002 (0.006)	-0.007 (0.011)	-0.014 (0.007)	0.021* (0.009)
Benchmark group mean	0.65	0.24	0.08	0.61	0.14	0.22
R ²	0.646	0.579	0.562	0.571	0.522	0.494
Observations	11,075	11,075	11,075	12,610	12,610	12,610
Participant fixed effects	✓	✓	✓	✓	✓	✓
Domain Cat. fixed effects	✓	✓	✓	✓	✓	✓

Notes: Regression results of Equation A1 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 A2 presents similar results for the decision to accept a subset of cookie types. * $p < 0.05$, ** $p < 0.01$; *** $p < 0.001$.

(a 61% increase). In the organic phase, the dominant response is to close the window, which increases by 6 percentage points (a 27% rise from an already high baseline). Hiding the “accept all” button has an even more pronounced effect, reducing acceptance rates by 47 percentage points in the survey phase (a 72% decline) and by 43 percentage points during organic browsing (a 71% decline). Participants shift primarily toward closing the window—53% do so in the survey phase and 75% in the organic phase—with a smaller share substituting toward rejecting cookies (38% in survey, 15% in organic).

In contrast, visual manipulations—i.e., reordering options and highlighting “accept all”—have more limited effects on user choices. Most coefficients are small and statistically insignificant, particularly in the survey phase. These nudges have somewhat larger effects during organic browsing, perhaps because users are not as focused on cookie preferences, but the magnitudes remain modest. For example, placing “accept all” at the top increases acceptance by only 3.8 percentage points (column 4, a 6% increase), while highlighting has no additional impact on acceptance. Similarly, placing “reject all” at the top marginally increases rejections by 2 percentages points (column 5, a 14% increase).

TABLE A2. Selective Cookie Choice by Dark Pattern

	Survey	Organic
	Accept Some (1)	(2)
Reject Hidden	0.059*** (0.010)	0.016* (0.007)
Accept Hidden	0.043*** (0.008)	0.041*** (0.008)
Accept Top	-0.003 (0.005)	0.005 (0.003)
Reject Top	-0.005 (0.005)	-0.002 (0.004)
Highlight Accept	-0.004 (0.004)	0.000 (0.003)
Benchmark group mean:	0.03	0.03
R ²	0.413	0.499
Observations	11,075	12,610
Participant fixed effects	✓	✓
Domain Cat. fixed effects	✓	✓

Notes: The table regressions of Equation A1, where the outcome is whether the user selects a subset of cookies. Otherwise the table is identical to Table A1. * $p < 0.05$, ** $p < 0.01$; *** $p < 0.001$.

Table A2 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 3% 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, increasing the probability of granular choices by 2-6 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 A9). This result suggests that targeted advertising is the least preferred use of consumer data, at least among the few users who make selective choices.

Appendix B. 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 when we increase the frequency of pop-ups.

The 10-minute treatment sees our banners in 53% of the domains they visited, while the 60-minute treatment sees these banners in 30% of the domains. 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.

$$y_{ij} = \beta_{10 \text{ minutes}} + \gamma * \text{time in study}_{ij} + \nu_{c(j)} + \epsilon_{ij}. \quad (\text{A2})$$

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.³⁴

Table A3 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. These null effects are precisely estimated, as the 95% confidence interval excludes effects greater than 7%. However, we acknowledge the caveat that the difference between exposing to banners 30% vs. 53% of the time may not be large enough compared to, say, comparing banner exposure between 30% and 100% of the time.

Second, time spent in the study has an effect 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, this implies that they are 14 percentage points more likely to close the window at the end of the study compared to the first day.

It is tempting to directly interpret the time in the study as another measure of choice fatigue, but it is not randomly allocated and could be correlated with underlying consumer characteristics and privacy preferences. To address this concern, in Appendix Table A16 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.

³⁴Adding this covariate does not affect whether we detect any treatment effects.

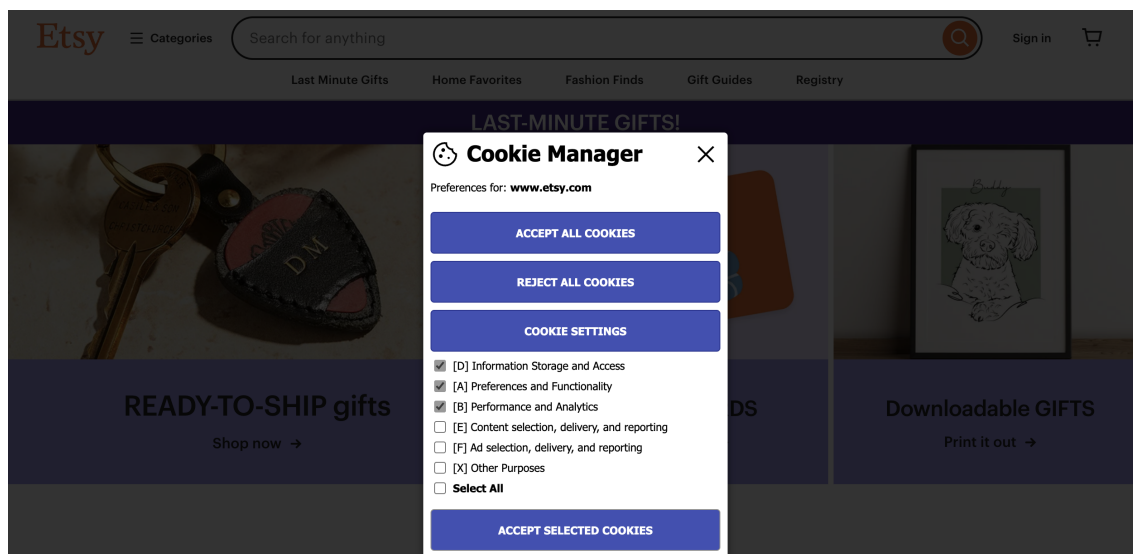
TABLE A3. Fatigue in Cookie Choices During Organic Browsing

	Accept All (1)	Reject All (2)	Close Window (3)
10 Min Pop-up	0.009 (0.037)	-0.005 (0.024)	-0.003 (0.031)
Time in Study (Days)	-0.009* (0.004)	-0.008** (0.003)	0.017*** (0.004)
Domain Rank (Log 10)	0.010** (0.004)	-0.008** (0.003)	-0.002 (0.003)
Pre-Exp Visit	0.041** (0.015)	-0.031* (0.013)	-0.015 (0.013)
R ²	0.008	0.007	0.009
Observations	12,610	12,610	12,610
Domain Cat. fixed effects	✓	✓	✓

Notes: This table shows estimates of Equation A2, 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. * $p < 0.05$, ** $p < 0.01$; *** $p < 0.001$.

Appendix C. Additional Tables and Figures

FIGURE A1. Cookie Manager's User Interface



Notes: Consent interface when a user clicks “cookie settings”. Note that the available blue buttons vary by treatment: for example, in the “accept-settings” design, the blue “Reject All Cookies” button does not show up when a user clicks “settings.” However, the list of specific cookies within the settings menu is identical across the six treatments, meaning they can always choose to reject or select all cookies upon coming to the “settings” menu.

TABLE A4. Websites Browsed during the Survey Phase

Domain	Domain Rank
facebook.com	3
youtube.com	8
amazon.com	28
yahoo.com	41
ebay.com	185
weather.com	325
duckduckgo.com	413
target.com	631
espn.com	278
etsy.com	301
nytimes.com	119
appleinsider.com	6319
seattletimes.com	3349
stockx.com	4547
funnyordie.com	16437
turo.com	16272
semafor.com	28266
thomannmusic.com	90809
truewerk.com	348372
merrysky.net	1000001

TABLE A5. Number of Participants across the Experimental Funnel

Stage	N	Percent	10 min	Percentage	60 min	Percentage
1) Start Survey	1227	100				
2) Eligible for Study	1227	100				
3) Study Consent	877	71.48				
4) Finished Survey	807	65.77				
5) Clicked All Links	808	65.85	359	100.00	418	100.00
6) Have Cookie Choice Data	767	62.51	350	97.49	410	98.09
7) After 15+ Domains Filter	687	55.99	316	88.02	371	88.76
8) After Mutual Presence Filter	602	49.06	282	78.55	320	76.56
9) Main Analysis Sample*	563	45.88	260	72.42	303	72.49
10) Finished Endline Survey	484	39.45	218	60.72	266	63.64

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 A6. 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: Banner design is randomized at the user X site level. * $p < 0.05$, ** $p < 0.01$; *** $p < 0.001$.

TABLE A7. 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: Banner frequency is randomized at the user level. We therefore exclude domain rank during the organic browsing, but include the number of banners exposed at the survey stage for covariate balance checks. * $p < 0.05$, ** $p < 0.01$; *** $p < 0.001$.

TABLE A8. Selective Cookie Choice by Experimental Condition

	Survey Accept Some (1)	Organic Accept Some (2)
Acc-Set	0.056*** (0.011)	0.021** (0.007)
Acc-GrRej-GrSet	-0.007 (0.004)	0.005 (0.003)
Acc-Rej-Set	-0.003 (0.005)	0.005 (0.003)
Rej-Acc-Set	-0.005 (0.005)	-0.002 (0.004)
Rej-Set	0.038*** (0.008)	0.039*** (0.007)
Benchmark group mean:	0.03	0.03
R ²	0.413	0.499
Observations	11,075	12,610
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 A9. Types of Cookies Accepted among People Making Granular Choices

Cookie Type	Percentage Selected
Preferences and functionality	0.826
Information storage and access	0.627
Performance and analytics	0.601
Content selection, delivery, and reporting	0.390
Ad selection, delivery, and reporting	0.070
Other purposes	0.048

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


TABLE A10. Heterogeneity of Dark Pattern Effect by Prior Visit

	Accept All (1)	Reject All (2)	Close Window (3)
Has Prior Visit \times Acc-GrRej-GrSet	0.011 (0.018)	-0.005 (0.013)	0.001 (0.015)
Has Prior Visit \times Acc-Rej-Set	0.017 (0.019)	-0.008 (0.014)	-0.008 (0.016)
Has Prior Visit \times Rej-Acc-Set	0.028 (0.020)	-0.009 (0.015)	-0.002 (0.015)
Has Prior Visit	0.015 (0.015)	-0.029* (0.012)	0.006 (0.012)
Has Prior Visit \times Acc-Set	0.007 (0.018)	0.038** (0.014)	-0.030 (0.016)
Has Prior Visit \times Rej-Set	0.011 (0.022)	-0.044* (0.018)	0.022 (0.020)
R ²	0.517	0.471	0.421
Observations	23,685	23,685	23,685
condition fixed effects	✓	✓	✓
Participant fixed effects	✓	✓	✓
Domain Cat. fixed effects	✓	✓	✓

Notes: The table shows regression estimates similar to Table 3, Panel *b*, except that the dummy for whether the participant visited the website in the days preceding the experiment is interacted with the banner design treatment dummies. “Condition fixed effect” refers to indicator variables for the 6 banner design conditions.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

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



Notes: The figure displays the treatment effects (point estimates and 95% confidence intervals) of the Acc-Set condition relative to the neutral interface for each domain included in the survey phase. Estimates are obtained from separate regressions of Equation 1, conducted individually for each domain in our survey. Each color (pink, blue, and grey) denotes a different outcome.

TABLE A11. Heterogeneity of Dark Pattern Effect by Domain Popularity

	Accept All (1)	Reject All (2)	Close Window (3)
Domain Rank (Log 10) \times Acc-GrRej-GrSet	0.000 (0.005)	-0.003 (0.004)	0.002 (0.004)
Domain Rank (Log 10) \times Acc-Rej-Set	-0.002 (0.005)	0.002 (0.004)	0.000 (0.004)
Domain Rank (Log 10) \times Rej-Acc-Set	-0.002 (0.005)	-0.001 (0.004)	0.001 (0.004)
Domain Rank (Log 10)	-0.010* (0.004)	0.007* (0.003)	0.003 (0.003)
Domain Rank (Log 10) \times Acc-Set	-0.007 (0.005)	0.004 (0.004)	0.005 (0.004)
Domain Rank (Log 10) \times Rej-Set	0.005 (0.006)	-0.015** (0.005)	0.013* (0.006)
R ²	0.518	0.470	0.422
Observations	23,685	23,685	23,685
condition fixed effects	✓	✓	✓
Participant fixed effects	✓	✓	✓
Domain Cat. fixed effects	✓	✓	✓


Notes: The table shows regression estimates similar to Table 3, Panel b, except that the domain rank (in logs) is interacted with the banner design treatment dummies. “Condition fixed effect” refers to indicator variables for the 6 banner design conditions. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

TABLE A12. Choice by Stated Privacy Choice Pattern

	Accept All (1)	Reject All (2)	Close Window (3)
Accepted Most Cookies	0.273*** (0.024)	-0.141*** (0.015)	-0.059* (0.023)
Rejected Most Cookies	-0.314*** (0.030)	0.363*** (0.038)	0.007 (0.036)
R ²	0.302	0.242	0.114
Observations	23,685	23,685	23,685
condition fixed effects	✓	✓	✓
Domain Cat. fixed effects	✓	✓	✓
Sample fixed effects	✓	✓	✓

Notes: This table presents regression estimates of cookie choices as in Equation 1, while adding participants' stated reasons for accepting or rejecting cookies, as reported in the endline survey. Each row corresponds to a binary indicator for a stated motivation (e.g., trust, functionality, distrust, unfamiliarity, privacy concerns). All regressions include fixed effects for interface condition, domain category, and study phase. Standard errors are clustered at the participant level. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

FIGURE A3. User Preferences for Global Privacy Controls



Notes: The figure plots the distribution of answers to the following question in the endline survey: 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.

TABLE A13. Choice by Stated Reasons to Accept/Reject

	Accept All (1)	Reject All (2)	Close Window (3)
Accepted for Trust	0.011 (0.038)	-0.032 (0.028)	-0.014 (0.027)
Accepted for Functionality	0.045 (0.044)	-0.018 (0.033)	-0.062* (0.029)
Rejected for Distrust	-0.048 (0.042)	-0.006 (0.028)	0.077** (0.030)
Rejected for Unfamiliarity	-0.092* (0.045)	0.007 (0.031)	0.106** (0.038)
Rejected for Privacy	-0.116* (0.049)	0.066 (0.045)	0.093** (0.034)
R ²	0.143	0.068	0.123
Observations	21,410	21,410	21,410
condition fixed effects	✓	✓	✓
Domain Cat. fixed effects	✓	✓	✓
Sample fixed effects	✓	✓	✓

Notes: This table presents regression estimates of cookie choices as in Equation 1, while adding participants' stated choices, as reported in the endline survey. The self-reported behavior aligns closely with actual choices: those who said they accepted most cookies are significantly more likely to accept and less likely to reject, while the reverse is true for those who reported rejecting most. All models include fixed effects for interface condition, domain category, and sample phase. Standard errors are clustered at the participant level. * $p < 0.05$, ** $p < 0.01$; *** $p < 0.001$.

TABLE A14. Differences Between Participants Who Always Accept, Always Reject, and Make Discerning Choices

Variable	Mean: Never Accept	Mean: Some Accept	Mean: Always Accept	F-value	p-value
Age	38.81	38.22	37.77	0.19	0.828
Female	0.48	0.41	0.45	0.61	0.546
Bachelor's or Above	0.26	0.12	0.17	3.43	0.033
Income > \$75,000	0.43	0.48	0.42	0.79	0.452
Prior Domains Visited	52.27	57.5	50.87	1.43	0.241

Notes: "Never Accept" and "Always Accept" indicate the participants who never and always choose "accept all", excluding instances of closing window. In this table, all choice probabilities are calculated within the neutral design treatment, thus choice variations in the "Some Accept" group reflect variation across sites, excluding the influence of dark patterns on choice.

TABLE A15. Consumer Surplus Under Counterfactual Policies (Utility Scale Results)

Counterfactual	Pooled Estimate	Accepters	Rejectors	Discerners
US Status Quo	1.92	2.87	-0.71	1.51
EU Norm	1.99	2.72	1	1.74
CS Maximizing Banner	2.16	3.1	0.71	1.8
Global Accept All	4.14	5.38	6.07	5.23
Global Reject All	3.22	3.26	8.39	4.78

Notes: The values represent unscaled consumer surplus *per choice* under various counterfactual policies in the original scale. “Pooled estimate” refers to the estimate across all subjects, and the other columns correspond to subset-specific estimates. “US status quo” refers to an accept-settings interface, combined with an accept-all default when consumers close window; “CS maximizing” refers to a neutral interface with an accept-all default when consumers close window; “EU norm” refers to a neutral interface with a reject-all default when consumers close window.

TABLE A16. Fatigue in Cookie Choices During Organic Browsing (Additional Fixed Effects)

	Accept All (1)	Reject All (2)	Close Window (3)
Visit Order / 10	-0.020* (0.009)	-0.004 (0.004)	0.023* (0.009)
Time in Study (Days)	0.000 (0.005)	-0.003 (0.003)	0.005 (0.005)
R ²	0.469	0.501	0.439
Observations	12,610	12,610	12,610
Domain Cat. fixed effects	✓	✓	✓
Participant fixed effects	✓	✓	✓
Hour fixed effects	✓	✓	✓

Notes: This table estimates a variant of Equation A2, which removes the banner frequency treatment and adds the order of which a domain is visited (“Visit Order”) and additional fixed effects. * $p < 0.05$, ** $p < 0.01$; *** $p < 0.001$.

TABLE A17. Effects on “Accept All”

	Base	With Covariates
Intercept	0.625*** (0.014)	0.510*** (0.072)
Acc-Rej-Set	-0.002 (0.008)	-0.005 (0.008)
Acc-Set	0.036*** (0.008)	0.033*** (0.008)
Rej-Acc-Set	-0.026** (0.008)	-0.030*** (0.008)
Rej-Set	-0.475*** (0.008)	-0.480*** (0.008)
Set-Acc-Rej	-0.031*** (0.008)	-0.032*** (0.008)
SD (Participant)	0.304	0.304
SD (Domain)	0.058	0.046
SD (Residual)	0.343	0.342
Num.Obs.	23685	23190

Notes: This table presents estimates of treatment effects models where random effects for participant and domain are included in the regression. The second column adds controls for website characteristics, demographics, and privacy beliefs presented in Table 3. The outcome is a dummy for whether a user accepts all cookies.

TABLE A18. Extra Time Spent When Clicking Settings

	Time Spent (Seconds)		
	Never Accept (1)	Some Accept (2)	Always Accept (3)
User Clicks Settings	2.507* (1.151)	2.882*** (0.803)	4.908*** (0.608)
Close Window	1.609 (1.816)	1.683** (0.595)	1.959*** (0.546)
Accept Selected	2.733* (1.186)	4.812*** (0.996)	4.927** (1.752)
R ²	0.087	0.103	0.105
Observations	2,702	7,534	12,606
Condition fixed effects	✓	✓	✓
Participant fixed effects	✓	✓	✓
Domain Cat. fixed effects	✓	✓	✓

Regression estimates: time spent to make a decision as a function of whether the user's action involved clicking settings. The baseline is the time users take to accept or reject cookies when those options are available on the main banner. The variable of interest (*User Clicks Settings*) is equal to 1 if a user's final choice includes partial cookie selection, if it is "accept all" when that option is hidden, or if it is "reject all" when that option is hidden. Treatment, user, and domain category fixed effects are included as controls. The observations include all choices in survey and organic data. Standard errors are clusters at the user level.

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

Appendix D. 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

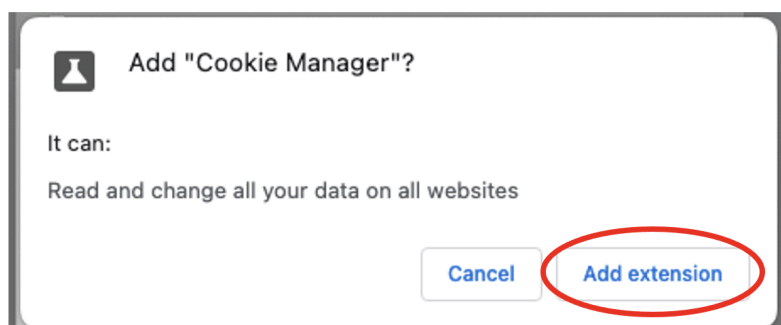
Cookie Manager Installation Instructions.

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


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



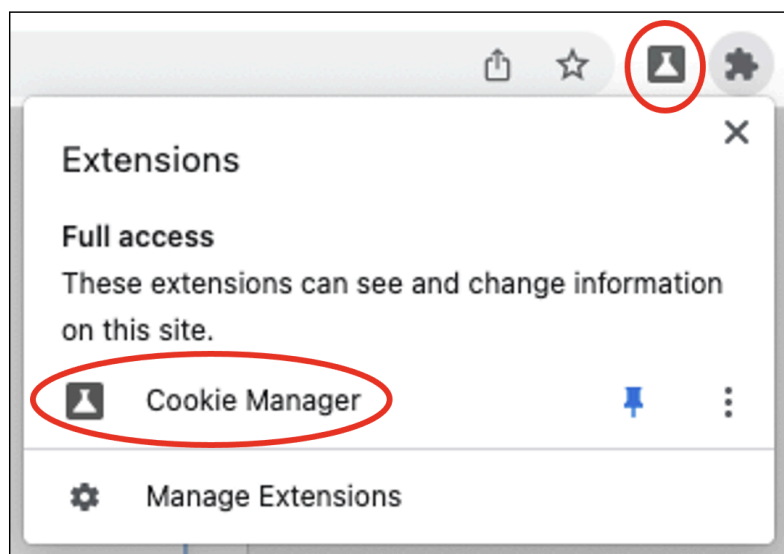
- When prompted, click “Add Extension.”



- 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

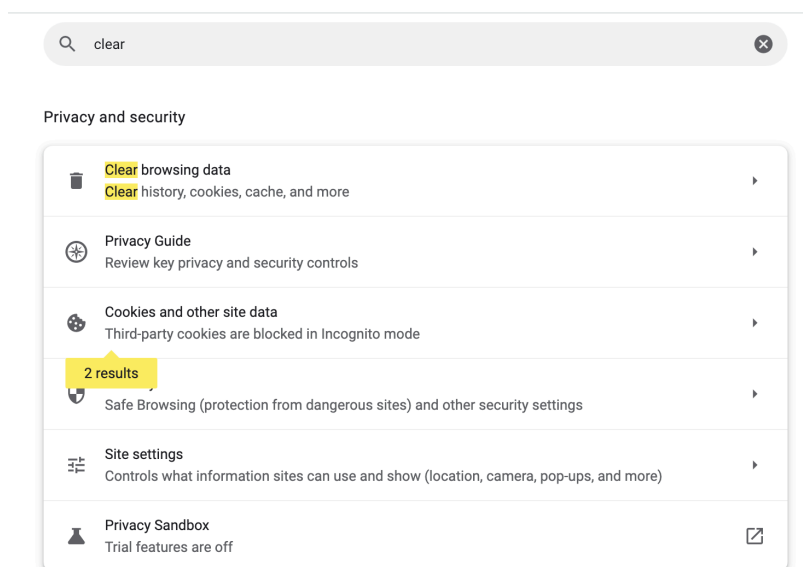
No

What difficulties have you encountered when installing the extension?

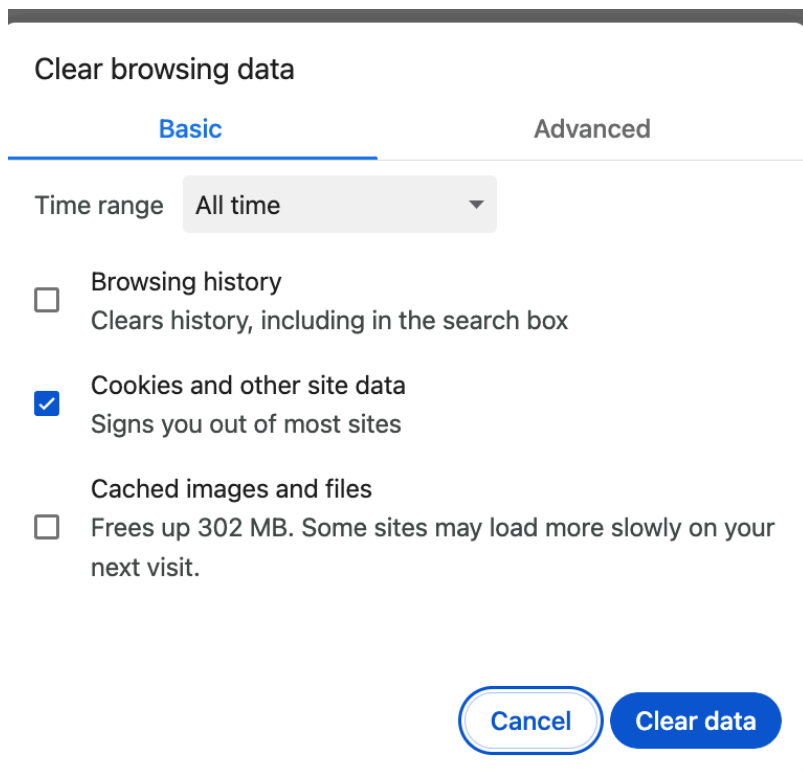
Clear Browsing History

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.

 **Cookie Manager** 

COOKIE SETTINGS

ACCEPT ALL COOKIES

REJECT ALL COOKIES

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 policies 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.

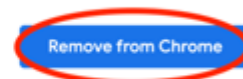


Now that you've clicked the cloud button, you can now proceed to uninstall the extension. Completing this step ensures that we stop collecting your browsing data going forward.


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