



Marketing Science Institute Working Paper Series 2026

Report No. 26-101

Designing Consent: Choice Architecture and Consumer Welfare in Data Sharing

Tesary Lin

“Designing Consent: Choice Architecture and Consumer Welfare in Data Sharing” © 2026

Tesary Lin

MSI Working Papers are Distributed for the benefit of MSI corporate and academic members and the general public. Reports are not to be reproduced or published in any form or by any means, electronic or mechanical, without written permission.

Designing Consent: Choice Architecture and Consumer Welfare in Data Sharing

January 27, 2026

Abstract

We study the welfare consequences of choice architecture for online privacy using a field experiment that randomizes cookie consent banners. We study three ways in which firms or policymakers can influence choices: (1) nudging users through banner design to encourage acceptance of cookie tracking; (2) setting defaults when users dismiss banners; and (3) implementing consent decisions at the website versus browser level. Absent design manipulation, users accept all cookies more than half of the time. Placing cookie options behind extra clicks strongly influences choices, shifting users toward more easily accessible alternatives. Many users dismiss banners without making an explicit choice, underscoring the importance of default settings. Survey evidence further reveals substantial confusion about default settings. Our structural estimates show that among consent policies requiring site-specific decisions, consumer surplus is maximized when consent interfaces clearly display all options and default to acceptance in the absence of an explicit choice. However, the welfare gains from optimizing banner design are smaller than those from adopting browser-level consent, which eliminates the time costs of repeated decisions.

Keywords: online privacy, choice architecture, consumer welfare, field experiments, consumer data

1. Introduction

Consumer data is a key input into the digital economy, powering everything from algorithmic recommendations to targeted advertising. Yet the collection and use of this data have sparked privacy concerns, prompting regulators from the European Union to California to demand that firms obtain explicit user consent before data collection and use. In response, firms now regularly solicit consent to use consumer data, most commonly through cookie consent banners.

Data consent requirements are controversial for at least three reasons. First, firms routinely design interfaces that steer users toward broad data collection and use, 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 consent, or using differential visual salience. Second, requiring each firm to obtain user consent can lead to frequent consent banners, disrupt the user experience, and increase cognitive burden. Lastly, consent requirements risk benefiting large and prominent firms due to consumers' higher propensity to consent to data use when interacting with familiar brands.²

In this paper, we examine how three key features of choice architecture (the design of consent interfaces, the use of default settings, and the frequency of consent prompts) shape online cookie consent decisions and influence consumer welfare. We conduct a large-scale field experiment in which we randomize cookie consent interfaces presented to users as they visit thousands of websites over one week. Our data captures both organic browsing behavior and website visits

¹<https://www.deceptive.design/>, accessed November 2025.

²For example, the FTC's report to OECD 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), accessed November 2025.

prompted by survey participation. The induced variation allows us to identify the causal effects of banner design on user consent decisions. We then estimate a structural model of consumer choice and use it to quantify welfare effects. We evaluate counterfactual policies that regulate design elements, set default settings, and mandate a uniform browser-level choice instead of site-by-site choices.

Our results reveal considerable differences in the effectiveness of choice architecture. Deliberate obstruction effectively deters users from choosing hidden options, whereas purely visual manipulations like reordering or highlighting with different colors have minimal effects. More popular websites achieve higher consent rates. However, choice architecture does not amplify the data-collection advantages of these sites. Users exhibit considerable heterogeneity in their baseline data-sharing preferences, and many frequently close banners without making a choice, underscoring the critical role of website default settings.

We embed these observed choices in a model of consumer decision-making. Using data on interaction times with consent banners, combined with an assumption about the value of time, we recover welfare in dollar terms under alternative consent policies. Our welfare analysis shows that policies designed to reduce choice burdens consistently improve consumer welfare. With site-by-site consent interactions, the optimal consent interface is one that presents all options without obstruction and defaults to “accept all cookies” when users close the banner without making a choice. This design substantially increases consumer surplus compared to the most common design in the U.S., which hides rejection options while also defaulting users to accepting cookies. A browser-level consent mechanism delivers the largest welfare gains even compared to the optimal banner design, because the time cost of repeated consent interactions substantially outweighs the benefit of site-specific customization.

Finally, to assess external validity, we conducted a follow-up survey across three groups: our main study participants, an independent U.S. sample, and a European sample. Privacy attitudes, preferences for consent interfaces, and willingness-to-pay measures are consistent across the three groups, reinforcing that our welfare estimates may apply more broadly than the experimental population.

Our experiment is enabled by Cookie Manager, a customized browser extension built on the Webmunk framework (Farronato, Fradkin and Karr 2024). The extension randomly assigns different consent interfaces to users while they browse the web, and enforces their decisions whenever technically feasible,³ making user choices incentive-compatible. We use Prolific to recruit U.S. consumers who consent to install Cookie Manager.

Our experiment randomizes six different interfaces across both users and web domains, varying the ease of choosing among three options: accept all cookies, reject all (non-essential) cookies, and customize preferences by cookie purpose. We test three types of choice architectures. *Deliberate obstruction* removes an option from the main banner, forcing users to make additional clicks to select this option. *Reordering* emphasizes an option by placing it at the top of the choice list. *Differential highlighting* uses different colors to make an option more salient than the others. Lastly, our treatments include a design free of any choice architectures tested (*baseline interface*), which allows us to observe consumer choices when not deliberately nudged. Together, these designs allow us to identify both preferences over cookie tracking and the effects of interface design on consumer choice.

³The extension identifies common HTML, CSS, and textual patterns indicating options to accept or reject cookies. It then selects those options detected based on user choice.

The study includes two phases. In the first phase (*survey browsing*), we direct participants to visit specific websites. This structured browsing allows us to evaluate users' privacy preferences across the same set of websites, regardless of whether they would visit them organically. In the second phase (*organic browsing*), we observe participants' natural browsing behavior for a week. In both phases, we randomize consent interface designs as users encounter new web domains.⁴

We find that when participants encounter the *baseline interface*, 65% accept all cookies during survey browsing, and 61% do so during organic browsing. Deliberate obstruction has a substantial influence on privacy choices, while visual changes have modest effects. Hiding the “reject all cookies” option reduces cookie rejection rates by 17.8 percentage points in survey visits and 9.4 percentage points in organic browsing. The sizable effect of deliberate obstruction is consistent with websites' strategic choices: deliberate obstruction is the most commonly used design manipulation in consent interfaces (Utz et al. 2019). In comparison, re-ordering options to prioritize “accept all cookies” only increases consent rates by up to 3.8 percentage points. Graying out options has similar null-to-small impacts.

Perhaps surprisingly, the effect of choice architecture does not vary substantially with website characteristics such as popularity or user familiarity. While users are more likely to share data with familiar or popular sites, choice architecture does not increase this tendency. If anything, it slightly reduces the advantage that popular sites enjoy during survey browsing. These findings challenge the hypothesis that choice architecture increases entry barriers or amplifies data-

⁴In the organic phase, we also randomly vary banner frequency: 50% of users see consent banners at most every 10 minutes, whereas the other 50% of users see consent banners at most every 60 minutes. This randomization allows us to test for choice fatigue induced by more frequent popups. We find no significant effects of banner frequency on consent choices (see Appendix B for details).

enabled network effects (Hagiu and Wright 2023), which would otherwise reinforce incumbent advantages in the data economy.

Beyond creating choice distortions, consent interfaces impose significant time costs on users. On average, participants spend 7.34 seconds interacting with a baseline banner during organic browsing. Extrapolating this number to scenarios where consent banners are present on every domain and a user's time is valued at \$36/hour (the average U.S. hourly wage), we estimate that consent interactions cost a user \$4 per week on average.⁵ Having to make an additional click to access options in the "settings" menu increases the time cost by over 50%.

Closing the banner without making a selection is the second most common choice during organic browsing, after accepting all cookies. Yet, 52.2% of participants mistakenly believe that doing so results in rejecting cookies, even though most websites default to acceptance in the U.S. This mismatch leads to discrepancies between users' intended and actual data-sharing outcomes, an issue we explicitly incorporate in our welfare analysis.

To quantify welfare implications, we estimate a model of consent decisions under choice architecture using the experimental variation in interface design. We account for consumer heterogeneity by estimating a distribution over latent user types. The welfare calculation requires us to overcome two challenges: capturing how incorrect beliefs translate into choice mistakes and welfare losses, and translating our utility parameters into dollar values. To tackle the first challenge, we leverage the fact that our survey measures participants' beliefs about the consequences of closing the banner, and incorporate these mistaken beliefs directly into the surplus calculation following Train (2015). To tackle the second challenge, we assume that the disutility from engaging with consent banners is proportional

⁵Our participants interact with 54 unique web domains on a given week.

to the time users spend on them. We then use each user's implied hourly wage, based on their reported income bracket, to convert this time into an opportunity cost in dollar terms.

Our welfare analysis allows us to compare regulatory approaches to online data sharing. One common approach is identifying and banning design manipulations, as mandated by the EU's Digital Services Act.⁶ In this regime, website defaults could matter, as many users close banners without making an explicit choice. We find that combining manipulation-free interfaces while defaulting consumers to the preference of the majority maximizes consumer surplus among banner-based approaches. This policy improves welfare by \$0.6 per user-week compared to current U.S. practices, where strategic choice architectures are predominant. Clarifying defaults, so that consumers have correct beliefs about the implications of closing consent banners, would further increase welfare by \$0.2 per user-week.

An alternative policy approach involves replacing site-by-site consent with browser-level privacy settings, commonly known as Global Privacy Control.⁷ The browser-level choice outperforms consent-based policies by a wide margin, improving consumer welfare by \$3.7 per user-week compared to the optimal site-by-site banner with correct consumer beliefs. The welfare gain primarily comes from removing time costs associated with repeated choice. Interestingly, even users with varying preferences across websites would benefit from a browser-level choice. This result is supported by participants' stated preferences: 62.4% prefer browser-level control tools to interacting with cookie banners on each site. This

⁶https://commission.europa.eu/strategy-and-policy/priorities-2019-2024/europe-fit-digital-age/digital-services-act_en.

⁷<https://oag.ca.gov/privacy/ccpa/gpc>

finding indicates that simplifying choice improves consumer welfare compared to approaches that impose substantial time costs on consumers.

These findings speak directly to the current policy debate on choice architecture related to online privacy. In the EU, both the Digital Services Act and the Artificial Intelligence Act regulate choice architecture;⁸ in the U.S., several states and the Federal Trade Commission have taken action against deceptive interface designs.⁹ Our results support these efforts: designs that deliberately obstruct choices reduce welfare and undermine meaningful choice. However, our analysis also indicates that design regulations alone are insufficient. Browser-level consent mechanisms—already implemented by Mozilla and supported by the California Privacy Rights Act—could deliver large additional welfare gains while preserving user choice.¹⁰

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; Que 2017; Acquisti, John and Loewenstein 2013) and examine the impact of data collection practices on consumers (Goldfarb and Tucker 2011; Miller and Tucker 2018; Zhao, Yildirim and Chintagunta 2021; Bian et al. 2023) and firms (Johnson, Shriver and Goldberg 2023; Johnson et al. 2024; Goldberg, Johnson and Shriver

⁸[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)

⁹The Federal Trade Commission has fined Epic Games and Amazon for user interface designs that induce accidental purchases and obstruct subscription cancellation (<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>). U.S. states such as California, Colorado, and Connecticut have enacted privacy regulations that ban companies from manipulating interface designs to increase data collection. (https://insightplus.bakermckenzie.com/bm/technology-media-telecommunications_1/united-states-consumer-protection-regulators-set-sights-on-dark-patterns).

¹⁰See: <https://blog.mozilla.org/netpolicy/2021/10/28/implementing-global-privacy-control/>; <https://globalprivacycontrol.org/press-release/20201007.html>

2024; Deisenroth et al. 2024; Wernerfelt et al. 2025). Importantly, we highlight that privacy choices in our setting are shaped not only by consumers' preferences over 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, thereby offering a policy-relevant view of privacy trade-offs.

Our analyses relate to the recent theoretical work examining the economics of information markets and privacy (Bergemann and Bonatti 2019, 2024; Chen 2025). A central theme in this literature is the presence of externalities, as individual data-sharing decisions influence the privacy and welfare of others (Acemoglu et al. 2022; Rhodes and Zhou 2024; Miklós-Thal et al. 2024) by affecting firms' inference abilities. In this work, we abstract away from these spillover effects and instead focus on individual choices, in line with existing regulatory frameworks (Goldfarb and Tucker 2024).

We also contribute to the broader literature on choice architecture, particularly in the context of privacy decisions. Most empirical work has focused on documenting the prevalence of manipulative consent designs online (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), Posner et al. (2023), 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, thus offering greater

real-world relevance. In addition, our analysis covers a wide range of websites, allowing us to assess how design manipulations interact with domain familiarity and popularity.

Finally, our work is related to the existing literature on behavioral biases and their implications 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. Interestingly, we find that the effectiveness of choice architecture is uniform across websites, regardless of popularity. This suggests that policies targeting choice architecture might not reduce incumbents' data advantage (Campbell, Goldfarb and Tucker 2015; Aridor et al. 2024; Johnson et al. 2024).

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. Section 5 compares the welfare impacts of policies regulating consent choice architecture and presents additional data on the external validity of our findings. Section 6 discusses the implications of our findings.

2. Study Design

In this section, we introduce the consent interfaces in our experiment, then discuss the study phases and randomization, and describe the sample of participants.

Consent Interfaces. Our experiment aims to identify how consumers make cookie tracking choices across 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). Participants install the extension on their Chrome browser. The extension randomizes choice architecture by displaying different 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 choice architecture to encourage data sharing (Habib et al. 2022; Bielova, Santos and Gray 2024). Although a full typology of consent design patterns is beyond the scope of this paper, existing audits make clear that manipulative interfaces are widespread. For example, Utz et al. (2019) crawled major EU websites after the introduction of the General Data Protection Regulation (GDPR) and found that 57.4% employed at least one manipulative element in their consent banners. Similarly, Nouwens et al. (2020) reported that more than 80% of the top 10,000 UK websites used some manipulative elements.

These practices broadly fall into three categories. Obstruction tactics create friction around privacy-protecting choices—for example, by hiding “reject all cookies” behind extra clicks or defaulting to “accept all cookies.” These are among the most common patterns observed online (Habib et al. 2022). Visual manipulation alters option salience through layout, color, or ordering, such as placing “accept all” first or graying out “reject all.” Persuasion tactics frame data sharing in a more favorable light using non-neutral language. We focus on obstruction and visual manipulation, excluding persuasion tactics due to their high-dimensional nature and the resulting requirement for more treatment arms than our participant pool allows.

Figure 1 displays our six consent interfaces. We treat Design C (“Set-Acc-Rej”) as the *baseline interface*. The other interfaces implement 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”) places the accept option on top (*reordering options*); and Design F (“Acc-GrayRej-GraySet”) emphasizes the accept button with a brighter color while graying out the other options (*differential salience*).

For the deliberate obstruction and reordering options, we also include designs that nudge users towards privacy-friendly options (Design B, “Rej-Set”; and Design E, “Rej-Acc-Set”). Though these are not designs that companies would prefer for the sake of maximizing data sharing, they have been featured in policy discussions (e.g., the “privacy by design” requirement in GDPR and the California Privacy Rights Act, or CPRA) and actual platform implementations (e.g., Apple’s App Tracking Transparency framework uses both Design E and persuasive tactics).

Users can select any displayed option, or close the banner without making an explicit choice by clicking the X in the top right corner of the banner. If they click “X,” the website will implement its default data-sharing setting, which is normally “accept all” for U.S. websites. When users click on “Cookie Settings,” they are presented with six cookie categories to choose from, such as information storage & access, performance & analytics, and ad selection delivery & reporting (see Appendix Table D.1). Selecting all options is equivalent to accepting all cookies; selecting none is equivalent to rejecting all (non-essential) cookies. To minimize choice friction under “Cookie Settings,” we allow consumers either to accept all cookies with only one additional click or to reject all cookies with similar ease (as the default on this page is selecting none of the category-specific cookies).

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-GrayRej-GraySet”: accept-gray reject-gray settings.

Our banner can appear on any website, regardless of whether the site asks for consent. When there is an organic cookie choice interface, our banner replaces it.¹¹ To ensure that data-sharing choices are incentive-compatible, Cookie Manager enforces participants’ choices by detecting and interacting with websites’ native consent forms. The extension uses detection rules adapted from open-source packages (e.g., DuckDuckGo’s AutoConsent) and custom scripts to automatically select the corresponding option on each website’s original consent interface. When participants close the consent banner without making a choice, Cookie Manager

¹¹Our participants are U.S. residents, so most of them do not see consent banners as often as EU consumers due to the lack of federal privacy regulations that require consent.

does not modify any website settings. We explain this enforcement mechanism to participants during the onboarding process.¹²

Study Phases and Randomization. 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.

In the survey phase, participants visit 20 pre-selected websites spanning different categories and popularity levels (see Appendix D.1 for the full list), presented in randomized order. As participants browse each site, the extension displays a randomly assigned cookie consent banner and records users' selection. This phase serves two purposes: first, it ensures sufficient observations per web domain to estimate choice architecture effects using fixed-effect specifications; second, it complements the organic phase data by providing consent choices when website visits are exogenously given.

In the organic phase, the extension remains active for one week while participants organically browse the web. During this period, the extension shows consent banners on domains where participants have not already interacted with the experimental consent banners. To reduce disruption to browsing, banners appear at most every 10 minutes, rather than on all domains that participants visit during this phase.¹³ At the end of the week, participants complete an exit

¹²We tell users “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.”

¹³ In this second phase, we also randomize users into two frequency treatments: frequent banners, where banners appear at most every 10 minutes; and infrequent banners, where banners appear at most every 60 minutes. To implement this, a countdown starts after each banner interaction, and a new banner appears after the 10- or 60-minute threshold is crossed and a new domain

survey and uninstall the extension. Each participant receives \$7.50 upon study completion. Survey instruments are provided in Appendix E.

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 records the corresponding choice. In pilot testing, we found no evidence of carryover effects from prior exposures to specific banner designs, so we chose this approach to increase statistical power.

Sample Description. We recruited participants on Prolific,¹⁴ and restricted our participants to adults residing in the U.S. who primarily speak English and use Google Chrome as their main browser. We pre-registered recruiting 800 participants and expected 640 of them to be included in the main study sample.¹⁵ Our actual participants are close to the pre-registered numbers (see Appendix Table D.2 for the conversion funnel). A total of 1,227 Prolific users started the study, of whom 74.7% were eligible. Among these, 877 consented to the study, and 563 generated valid data points.

Our final sample of 563 users includes participants who completed the baseline survey and generated valid data during the organic browsing phase, regardless of whether they completed the exit survey.¹⁶ We also exclude a small number of participants who reported making choices based on their expectations about the study rather than genuine privacy preferences. These restrictions ensure a consistent sample across both survey and organic phases.

is visited, whichever occurs later. We do not find differences in consumer behavior across these two treatments (see Appendix B).

¹⁴<https://www.prolific.com/>.

¹⁵Our pre-registration is available at <https://www.socialscienceregistry.org/trials/12862>.

¹⁶Approximately 3% of users were not assigned to either of the intended frequency conditions (i.e., seeing consent banners at most every 10 versus 60 minutes) and saw banners for every new domain throughout the organic phase. Our main analysis also excludes these participants.

Table 1 presents descriptive statistics for the main sample. We have a balanced sample between men (55%) and women (45%), with an average age of 38 years. The median household income is \$50,000-\$74,999, with substantial variation, including 11% of households with an annual income above \$150,000. Participants visited an average of 53 unique domains in the week prior to the study.¹⁷ During the study week, participants visited an average of 54 unique domains (excluding the survey-assigned websites), suggesting minimal change in their browsing behavior in response to being part of the study. Our consent interfaces appear on 41.2% of the organically visited domains.

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
	Female	0.45	0.00	0.50
	Bachelor's or Above	0.52	1.00	0.50
	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 banners in two ways. First, we run a proportion test on the distribution of banner designs per website. The

¹⁷This information is calculated from participants' Chrome browsing history, which the browser extension collects upon participant consent.

proportion test across the survey websites has a p-value of 0.96, which fails to reject the null hypothesis of balanced proportions across treatments. Second, we perform covariate balance tests by regressing user- and domain-level covariates on treatment conditions (Appendix Table D.3). We find no statistically significant difference across designs.

3. Experimental Results

In this section, we present reduced-form evidence on the causal effects of choice architecture, 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 baseline interface, and designs that increase choice friction significantly shift consent behavior. While domain-level factors such as popularity have modest effects, user-level heterogeneity is substantial. Survey responses indicate that participants generally understand cookie functionality and the consequences of data sharing, suggesting that consent decisions are reasonably informed. However, beliefs vary considerably regarding default settings when they close the consent banner without making an explicit choice. We conclude by quantifying the time costs of cookie preference decisions.

3.1. The Effect of Choice Architecture on Data Sharing Choices

Figure 2 presents the choice distribution across treatment conditions, separately for the survey (Panel A) and organic phases (Panel B). We highlight three findings. First, participants share their data with websites more than 50% of the time absent deliberate nudging, with accept-all rates at 65% during the survey phase and 61%

during the organic phase.¹⁸ The exception is the banner condition where “accept all” is hidden from the main screen, where 19% of consumers or fewer during both phases choose to accept all cookies through “cookie settings.”


Second, granular choices, defined as accepting only specific cookie categories, are infrequent across all conditions. Granular choices occur at a rate from 3% in the baseline interface to 8% when “reject all” is deliberately hidden during the survey phase.

Third, participants respond similarly to choice architecture across both survey and organic phases, with one key exception: they close consent windows more frequently during organic browsing. This difference in behavior highlights the role of attention over repeated consent prompts and the importance of default settings when users organically browse the web.

Despite different propensities to close the consent banner, incorporating users’ beliefs about what happens when they close the window (discussed in Section 3.2) reveals similar underlying preferences across the two phases. After imputing passive choices using belief data, we estimate that 63% of users in the survey phase and 62% in the organic phase intend to accept cookies. Similarly, 30% in both phases intend to reject them. Together, these findings underscore two points. First, survey-based choices, though potentially more artificial, still reliably capture choice architecture effects. Second, default settings play a critical role when users shift their attention from privacy decisions to other browsing activities.

¹⁸Our acceptance rates are high but lower than prior evidence in lab studies that also examine privacy choices absent nudging designs, which document the probability of choosing “accept all” to be 77% (Vásquez 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/>, accessed November 2025. However, Apple uses wording that discourages sharing (Baviskar et al. 2024), and we suspect that some mobile data (such as location) can be more sensitive compared to web behavior data.

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 banner 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 banner design on consumer choice, we estimate the following type of regressions:¹⁹

$$y_{ij} = \beta_{acc-set;ij} + \beta_{acc-grrej-grset;ij} + \beta_{acc-rej-set;ij} + \beta_{rej-acc-set;ij} + \beta_{rej-set;ij} + \mu_i + \nu_{c(j)} + \epsilon_{ij}. \quad (1)$$

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 obtained using large language models.²⁰ Each β coefficient measures the effect

¹⁹Our treatments often combine different types of manipulations into a single interface. For example, “acc-set” simultaneously prioritizes “accept all” and hides “reject all.” Appendix A shows the separate effects of individual choice architecture (re-ordering, obstruction, and highlighting) on consumer choices.

²⁰We use ChatGPT 4o to classify the websites via 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’,

of a specific treatment condition relative to the baseline interface (set-acc-rej in Figure 1).²¹

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. We separate the analysis for survey and organic choices and focus on three outcomes: accepting all cookies, rejecting all cookies, and closing the window without choosing.²²

Treatments involving deliberate obstruction are the most effective at steering consumer choices. Hiding “reject all” (the *Acc-Set* treatment) decreases rejection

‘Other’. Each domain is separated by ‘+ - - - +’. Please always return a 10-element sequence of classifications separated by ‘- - -’.

²¹Relative to the pre-registered specification, we have changed the comparison benchmark to be a design without nudges.

²²Given its small share, the analysis that focuses on granular choice as the outcome is left to Appendix Table D.6.

rates by 17.8 percentage points (a 74% decline). Acceptance 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.²³

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 produce modest effects, concentrated in the organic phase, perhaps because users pay less attention to cookie decisions while naturally browsing the web. In the organic phase, placing “accept all” at the top (*Acc-Rej-Set*) increases acceptance rates by 3.8 percentage points (a 6% increase), while placing “reject all” at the top (*Rej-Acc-Set*) slightly increases rejections (up 2 percentage points, a 14% increase); neither design has a meaningful effect in the survey phase. The treatment that highlights “accept all” and ranks it on top (*Acc-GrRej-GrSet*) modestly increases acceptance in the organic phase by 3.1 percentage points, but this effect is indistinguishable from the effect 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.

²³Our consent banners that incorporate deliberate obstruction also position specific options more prominently, making it difficult for the regressions in Table 2 to isolate the effects of each individual choice architecture element. Appendix Table A.1 addresses this by providing separate estimates, and shows that the effects in the main analysis stem primarily from deliberate obstruction.

Appendix Table D.6 shows that consumers rarely make granular cookie choices unless prompted by design interventions. In the baseline 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. Among those making selective choices, 83% accept cookies for *preferences and functionality*, while only 7% consent to *ad selection, delivery, and reporting* (see Appendix Table D.5). This pattern suggests that targeted advertising represents consumers’ least preferred data use, at least among users who selectively consent to cookie tracking.

These results point to three conclusions. First, users often accept cookie tracking when deliberate nudging is absent. Second, obstruction-based choice architecture is more effective than visual manipulation.²⁴ Third, dismissal of consent windows without an active choice occurs frequently, even more so in the wild than in survey-based settings.

3.2. Heterogeneity across Websites and Consumers


In addition to the effects of choice architecture, we examine how cookie choices vary across websites and individuals to assess the potential competitive implications of consent design. Prior work shows that privacy decisions are highly context-dependent (Nissenbaum 2004; Lin 2022), and our results extend this literature by documenting substantial user-level heterogeneity even when holding contexts fixed. Overall, 68.9% of participants change their cookie-sharing decisions across sites at least once. This variation reflects a combination of choice

²⁴These 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 minimal effect on choices.

architecture effects and systematic variation in preferences across websites and users.

Website-Level Heterogeneity. We examine whether consent behavior varies across websites by content category, prior familiarity, and domain popularity, and whether choice architecture amplifies these differences. 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 pattern suggests that participants are more hesitant to share data with websites perceived as more sensitive compared to those in the e-commerce and entertainment categories.

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.

Next, we examine how a user’s familiarity with a website or the website’s popularity shapes consent decisions. We draw on two data sources to measure participants’ familiarity with each site. First, we use participants’ stated reports about whether they had heard of and regularly visited each of the 20 assigned survey 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 site popularity, we use log domain ranks from Tranco,²⁵ which aggregates rankings from multiple public sources.

To examine how cookie choices vary with user familiarity, we add explanatory variables to Equation 1 and report their coefficient estimates in Table 3. Panel *a* focuses on the survey phase, using self-reported measures of website experience. Participants are 3 percentage points more likely to accept cookies on websites they have heard of, and 7 percentage points more likely on websites they normally visit.

Panel *b* combines data from both the survey and organic phases, using observed behavioral 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 is associated with greater willingness to accept cookies, suggesting greater willingness to share data with well-known sites.

Together, these results confirm that user familiarity, whether self-reported or inferred from usage data, is associated with higher cookie acceptance rates. 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.

²⁵<https://tranco-list.eu/>, accessed November 2025.

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: User characteristics</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)
Prior Domains Visits (Log)	-0.012 (0.013)	-0.013 (0.011)	0.019* (0.009)
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 intake survey, so we only include consent choices during the survey phase). 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 indicating whether the study participant visited the website in the two weeks preceding the study (obtained from their Chrome browsing history) and by the website's popularity rank (from Tranco logs). Panel *c* includes user characteristics (age, gender, income, education, and the number of prior domain visits from their browsing history). Panel *d* focuses on self-reported and observed privacy attitudes. We include indicators of participants' beliefs about the meaningfulness of privacy decisions, whether they highly value privacy, and whether they left the study before completing the final survey. Regressions in Panels *c* and *d* require us to exclude individual fixed effects. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

We see limited heterogeneity in the effectiveness of choice architecture across site popularity and familiarity (see Appendix Figure D.3 and Appendix Tables D.7 and D.8). If anything, nudges appear to dampen participants' tendency to share more data with popular and familiar websites, suggesting they are unlikely to reinforce data-driven competitive advantages.

Individual-Level Heterogeneity and User Beliefs. Next, we examine variations in data-sharing preferences across users. While earlier sections documented heterogeneity in behavior across websites, Panel c of Table 3 shifts the focus to user characteristics as predictors of privacy choices. Overall, most demographic factors have limited predictive power, but some meaningful patterns emerge. Women are modestly less likely to accept all cookies, and older participants are slightly less likely to close the banner without making a choice. Users who browse the internet more frequently are more likely to close consent windows.

While these patterns suggest some demographic tendencies in privacy behavior, they raise a further question: to what extent are users' choices affected by an understanding of what cookies do and how they affect data sharing? 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.

Survey responses suggest that participants made reasonably informed decisions. Among those who accepted all cookies, the most common motivations were website trust (38.6% reported "It was a website I totally trusted"), convenience (21.9%: "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”), and functionality (17.6%: “To make the websites function well”). Cookie rejection reflected similar reasoning: participants cited concerns about unfamiliar or untrustworthy websites (26.2%: “Did not trust site”), tracking or advertising concerns (7.4%: “I didn’t want to be tracked by advertisers on that website”), and general privacy preferences (13.2%: “in some website i felt not to accept some cookies because i may not want the website to have my data”).

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 70% identified personalized ads and 15% identified better site functionality as benefits, while over 75% cited privacy loss (43%), security risks (18%), or increased advertising (15%) as key concerns. Only less than 4% reported being unsure about what cookies do.

However, participants held varying beliefs about what would happen if they dismissed consent banners. 52.2% believed that closing the banner using the “X” button meant rejecting cookies, while 21.7% believed it meant accepting cookies. The remaining respondents expressed uncertainty. Given that U.S. websites typically default to collecting cookies, most users misunderstood actual website practices. We formally account for these belief differences in our structural model and welfare calculation in the next section.

Stated privacy preferences are weakly correlated with revealed preferences. Panel d of Table 3 shows that participants who say they highly value privacy are 7 percentage points more likely to reject all cookies. However, the majority coefficients are indistinguishable from zero.²⁶

²⁶Additional analyzes of survey-based privacy preferences and decisions are presented in Appendix Tables D.10 and D.9).

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 D.11 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 consent banners degrade user experience and impose time costs. Here, we quantify these time costs and how they vary across interface designs.

We measure time spent on consent banners as the elapsed time between banner appearance and final action, including intermediate clicks and any back-and-forth interactions. We censor the time spent at 60 seconds, well above the 99th percentile, to exclude likely task-switching behavior rather than genuine interaction time.

Like in Section 3.1, we estimate Equation 1 but instead use the time spent interacting with the banner as the outcome. Table 4 displays the results. In the baseline 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 substantially increases interaction time is the *rej-set* design, which hides “accept all,” the most commonly chosen option.²⁷

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

Additional clicks substantially increase time costs. Table 5 shows that accessing the settings menu, required when preferred options are hidden, accounts for over 50% of average banner interaction time. This finding explains why hiding options behind the main screen creates substantial choice friction that deters users from their preferred choices.

²⁷This additional time cost may stem from user confusion or unfamiliarity, since interfaces where “accept all” is hidden are not used by websites.

TABLE 5. Extra Time Spent When Clicking on Settings

	Time Spent (Seconds)
	(1)
User Clicks Settings	4.069*** (0.442)
Close Window	1.825*** (0.409)
Accept Selected	3.282*** (0.673)
R ²	0.102
Observations	22,842
Condition fixed effects	✓
Participant fixed effects	✓
Domain Cat. fixed effects	✓

Notes: The table presents regression estimates of time spent to make a decision as a function of whether the user’s action involved clicking on 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 the survey and organic data. Standard errors are clustered at the user level. * $p < 0.05$, ** $p < 0.01$; *** $p < 0.001$.

To translate the time costs into monetary terms, we perform a back-of-the-envelope calculation using the U.S. average hourly wage of \$36 as a baseline.²⁸ With users spending 7.34 seconds per domain and visiting 54 unique domains per week on average, this implies a weekly time cost of approximately \$4 per user. This calculation previews the analysis in the next section, which incorporates both choice architecture effects and time costs into an evaluation of consumer welfare.

²⁸The \$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>, accessed November 2025).

4. Structural Model and Estimation

To quantify the welfare implications of different consent choice architectures, we develop a structural model of consumer data sharing decisions that incorporates the influence of consent design. The model translates observed choices into welfare measures and quantifies the costs imposed by design frictions. 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.

4.1. Model Setup

Given a cookie consent banner, a consumer can choose among four actions: {accept all, reject all, customize settings, close window}. We model consumers' decision utility as a function of the selected option and the frictions introduced by the design of the consent interface. Here, we use the notion of "decision utility" (Kahneman and Thaler 1991) to recognize that consumers may not always choose the option that maximizes their welfare, since consumers may have mistaken beliefs and their choices can be affected by salience manipulations. For notational simplicity, we omit consumer-specific subscripts below.

Let θ_k denote the value that the consumer assigns to data-sharing option k . Specifically, θ_{acc} corresponds to the value of accepting all cookies, θ_{rej} to rejecting all cookies (normalized to zero for identification), and θ_{set} to customizing settings.²⁹

²⁹Our design does not separately identify the benefit of granular cookie selection from the hassle cost of doing so. As a result, θ_{set} is the net utility of the customization benefit and the hassle cost of making granular selections.

The perceived value of closing the consent banner depends on a consumer’s belief about the website’s data collection default. Our endline survey elicits each participant’s subjective probability ρ_{belief} that closing the window results in accepting all cookies. Their perceived utility is thus $\rho_{\text{belief}}\theta_{\text{acc}}$ (recall $\theta_{\text{rej}} = 0$). For respondents who believe that closing the window means accepting cookies (52%), we set $\rho_{\text{acc}} = 1$; for those who believe it means rejecting cookies (22%), we set $\rho_{\text{rej}} = 0$. For participants who report uncertain beliefs (12%) or did not answer the belief question (14%), we estimate their beliefs as ρ_{unc} and ρ_{miss} , respectively. Under the assumption of risk neutrality, ρ_{unc} and ρ_{miss} can be interpreted as the probability that these participants assign to the event that closing the window implies full cookie acceptance.³⁰

We incorporate a cost of closing the consent banner because people are more likely to make an explicit choice than to close the window. The cost consists of two terms: a constant term κ_1 and a time-dependent term κ_2/t_i , where t_i represents the number of days since user i ’s enrollment in the study (with the enrollment day counted as day 1). The time-dependent term is motivated by the fact that users increase the frequency with which they close consent banners over time (Appendix Table B.1).

Choice architecture affects decisions through two channels. First, deliberate obstruction creates friction through click costs: consumers incur an additive cost C_{set} to access the settings menu when options are hidden. Second, visual manipulations alter salience and thus the decision utility of affected options. We model their effects as multiplicative factors applied to the baseline utility. For positive baseline utilities $\theta_k > 0$, placing the “accept all” option at the top

³⁰We do not apply any discount for participants who 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.

increases its utility to $(1 + \delta_{\text{top}})\theta_k$, while graying out an option reduces its utility to $(1 + \delta_{\text{gray}})\theta_k$, where $\delta_{\text{gray}} < 0$. To accommodate potentially negative baseline utilities while ensuring proper estimation of the δ parameters, we specify the salience-adjusted utilities as $\theta_k + \delta_{\text{top}}\theta_k^2$ for top placement and $\theta_k + \delta_{\text{gray}}\theta_k^2$ for graying out. This formulation ensures that salience effects scale proportionally with the magnitude of baseline utility, regardless of sign.

For ease of comparison, Table 6 summarizes the complete utility specification across all treatment conditions.

TABLE 6. Decision Utility of Each Option Across Treatment Conditions

Treatment Design	Option			
	Accept All	Reject All	Customize Settings	Close Window
Acc-Set	$\theta_{\text{acc}} + \delta_{\text{top}}\theta_{\text{acc}}^2$	$-C_{\text{set}}$	$\theta_{\text{set}} - C_{\text{set}}$	$\rho_{\text{belief}}\theta_{\text{acc}} + \kappa_1 + \frac{\kappa_2}{t_i}$
Rej-Set	$\theta_{\text{acc}} - C_{\text{set}}$	0	$\theta_{\text{set}} - C_{\text{set}}$	$\rho_{\text{belief}}\theta_{\text{acc}} + \kappa_1 + \frac{\kappa_2}{t_i}$
Set-Acc-Rej	θ_{acc}	0	$\theta_{\text{set}} + \delta_{\text{top}}\theta_{\text{set}}^2 - C_{\text{set}}$	$\rho_{\text{belief}}\theta_{\text{acc}} + \kappa_1 + \frac{\kappa_2}{t_i}$
Acc-Rej-Set	$\theta_{\text{acc}} + \delta_{\text{top}}\theta_{\text{acc}}^2$	0	$\theta_{\text{set}} - C_{\text{set}}$	$\rho_{\text{belief}}\theta_{\text{acc}} + \kappa_1 + \frac{\kappa_2}{t_i}$
Rej-Acc-Set	θ_{acc}	0	$\theta_{\text{set}} - C_{\text{set}}$	$\rho_{\text{belief}}\theta_{\text{acc}} + \kappa_1 + \frac{\kappa_2}{t_i}$
Acc-GrRej-GrSet	$\theta_{\text{acc}} + \delta_{\text{top}}\theta_{\text{acc}}^2$	0	$\theta_{\text{set}} + \delta_{\text{gray}}\theta_{\text{set}}^2 - C_{\text{set}}$	$\rho_{\text{belief}}\theta_{\text{acc}} + \kappa_1 + \frac{\kappa_2}{t_i}$

Notes: The table reports the decision utility associated with each option (Accept All, Reject All, Customize Settings, and Close Window) across the six treatment conditions. Although not explicitly indicated with subscripts, the parameters $(\theta_{\text{acc}}, \theta_{\text{set}}, C_{\text{set}}, \delta_{\text{top}}, \delta_{\text{gray}})$ are functions of the consumer type ι . In contrast, the costs of closing the window (κ_1, κ_2) and the belief parameters $\rho_{\text{belief}} = (\rho_{\text{unc}}, \rho_{\text{mis}})$ are assumed to be constant across consumer types.

A consumer i picks the option k that maximizes their utility given the banner design on website j :

$$\text{Max}_{k \in \{\text{acc}, \text{rej}, \text{set}, \text{close}\}} U_{ijk|design_{ij}},$$

where $U_{ijk|design_{ij}} = u_{ik|design_{ij}} + \epsilon_{ijk}$. The first component is the user i 's decision utility from option k (which varies by treatment banner design as presented in Table 6), whereas the random component ϵ_{ijk} is i.i.d. distributed according to a Type-1 extreme value distribution.

To capture the preference heterogeneity documented in Section 3.2, we estimate a latent class model (Kamakura and Russell 1989), where we allow for three latent consumer types ι with shares below:

$$s_{i\iota} = \frac{\exp(\lambda_{i\iota})}{\sum_{\iota' \in \{1,2,3\}} \exp(\lambda_{i\iota'})}. \quad (2)$$

We let the shares depend on the user's education and browsing activity in the week preceding the study: $\lambda_{i\iota} = \lambda_{\iota}^1 + \lambda_{\iota}^2 \mathbf{1}\{\text{education}_i > \text{bachelor degree}\} + \lambda_{\iota}^3 \log(1 + \text{prior domain visits}_i)$, with normalization $\lambda_3 = 0$.

Each consumer type ι has distinct preferences and choice architecture sensitivity parameters $(\theta_{acc}, \theta_{set}, C_{set}, \delta_{top}, \delta_{gray})$, while belief parameters (ρ_{unc}, ρ_{mis}) and closing window costs (κ_1, κ_2) are common across types. Conditional on a user being type ι , the probability of choosing option k is

$$P_{\iota jk} = \frac{\exp(u_{\iota k | design_{ij}})}{\sum_{k \in \{acc, rej, set, close\}} \exp(u_{\iota k | design_{ij}})}. \quad (3)$$

The unconditional probability of selecting option k by consumer i is thus its probability of being a certain type and the type-specific choice probability:

$$P_{ijk} = \sum_{\iota \in \{1,2,3\}} s_{i\iota} P_{\iota jk}. \quad (4)$$

4.2. Identification and Estimation

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 options are identified from choices under the baseline treatment without deliberate nudging. Friction costs are identified by comparing choice

probabilities when options are hidden versus visible on the main screen. Saliency effects are identified by comparing choices between interfaces that reorder or gray out options while maintaining the same choice set. Beliefs about what happens upon banner closure are separately elicited through the endline survey, which helps us infer the potential mistakes and costs associated with closing the banner. The panel structure across dates helps pin down the time-varying component of the utility from closing the window.

We use Equation 4 to form our likelihood function, and estimate the model parameters by maximum likelihood using data from both the survey and organic phases. We cluster standard errors at the participant level.

4.3. Results

Panel A of Table 7 presents the estimates for consumer utility parameters. Type-1 users have the highest utility from accepting all cookies ($\theta_{acc} = 2.66$), so we label them “acceptors.” Type-2 users strongly dislike data sharing ($\theta_{acc} = -4.96$), so we label them “rejectors.” Type-3 users, labeled “discerners,” have a moderate utility for accepting all cookies ($\theta_{acc} = 1.21$) and a greater willingness to customize settings compared to the other two groups ($\theta_{set} = 0.62$).

The friction parameters reveal that interface designs impose substantial behavioral costs. The cost of clicking on the “settings” button (C_{set}) ranges from 1.7 to 2.9 across consumer types, matching the magnitude of the utility from accepting cookies for acceptors and discerners. In comparison, visual saliency effects are more modest. While the signs of the parameters indicate that ranking an option on top increases its choice probability and graying out an option reduces it, the

TABLE 7. Parameter Estimates Across Consumer Subsets

Parameter	Explanation	Estimates		
Panel A: Utility Parameters				
<u>Type-Specific Parameters</u>		<u>Type 1</u>	<u>Type 2</u>	<u>Type 3</u>
θ_{acc}	Utility of accepting all	2.660*** (0.157)	-4.960*** (0.736)	1.210*** (0.085)
θ_{set}	Utility of granular choice	0.407 (0.355)	-0.280* (0.118)	0.615*** (0.158)
C_{set}	Cost of clicking settings	2.900*** (0.200)	2.650*** (0.166)	1.660*** (0.116)
δ_{top}	Effect of ranking on top	0.422*** (0.082)	0.029 (0.042)	0.076 (0.077)
δ_{gray}	Effect of being grayed out	-1.610 (10.200)	-4.410 (5.030)	-0.041 (0.412)
<u>Pooled Parameters</u>				
κ_1	Utility of closing window		0.752*** (0.045)	
κ_2	Inclination of closing over time		-1.670*** (0.070)	
ρ_{unc}	Pr(accept close) for uncertain		0.352*** (0.025)	
ρ_{miss}	Pr(accept close) for missing		0.469*** (0.138)	
Panel B: Latent Class Probability Parameters				
λ^1	Intercept	0.573*** (0.165)	-0.510** (0.194)	
λ^2	Bachelor degree or higher	-1.720*** (0.426)	-0.874** (0.305)	
λ^3	Num. prior domain visits	-1.470*** (0.238)	-1.520*** (0.214)	
Share of Users		0.454	0.176	0.369

Notes: The table presents parameter estimates from the utility model. Panel A reports estimates for consumer utility parameters, while Panel B reports estimates for the latent model parameters. Given the parameter estimates from Panel A, we refer to Type 1 users as “acceptors,” Type 2 users as “rejectors,” and Type 3 users as “discerners.” Standard errors are clustered at the participant level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

effects are not statistically significant for most types. Overall, these results align with our reduced-form findings that obstruction dominates visual manipulation.

The cost of closing the banner without making a choice ($\kappa_1 + \frac{\kappa_2}{t_i}$) is initially high relative to rejecting cookies, estimated as -0.92 on the first day of the study. However, the closing option becomes increasingly attractive over time as the time-dependent component $\frac{\kappa_2}{t_i}$ increases. By day 3, closing the banner becomes more attractive than explicitly rejecting cookies. Finally, the belief parameters (ρ) suggest that respondents who are uncertain about the website closing default assign a 35% probability that closing the window implies accepting cookies, while for participants who did not complete the endline, this probability is estimated to be 47%.

Panel *B* of Table 7 shows how consumer types vary across the population. Type-1 users (acceptors) form the largest group at 45% of users, followed by type-3 (discerners) at 37% and type-2 (rejectors) at 18%. Both education and prior browsing intensity significantly predict consumer type. College-educated users are the least likely to be acceptors and most likely to be discerners (as indicated by the estimates of λ^2). Similarly, users with more extensive browsing histories (higher number of domains visited) are significantly less likely to be classified as acceptors or rejectors than discerners (from the λ^3 estimates).

5. Consumer Surplus under Counterfactual Policies

The structural model estimates enable us to assess the welfare impacts of various consent policy designs. Existing policies or policy proposals vary in three dimensions: the type of banner designs they allow, website defaults upon inaction, and the frequency with which users are asked to consent. The fact that choice frictions

significantly deter users from selecting their preferred option justifies policies that reduce those frictions. One common approach is to mandate interfaces free of obstruction while maintaining site-by-site consent. With site-specific consent, the website's default can be critical, as many users close the banner without making a choice; we therefore compare policies with “pro-privacy” and “pro data-sharing” defaults upon banner closure. A different policy proposal, browser-level consent, reduces the frequency of consent decisions to the browser level. This policy, also known as *global privacy control* (GPC for short),³¹ reduces interaction costs but limits customization, making its welfare impact ex-ante ambiguous. GPC is gaining momentum, supported not just by US state privacy laws but also more recently by a proposal for EU privacy law reform.³² Several browsers, including Firefox, Brave, DuckDuckGo, and Chrome extensions such as Disconnect, already enable users to set preferences under the GPC framework.

We evaluate consumer welfare under the following policies:

1. **U.S. status quo:** An “acc-set” design (i.e., a design that uses deliberate obstruction and reordering), with websites defaulting to collect all cookies when the user does not make an explicit choice;
2. **EU norm:** An interface without deliberate nudging (“set-acc-rej”), with websites defaulting to no cookie tracking upon inaction;
3. **Optimal Banner Interface:** An interface without nudging (“set-acc-rej”), with the default set to collect all cookies upon inaction;

³¹<https://globalprivacycontrol.org/>, accessed November 2025.

³²See https://www.nelsonmullins.com/insights/alerts/privacy_and_data_security_alert/all/get-to-know-the-global-privacy-control-gpc-state-attorneys-general-remind-consumers-of-their-right-to-opt-out-of-the-sale-or-sharing-of-personal-information and <https://brusselssignal.eu/2025/11/eu-proposes-to-loosen-ai-and-data-privacy-rules-cut-cookie-banners> (accessed November 2025).

4. **Optimal Banner Interface with Correct Beliefs:** An interface without nudging (“set-acc-rej”), with the default set to collect all cookies upon inaction, and with consumers having correct beliefs about the default settings;³³
5. **Global Accept:** A browser-level configuration that automatically accepts all cookies across websites;
6. **Global Reject:** A browser-level configuration that automatically rejects all cookies across websites;
7. **Global Privacy Control:** A browser-level setting that lets users choose one of three global options for all websites: accept all cookies, reject all cookies, or apply a uniform customized preference.

Our welfare calculation needs to address three key features specific to our setting. First, consumers’ mistaken beliefs about default settings and designs that alter the visual salience of options can lead to suboptimal decisions and welfare losses. Second, the time cost of interacting with consent banners is an essential component of consumer welfare, which becomes relevant for global privacy control policies. Lastly, although our experiment does not directly elicit monetary valuations for privacy or time, we must translate the estimated utility into dollar terms. We describe our approach to each of these components below.

Accounting for belief mistakes and visual salience distortion. Behavioral economics typically distinguishes between welfare utility, the utility that reflects actual welfare outcomes (Kahneman and Thaler (1991)), and decision utility (or perceived utility), which governs observed choices and can be shaped by choice architecture or incorrect beliefs. In this terminology, we treat consumer choices in the baseline treatment, where no deliberate nudges are present and beliefs are correct, as reflecting welfare-relevant utility. We also differentiate how various

³³Due to the absence of belief mistakes and salience distortion, consumers’ utility in this state would be the “welfare relevant” utility per Kahneman and Thaler (1991).

elements of choice architecture affect utility. Specifically, we regard time costs as real costs that reduce welfare utility, whereas visual-salience manipulations distort decision utility without altering the underlying welfare utility.

For site-by-site consent policies (policies 1-4), the welfare calculation must account for the discrepancy between the welfare and decision utilities. To this end, we allow choice probabilities to be guided by consumer beliefs and salience parameters, but compute the welfare utility ignoring salience parameters and using the actual website default, following Train (2015). The surplus received by consumer i is as follows:

$$CS_i(\text{banner-level policy}) = \sum_{\iota} s_{i\iota} \left[\log \left(\sum_k e^{u_{ik\iota} | \text{policy}} \right) + \left(\sum_k P_{ik\iota} | \text{policy} \cdot d_{ik\iota} | \text{policy} \right) \right].$$

In this equation, $\iota \in \{1, 2, 3\}$ denotes the latent consumer type and $s_{i\iota}$ indicates the probability that consumer i is in segment ι ; $k \in \{acc, rej, set, close\}$ denotes the choices. The first term is the standard logit consumer surplus based on decision utilities $u_{ik\iota} | \text{policy}$, which vary across individuals due to latent types and beliefs about default settings. $P_{ik\iota} | \text{policy}$ is the choice probability for option k , and $d_{ik\iota} | \text{policy}$ is the difference between the welfare and the decision utility. The difference arises from two sources: mistaken beliefs about the default upon closing the window (ρ_{belief}) and salience effects when consent banners visually vary how options are displayed (δ_{top} and δ_{gray}).

Accounting for the value of time in browser-level control. Browser-level policies (policies 5-7) eliminate repeated banner interactions but restrict site-by-site customization. Consumer surplus becomes:

$$CS_i(\text{browser-level policy}) = \sum_{\iota} s_{i\iota} u_{ik^* \iota} | \text{policy} + V_{\text{banner time}},$$

where $k^* = \text{acc}$ for global accept, $k^* = \text{rej}$ for global reject, and $k^* = \arg \max_k u_{ikl}$ for individualized global privacy control. The consumer surplus reflects only the mean utility of the selected option, ignoring idiosyncratic utility shocks since users make choices before interacting with specific websites. We also assume that browser-level interfaces are free of nudging patterns, and consumers do not close the browser-level consent without making a choice.³⁴

To calculate the utility of time saving from bypassing banner interaction, we assume that the (dis)utility is proportional to time spent, and use the C_{set} estimate to back out this utility term: $V_{\text{banner time}} = C_{set} \times (\text{time spent per banner}/\text{time to click "settings"})$. The consumer surplus from global privacy choices thus reflects three components: utility loss from the inability to customize cookie preferences for each website ($u_{ik^*l|policy} < \log(\sum_k e^{u_{ikl|policy}})$); gains from avoiding repeated interactions with banners ($V_{\text{banner time}} > 0$); and the absence of choice distortions induced by choice architecture ($\sum_k P_{ikl|policy} \cdot d_{ikl|policy} < 0$).

In assuming that disutility is proportional to time spent, we abstract from additional annoyance or cognitive costs that may be generated by deliberate obstruction or by the presence of consent banners themselves (see Bernheim, Kim and Taubinsky (2024) for a detailed discussion). We also abstract from the possibility that time spent browsing is valued differently from working time. These two omissions likely work in opposite directions: the former would imply our estimate underestimates the friction costs, while the latter implies that we could potentially overestimate the opportunity cost of time. Although we do not attempt to adjust for either effect directly, we later conduct a break-even analysis to quantify the time valuation that would make consumers indifferent across policy regimes.

³⁴One may want to consider the possibility that choice architecture is present even in global privacy control. To this end, global accept/reject policies can be interpreted as the extreme case when global privacy control is given but hidden (thus hard for consumers to choose), and websites default to collecting/rejecting cookies, respectively.

A final consideration is that the value of closing the banner evolves over time. Rather than extrapolating the utility of closing to a steady state, we use the value estimated during the week of the study. Note that participants' inclination to close window increases over time but gets stable by the end of the third day. If the value of closing continues to increase beyond what we observe at the end of the study, then resolving mistaken beliefs about closing would generate even larger welfare gains than those we report.

Expressing welfare estimates in dollar terms. Note that our data does not have consumers' willingness to pay for privacy or to avoid additional clicks. To address this challenge, we use participants' self-reported income to estimate their dollar value of time spent clicking on settings.³⁵ Specifically, we define the monetary value of the clicking cost (C_{set}) as the average time spent clicking on settings multiplied by a calibrated opportunity cost of time. Table 5 shows that users take 4.07 additional seconds per domain when clicking on settings. To get the opportunity cost of time, we use participants' self-reported annual household income to assign hourly wage values based on bracket midpoints. We then multiply this value by the time it takes to click on settings per domain and the number of unique domains visited per week to convert C_{set} into a weekly dollar value. For example, a participant in the \$75–100k income bracket (\$42 implied hourly wage) would incur a weekly time cost C_{set} of \$2.6, calculated as $(\$42/\text{hour}) \times (4.07/3,600 \text{ hours/domain}) \times 54 \text{ domains}$. We then use the ratio between the dollar-value opportunity costs of time and C_{set} to convert consumer welfare values to dollars.

³⁵During 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.

Results. Table 8 reports our consumer surplus estimates in dollar terms. The first column displays results averaged across types, while the remaining columns highlight the heterogeneity among the three consumer types.³⁶ Although the per-user-week effects appear modest in absolute terms, they represent meaningful welfare changes that compound substantially over time.

TABLE 8. Consumer Surplus Under Counterfactual Policies (\$, Weekly)

Counterfactual	Average	Type 1 Acceptors	Type 2 Rejectors	Type 3 Discerners
U.S. Status Quo	2.03	2.69	-2.63	3.43
EU Norm	2.58	2.66	0.91	3.28
Optimal Banner Interface	2.65	2.92	-0.12	3.63
Optimal Banner Interface (with correct beliefs)	2.89	3.1	0.6	3.73
Global Accept	5.81	6.94	0.37	7.01
Global Reject	4.98	4.76	5.03	5.24
Global Privacy Control	6.63	6.94	5.03	7.01

Notes: The table reports consumer surplus per week per user in dollar units under various counterfactual policies. The U.S. status quo refers to an *acc-set* interface, combined with an “accept all” default when consumers close window; the EU norm refers to an interface without nudging designs with a “reject all” default when consumers close window; the optimal banner interface refers to an interface without nudging with an “accept all” default when consumers close window; the optimal banner interface with correct beliefs ensures that people closing the window have correct beliefs about default settings. “Global Accept” and “Global Reject” force each individual to either always accept or always reject all cookies. The last row, “Global Privacy Control,” allows each individual to make their preferred global choice. The average column refers to the estimate across all subjects, and the other columns correspond to estimates by user type.

We start by comparing consent-based policies. Since consumer welfare is estimated up to a constant (Train 2009), we only discuss level differences. The difference between the U.S. status quo and the optimal banner interface reflects the impact of choice architecture, while the difference between the EU norm and the optimal banner interface reflects the impact of default settings. The optimal banner interface, which is a no-nudging design with an accept-all default, yields an average consumer surplus of \$2.65 per person-week. Relative to the optimal banner interface, the U.S. status quo delivers \$0.6 lower consumer surplus per user-week, with much larger losses for type-2 users (“rejectors”). Correcting beliefs

³⁶Appendix Table D.12 reports welfare in utility units.

about default settings boosts surplus by \$0.2 per user-week, beyond the gains from removing deliberate obstruction alone. Rejectors benefit the most, since they no longer believe that closing the banner rejects cookies.

The EU Norm, defined as a no-nudging interface paired with a reject-all default, results in a small difference (\$0.1 lower consumer surplus per week) compared to the optimal banner interface. Defaulting users to rejecting cookies has two opposing effects on consumer welfare. On one hand, it could improve average welfare by conforming to the belief of the majority. On the other hand, it can reduce welfare by deviating from the option actually preferred by the majority of our participants. Our analysis suggests that the net effect is slightly negative across consumers, though it benefits the rejectors over alternative policies, who gain from having the default aligned with their preferences. The EU norm also outperforms the U.S. status quo, increasing average surplus by \$0.6 per user-week, due to its use of a baseline design that reduces choice friction.

Remarkably, all global privacy control policies (last three rows in Table 8) substantially outperform even the optimal banner-based approach with correct beliefs. Acceptors and discerners benefit the most from a global accept policy, with surplus gains increasing by \$3.3-3.8 per week compared to the optimal consent-based policy. Rejectors benefit the most from a global reject policy, enjoying a surplus of \$5 per week compared to \$0.9 under the EU norm, their preferred banner policy. Global choice benefits even the discerners, who would customize data sharing decisions when site-specific banners are given, because the time-saving benefits from avoiding repeated interactions outweigh the costs of uniform choices. The large magnitude of the time-saving gains is consistent with our reduced-form evidence, where removing an option from the main consent interface decreases selection of that option by at least 70%. On average, the surplus gain from adopt-

ing the global privacy control is \$3.7 per week compared to the optimal banner interface.

So far, our analysis points to the conclusion that global privacy control consistently dominates site-by-site consent policies. We next assess how sensitive this conclusion is to our assumption about the value of time. To do so, we calculate the banner interaction time that would make the optimal site-by-site consent policy (given existing consumer beliefs) welfare-equivalent to the global privacy control. This break-even point occurs at just 0.64 seconds per banner, approximately 8.7% of users' current average interaction time. Put differently, on average consumers would need to value their time at only \$3.10 per hour for the two policies to yield equal welfare.

Second, we validate these policy evaluation results using stated preferences for global privacy control tools in the endline survey. 62.4% of participants report preferring such a tool over site-by-site choice, while only 13.4% prefer the site-by-site approach (Appendix Figure D.2). The average stated valuation for a global privacy control tool is \$4, which falls within our estimated surplus gain from moving from the U.S. status quo to global privacy control: accepters gain \$4.2, rejectors gain \$7.7, and discerners gain \$3.6. While stated and revealed preferences need not coincide perfectly, the similarity in magnitudes provides reassurance that our surplus estimates are sensible.

Overall, these results indicate substantial welfare gains from allowing users to specify global privacy preferences that match their individual data-sharing attitudes. Current regulatory approaches that focus on banner design improvements, while beneficial, may miss larger opportunities for welfare enhancement through architectural changes to how consent is collected and managed.

5.1. Additional Evidence from a Follow-up Survey

As discussed above, translating consumers' privacy choices into welfare measures presents several challenges. Participants in our experiment may differ systematically from the broader population; in addition, U.S. consumers, who constitute the primary sample for our main study, may have different privacy preferences than European consumers, where most privacy regulations originate. To partially address these concerns, this subsection presents evidence from a follow-up survey designed to validate and benchmark our welfare measures. Broadly, the privacy preferences observed in our sample mirror those in other populations, reinforcing the external validity of our findings.

In the fall of 2025, we conducted an additional survey across three distinct Prolific populations: (i) participants from our main study, (ii) U.S. Prolific users who did not participate in the main study, and (iii) EU participants. The survey consisted of three main blocks of questions. The first block collected demographic information,³⁷ which was always presented first. The second block focused on privacy attitudes: what privacy means to respondents and their perceptions of what advertisers know about them. The third block involved three pairwise comparisons of consent interfaces: (1) acc-set versus acc-rej-set, (2) acc-rej-set versus acc-rej-set with all settings options immediately visible, and (3) respondents' preferred banner versus a global privacy control (a similar comparison is contained in the outtake survey of the main study, as discussed previously). More details are provided in Appendix C. We recruited 320 participants from our main study, 471 additional participants residing in the U.S., and another 471 participants from the

³⁷ Requesting demographic information allows us to account for how demographic compositions affect the differences in the responses across samples. This is necessary because by requiring English fluency to complete the survey, our EU sample skews younger and more highly educated.

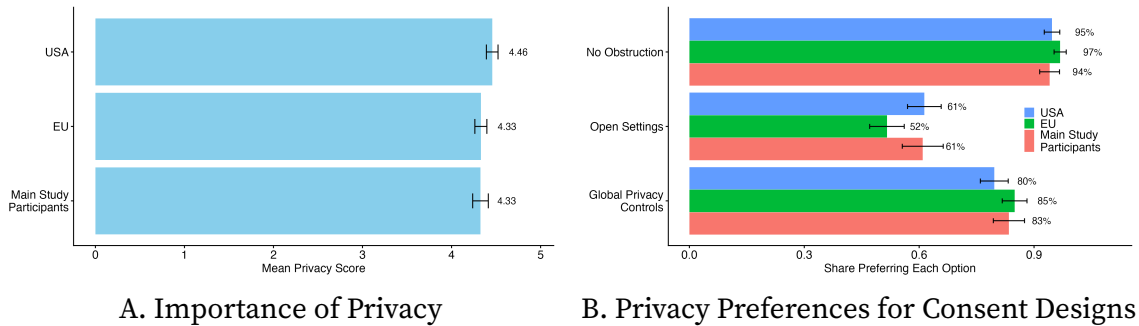
EU. The 320 participants from our prior study are statistically indistinguishable from main study participants who did not answer the new survey (Appendix Table C.1).

To test for order and framing effects, participants were randomly assigned to one of four conditions in a factorial design that varied (a) the order of privacy attitudes and consent-comparison blocks, and (b) whether valuations were elicited as willingness-to-pay (WTA) or willingness-to-accept (WTP). We don't find robust meaningful effects of question order on their stated valuations. In contrast, framing matters substantially, generating a wide gap between reported WTAs and WTPs.

There are only limited differences between prior study participants and the broader survey sample. The left panel of Figure 4 reports respondents' stated valuation for online privacy. On average, participants assign a value of 4.38 (the maximum value of 5 denotes *very important*), with only modest variation across consumer groups. Although some differences become statistically significant after controlling for demographic characteristics (Appendix Table C.2), the corresponding effect sizes remain very small.

There is also a striking similarity across groups in the types of information respondents believe are tracked about them online. Browsing history and device information are the most commonly mentioned, followed by location data and time spent on apps or websites, and then clicks and likes on social media. Demographics, contact information, credit data, health information, and national ID details are mentioned less for all groups; this is consistent with Neumann, Tucker and Whitfield (2019), who show that data brokers are generally unable to predict demographic characteristics with high accuracy.

FIGURE 4. Preferences by Consumer Groups



Notes: Figure plots privacy preferences across three consumer groups: participants in our main study, U.S. survey respondents, and EU survey respondents. All respondents were surveyed in Fall 2025. The left panel reports the average response to a five-point Likert-scale question: “Privacy means different things to different people. In thinking about your online browsing, how important is it for you to be in control of who can access information about you?” A value of 5 corresponds to very important. The right panel reports the share of respondents who prefer (i) a banner with no obstruction over one that hides the “reject all” option, (ii) a banner that displays all settings upfront over one that requires an additional click, and (iii) global privacy controls over their preferred website-level consent banner. Bars indicate 95% confidence intervals.

To assess whether preferences for consent designs differ meaningfully across respondents, we compare their stated choices over alternative banners. Consumers express a clear dislike for designs that deliberately obstruct choices: 95.3% prefer the accept–reject–settings banner over the accept–settings version. Preferences regarding the visibility of cookie settings are mixed, with 57.6% favoring banners that display all settings on the first page. Support for global privacy controls is higher, with 82.6% of respondents rating them comparable (9.3%) or better (73.3%) than their preferred site-level consent banner.

The right panel of Figure 4 shows that these patterns are similar across groups. If anything, European respondents are more likely to prefer designs that make choices easier (e.g., by not exposing all settings immediately) relative to U.S. participants and our main-study sample (Appendix Table C.2). We also asked participants for their willingness to pay (WTP) or willingness to accept (WTA) per month for different consent interfaces. WTP questions yield substantially lower monetary

valuations than WTA questions. Although stated preferences for privacy are notoriously overstated (Athey, Catalini and Tucker 2017; Kokolakis 2017), overall our welfare estimates are closer to the WTP valuations (see Appendix Table C.3 for additional analyses).

6. Discussion

The debate over privacy policy is characterized by competing interests: businesses championing increased data collection versus regulators who assume consumers prioritize data minimization. Contrary to the assumptions underlying much of this discourse, our data show that many consumers willingly share their data even without nudging. That said, consumers' choices to share data are sensitive to how consent is presented. Deliberate obstruction, which makes privacy-preserving options harder to access, is the design pattern with the largest effects among the designs we study. Visual manipulations such as reordering or highlighting have limited effects. Our evidence suggests that even though many websites use both obstruction and visual manipulations (Utz et al. (2019)), the benefits of visual manipulations may be small.

Many consumers misunderstand what happens when they close consent banners. This confusion matters because consumers often avoid explicit choices and close consent banners instead. Ambiguity about defaults extends beyond consumers to regulatory enforcement. Data Protection Authorities across EU member states interpret user inaction differently. For example, Italy and Sweden view it as indecision while France and Luxembourg consider it rejection (Bielova, Santos and Gray 2024).

Our welfare analysis reveals large benefits from eliminating site-specific consent banners in favor of browser-level consent management. Browser-level consent management is endorsed by the CPRA and has been implemented by Mozilla's Firefox and the Brave browser. We find this policy change would boost consumer welfare by \$3.7 per user-week relative to the optimal site-specific banner and \$4.6 relative to the U.S. status quo. While seemingly modest at the individual-week level, these gains scale to substantial societal benefits across all internet users.

Our revealed-preference approach assumes that, in the absence of nudges, observed decisions reflect true utility, conditional on observed beliefs. This assumption may not hold if users lack a complete understanding of the consequences of sharing cookies. While we incorporate consumers' belief mistakes about banner closure outcomes in our welfare calculations, broader information gaps may remain. Our survey results suggest that participants have a basic understanding of cookie functionality, but additional information has the potential to further alter consumer privacy choices.

Despite featuring a more naturalistic environment than typical lab studies, our experiment still differs from a fully natural environment in several respects. First, participants' self-selection into our experiment could in principle introduce biases that limit external validity. Encouragingly, a follow-up survey comparing our main study participants with U.S. and EU samples reveals no large differences in privacy valuations, alleviating selection concerns. Second, participants' awareness of being studied may itself introduce biases. Finally, our focus on short-term responses leaves open the possibility of longer-term behavioral adaptations.

Overall, our findings underscore the need for privacy policies grounded in empirical evidence of actual consumer behavior rather than assumptions about privacy preferences. As digital markets evolve, understanding how consumers

actually make privacy decisions, rather than how they should make them, will prove essential for designing effective and welfare-enhancing privacy regulations.

References

- Acemoglu, Daron, Ali Makhdoumi, Azarakhsh Malekian, and Asu Ozdaglar.** 2022. “Too much data: Prices and inefficiencies in data markets.” *American Economic Journal: Microeconomics*, 14(4): 218–256.
- Acquisti, Alessandro, Leslie K John, and George Loewenstein.** 2013. “What is privacy worth?” *The Journal of Legal Studies*, 42(2): 249–274.
- Aridor, Guy, et al.** 2024. “Evaluating The Impact of Privacy Regulation on E-Commerce Firms: Evidence from Apple’s App Tracking Transparency.” SSRN Working Paper 4698374.
- Athey, Susan, Christian Catalini, and Catherine Tucker.** 2017. “The digital privacy paradox: Small money, small costs, small talk.” *NBER Working Paper No. 23488*.
- Baviskar, Sagar, Iffat Chowdhury, Daniel Deisenroth, Beibei Li, and D Daniel Sokol.** 2024. “ATT vs. Personalized Ads: User’s Data Sharing Choices Under Apple’s Divergent Consent Strategies.” *Personalized Ads: User’s Data Sharing Choices Under Apple’s Divergent Consent Strategies (June 28, 2024)*.
- Bergemann, Dirk, and Alessandro Bonatti.** 2019. “Markets for information: An introduction.” *Annual Review of Economics*, 11(1): 85–107.
- Bergemann, Dirk, and Alessandro Bonatti.** 2024. “Data, competition, and digital platforms.” *American Economic Review*, 114(8): 2553–2595.
- Bernheim, B Douglas, Kristy Kim, and Dmitry Taubinsky.** 2024. “Welfare and the Act of Choosing.” National Bureau of Economic Research.
- Bian, Bo, Michaela Pagel, Huan Tang, and Devesh Raval.** 2023. “Consumer surveillance and financial fraud.” National Bureau of Economic Research.
- 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).

Bielova, Nataliai, Laura Litvine, Anysia Nguyen, Mariam Chammat, Vincent Toubiana, and Estelle Harry. 2024. “The Effect of Design Patterns on (Present and Future) Cookie Consent Decisions.”

Campbell, James, Avi Goldfarb, and Catherine Tucker. 2015. “Privacy regulation and market structure.” *Journal of Economics & Management Strategy*, 24(1): 47–73.

Chen, Zhijun. 2025. “Paying Consumers for Their Data: An Economic Analysis of Data Acquisition and Digital Privacy.” *Available at SSRN 5109828*.

Collis, Avinash, Alex Moehring, Ananya Sen, and Alessandro Acquisti. 2021. “Information frictions and heterogeneity in valuations of personal data.” *Available at SSRN 3974826*.

D’Assergio, Caterina, Puneet Manchanda, Elisa Montaguti, and Sara Valentini. 2022. “The race for data: Gaming or being gamed by the system?” *Available at SSRN 4250389*.

Decarolis, Francesco, Muxin Li, and Filippo Paternollo. 2023. “Competition and Defaults in Online Search.” *Available at SSRN 4660406*.

Deisenroth, Daniel, Utsav Manjeer, Zarak Sohail, Steven Tadelis, and Nils Wernerfelt. 2024. “Digital advertising and market structure: Implications for privacy regulation.” National Bureau of Economic Research.

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.” 1–14.

Farronato, Chiara, Andrey Fradkin, and Chris J. Karr. 2024. “Webmunk: A New Tool for Studying Online Consumer Behavior.” White paper.

Goldberg, Samuel G, Garrett A Johnson, and Scott K Shriver. 2024. “Regulating privacy online: An economic evaluation of the GDPR.” *American Economic Journal: Economic Policy*, 16(1): 325–358.

Goldfarb, Avi, and Catherine E Tucker. 2011. “Privacy regulation and online advertising.”

Management science, 57(1): 57–71.

Goldfarb, Avi, and Catherine E Tucker. 2024. *The Economics of Privacy*. University of Chicago Press.

Habib, Hana, Megan Li, Ellie Young, and Lorrie Cranor. 2022. ““Okay, whatever”: An evaluation of cookie consent interfaces.” 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 A, Scott K Shriver, and Samuel G Goldberg. 2023. “Privacy and market concentration: intended and unintended consequences of the GDPR.” *Management Science*, 69(10): 5695–5721.

Johnson, Garrett, Tesary Lin, James C Cooper, and Liang Zhong. 2024. “COPPAcalypse? The Youtube Settlement’s Impact on Kids Content.”

Kahneman, Daniel, and Richard Thaler. 1991. “Economic analysis and the psychology of utility: Applications to compensation policy.” *The American Economic Review*, 81(2): 341–346.

Kamakura, Wagner A, and Gary J Russell. 1989. “A probabilistic choice model for market segmentation and elasticity structure.” *Journal of marketing research*, 26(4): 379–390.

Ke, T Tony, and K Sudhir. 2023. “Privacy rights and data security: GDPR and personal data markets.” *Management Science*, 69(8): 4389–4412.

Kokolakis, Spyros. 2017. “Privacy attitudes and privacy behaviour: A review of current research on the privacy paradox phenomenon.” *Computers & security*, 64: 122–134.

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.
- Miklós-Thal, Jeanine, Avi Goldfarb, Avery Haviv, and Catherine Tucker.** 2024. “Frontiers: digital hermits.” *Marketing Science*, 43(4): 697–708.
- 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*.
- Neumann, Nico, Catherine E Tucker, and Timothy Whitfield.** 2019. “Frontiers: How effective is third-party consumer profiling? Evidence from field studies.” *Marketing Science*, 38(6): 918–926.
- Nissenbaum, Helen.** 2004. “Privacy as contextual integrity.” *Wash. L. Rev.*, 79: 119.
- Nouwens, Midas, Iaria Liccardi, Michael Veale, David Karger, and Lalana Kagal.** 2020. “Dark patterns after the GDPR: Scraping consent pop-ups and demonstrating their influence.” 1–13.
- Posner, Nathaniel, Andrey Simonov, Kellen Mrkva, and Eric J Johnson.** 2023. “Dark defaults: How choice architecture steers political campaign donations.” *Proceedings of the National Academy of Sciences*, 120(40): e2218385120.
- Que, Verina F.** 2017. “How Do Past Privacy Choices Shape the Future?” *Working Paper*.
- Rhodes, Andrew, and Jidong Zhou.** 2024. “Personalization and privacy choice.” *Available at SSRN 4795118*.
- Spiegler, Ran.** 2014. “Competitive framing.” *American Economic Journal: Microeconomics*,

6(3): 35–58.

Tang, Huan. 2023. “The value of privacy: Evidence from online borrowers.” *Available at SSRN*.

Tomaino, Geoff, Klaus Wertenbroch, and Daniel J Walters. 2023. “Intransitivity of consumer preferences for privacy.” *Journal of Marketing Research*, 60(3): 489–507.

Train, Kenneth E. 2009. *Discrete choice methods with simulation*. Cambridge university press.

Train, Kenneth E. 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.” 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.” 107–121.

Wernerfelt, Nils, Anna Tuchman, Bradley T Shapiro, and Robert Moakler. 2025. “Estimating the value of offsite tracking data to advertisers: Evidence from meta.” *Marketing Science*, 44(2): 268–286.

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:

$$Y_{ij} = \gamma_{reject\ hidden_{ij}} + \gamma_{accept\ hidden_{ij}} + \gamma_{accept\ on\ top_{ij}} + \gamma_{reject\ on\ top_{ij}} + \gamma_{highlight\ accept_{ij}} + \mu_i + \nu_{c(j)} + \epsilon_{ij}, \quad (\text{A.1})$$

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 taxonomy³⁸ and extended to the full set of websites in our study using ChatGPT.³⁹

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 baseline interface (Treatment C), while $\gamma_{reject\ on\ top_{ij}}$ captures the analogous effect for prioritizing "reject all." The remaining γ coefficients reflect the

³⁸<https://www.ghostery.com/>, accessed November 2025.

³⁹We 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'.*

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 A.1 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.

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

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 (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 percentage points (column 5, a 14% increase).

Table A.2 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 baseline condition, 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 A.2. 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 A.1, where the outcome is whether the user selects a subset of cookies. Otherwise the table is identical to Table A.1. * $p < 0.05$, ** $p < 0.01$; *** $p < 0.001$.

Table D.5). 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 see more consent banners. 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 banners.

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.

$$y_{ij} = \beta_{10 \text{ minutes}} + \gamma * \text{time in study}_{ij} + \nu_{c(j)} + \epsilon_{ij}. \quad (\text{B.1})$$

The baseline is the condition where a user sees the banner every 60 minutes, while the alternative condition displays a banner 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.⁴⁰

Table B.1 displays the results. We highlight two findings. First, we do not find a differential impact of banner 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 banner 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

⁴⁰Adding this covariate does not affect whether we detect any treatment effects.

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.

TABLE B.1. 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 B.1, where ‘10 Min Pop-up’ is an indicator for whether the user was in the treatment where banners occurred at a frequency of once every 10 minutes. * $p < 0.05$, ** $p < 0.01$; *** $p < 0.001$.

Second, time spent in the study has an effect on choice. Each additional day in the study increases the share of people closing the banner 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 Table B.2 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.

TABLE B.2. 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 B.1, which removes the banner frequency treatment and adds the order of which a domain is visited (“Visit Order”) and additional fixed effects. * $p < 0.05$, ** $p < 0.01$; *** $p < 0.001$.

Appendix C. Follow Up Survey

In fall 2025, we ran a supplemental Prolific survey across three groups: participants from our main study, additional U.S. respondents, and EU respondents. The survey included three parts: demographics, questions about privacy attitudes and online behavior, and three pairwise comparisons of consent-interface designs. To examine order and framing effects, respondents were randomized in a 2×2 design that varied the sequence of the privacy and interface blocks and whether valuations were elicited using willingness-to-pay or willingness-to-accept questions. We recruited 320 participants to the main study, 471 new U.S. participants, and 471 EU participants. On average, participants took 5.3 minutes to complete the survey. We paid \$1.20 for each completed survey. The full survey is presented in Appendix E. This appendix presents tables and plots that summarize the main findings.

TABLE C.1. Selection into Follow-Up Survey

	Accept all (1)	Privacy importance (2)	Global Privacy Controls (3)	WTP for Global Privacy Controls (4)
Constant	0.602*** (0.017)	3.265*** (0.041)	0.872*** (0.091)	3.814*** (0.465)
Respondent to New Survey	-0.038 (0.027)	0.034 (0.059)	0.055 (0.127)	0.404 (0.706)
Mean of Y	0.58	3.28	0.9	4.01
R ²	0.004	0.001	0.000	0.001
Observations	563	483	484	484

Notes: The table reports linear regressions comparing main study participants who completed the Fall 2025 follow-up survey with those who did not. All outcomes are measured at the time of the main study. The key independent variable is an indicator for completing the follow-up survey. Column (1) corresponds to the average share of accept all clicks recorded by the browser extension during the main study. Column (2) corresponds to responses on a five-point Likert scale to the question: “Privacy means different things to different people. In thinking about all of your online browsing, please state how important it is for you to be in control of who can get information about you.” A value of 5 corresponds to *very important*. Column (3) is a dummy for whether the respondent said they prefer global privacy controls to website-level consent banners. Column (3) corresponds to the willingness to pay for a global-privacy-control tool. The number of observations in column (1) reflects the total number of participants in the main study (line 9 in Table D.2). All other columns include only participants who completed the outtake survey of the main study (line 10 in Table D.2). * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

TABLE C.2. Online Privacy—Stated Preferences

	Privacy importance (1)	No Obstruction (2)	Open Settings (3)	Global Privacy Controls (4)
EU	0.002 (0.057)	0.008 (0.015)	-0.100*** (0.038)	0.009 (0.030)
Main Study Participant	-0.096* (0.055)	-0.011 (0.017)	0.005 (0.037)	0.022 (0.029)
Privacy first	0.067 (0.042)	0.005 (0.012)	0.078*** (0.028)	-0.015 (0.022)
Mean of Y	4.38	0.95	0.58	0.83
R ²	0.075	0.044	0.042	0.041
Observations	1,259	1,259	1,259	1,259

Notes: The table reports linear regressions of outcomes capturing online privacy preferences. Column (1) corresponds to responses on a five-point Likert scale to the question: “Privacy means different things to different people. In thinking about all of your online browsing, please state how important it is for you to be in control of who can get information about you.” A value of 5 corresponds to *very important*. Columns (2)-(4) show outcomes reflecting respondents’ preferences across three consent design comparisons. Each outcome is a binary indicator equal to one if the respondent preferred the second option in the respective pairwise comparison. Column (2) corresponds to the comparison between a banner with only “Accept all cookies” and “Settings” options (acc-set) versus one with both “Accept all” and “Reject all” options (acc-rej-set). Column (3) compares the acc-rej-set banner to a version where all settings options are visible on the first page. Column (4) compares respondents’ preferred cookie banner from these designs to a global privacy control, as in the outtake survey from the main study. *EU* is a dummy for whether a respondent is from the European Union; *Main Study Participant* indicates whether the respondent participated in our main study (the reference group is U.S. respondents who did not participate in the main study); and *Privacy first* equals one if the privacy block appeared before the banner comparison block (the order was randomized). All specifications include fixed effects for age, race/ethnicity, household size, presence of children in the household, education, and income group. Heteroskedasticity-robust standard errors are reported in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

TABLE C.3. Valuations for Consent Designs

	Willingness to Pay			Willingness to Accept		
	No Obstruction (1)	Open Settings (2)	Global Privacy Controls (3)	No Obstruction (4)	Open Settings (5)	Global Privacy Controls (6)
EU	2.053 (5.542)	-0.543 (1.124)	3.642* (2.137)	27.779 (20.198)	-1.311 (4.577)	14.892** (6.805)
Main Study Participant	-3.067 (2.574)	-0.488 (0.564)	-0.495 (0.658)	5.552 (22.817)	4.556 (4.653)	-0.884 (6.299)
Privacy first	-0.577 (2.958)	-0.865 (0.734)	-1.060 (1.389)	22.011 (14.710)	3.866 (3.415)	3.634 (5.368)
Mean of Y	3.990	1.050	3.370	74.790	7.880	23.820
R ²	0.036	0.040	0.044	0.060	0.043	0.064
Observations	631	643	644	563	551	552


Notes: The table reports linear regressions of respondents' stated monetary valuations for different consent designs. Each outcome corresponds to the amount a respondent is either willing to pay for their preferred banner over the alternative (columns 1–3) or willing to accept to be assigned their least preferred banner (columns 4–6). Positive values indicate that the second option in each pairwise comparison is valued more highly, while negative values indicate that the first option is valued more highly. Observations with valuations in the extreme tails (top 2.5% and bottom 2.5%) are removed. The first three columns correspond to willingness-to-pay (WTP) valuations for the three pairwise comparisons: (1) banners with “Accept all” and “Settings” options (acc–set) versus banners with “Accept all” and “Reject all” options (acc–rej–set); (2) acc–rej–set banners versus acc–rej–set banners with all settings visible on the first page; and (3) respondents' preferred banner versus a global privacy control. The last three columns repeat these comparisons for willingness-to-accept (WTA) valuations. We report WTP and WTA separately given the large differences in the magnitude of stated values. *EU*, *Main Study Participant*, and *Privacy first* are defined as in Table C.2. All specifications include fixed effects for age, race/ethnicity, household size, presence of children in the household, education, and income group. Heteroskedasticity-robust standard errors are reported in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

TABLE C.4. Belief about Advertisers' Tracking

	Device (1)	Browsing (2)	Location (3)	Time (4)	Clicks & Likes (5)	Demogr. (6)	Contact Info (7)	Credit (8)	Health (9)	National ID (10)
EU	0.064** (0.027)	-0.059** (0.028)	0.039 (0.031)	0.034 (0.035)	0.013 (0.037)	-0.004 (0.039)	-0.062* (0.033)	-0.053*** (0.015)	-0.013 (0.015)	0.001 (0.014)
Main Study Participant	0.027 (0.027)	0.021 (0.025)	0.044 (0.030)	0.075** (0.033)	0.047 (0.035)	0.084** (0.037)	0.019 (0.034)	0.010 (0.019)	0.034* (0.018)	0.006 (0.014)
Privacy first	0.020 (0.020)	-0.004 (0.021)	-0.008 (0.023)	0.012 (0.026)	-0.024 (0.027)	-0.052* (0.029)	-0.031 (0.025)	0.003 (0.012)	0.005 (0.012)	0.005 (0.009)
Mean of Y	0.85	0.84	0.79	0.72	0.66	0.54	0.27	0.05	0.04	0.03
R ²	0.039	0.033	0.039	0.028	0.035	0.030	0.033	0.040	0.024	0.027
Observations	1,259	1,259	1,259	1,259	1,259	1,259	1,259	1,259	1,259	1,259

Notes: The table reports linear regressions of outcomes measuring participants' beliefs about what types of information advertisers typically know about them. Each column corresponds to one option from the survey question: "As you browse the internet, which of the following do you think advertisers typically know about you? (Select all that apply)." The outcomes are binary indicators equal to one if a respondent selected that option. Specifically, column (1) corresponds to "The device I'm using (phone, computer, etc.)," column (2) to "Websites I visit and things I search for (browsing history)," column (3) to "My location," column (4) to "How much time I spend on websites or apps," column (5) to "My likes, clicks, and interactions on social media," column (6) to "My age, gender, or other demographic information," column (7) to "My contact information (email, phone number, etc.)," column (8) to "My credit score," column (9) to "My health information," and column (10) to "My national identification number (for example, social security number)." The columns are ordered from most likely to least likely to be selected. *EU* is a dummy for whether a survey respondent is from the European Union; *Main Study Participant* is a dummy for whether the respondent participated in our main study (so the benchmark group is U.S. respondents who did not participate in the main study); and *Privacy first* is a dummy for whether the privacy block was shown before the banner comparison block (the order was randomized). All specifications include fixed effects for age, race/ethnicity, household size, presence of children in the household, education, and income group. Heteroskedasticity-robust standard errors are reported in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

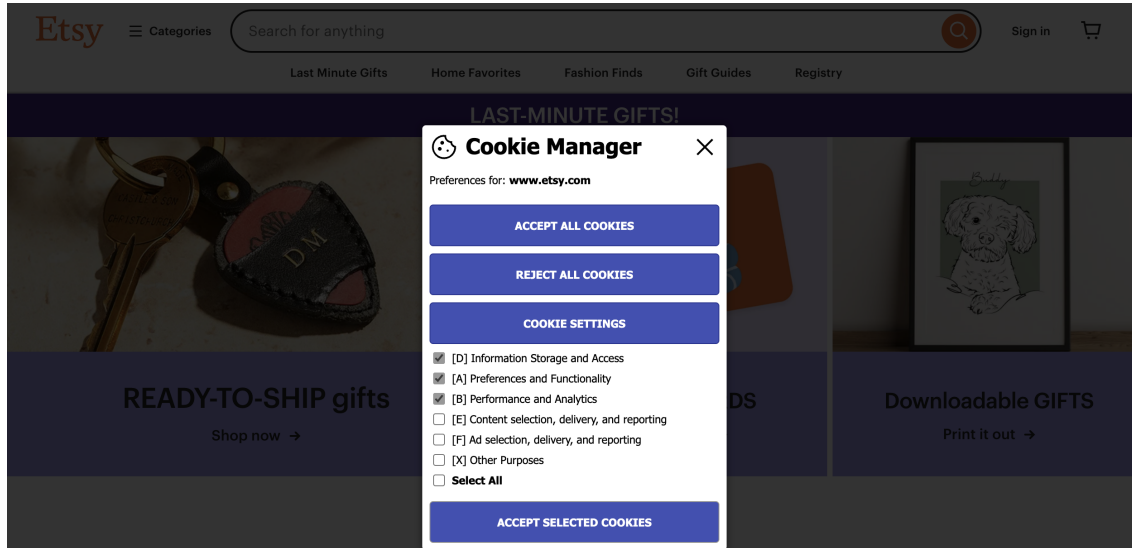
FIGURE C.1. Belief By Consumer Groups



Notes: The figure reports the share of respondents who believe advertisers typically know various pieces of information about them. Each column corresponds to one option from the survey question, “As you browse the internet, which of the following do you think advertisers typically know about you?” with outcomes coded as binary indicators. Columns span (1) device type, (2) browsing history, (3) location, (4) time spent on websites/apps, (5) social-media interactions, (6) demographic information, (7) contact information, (8) credit score, (9) health information, and (10) national ID number; columns are ordered by frequency of selection. "USA" identifies non-main-study US respondents, "EU" identifies European respondents, "Main Study Participants" indicates respondents who took part in the main study. Error bars show 95% confidence intervals.

Appendix D. Additional Exhibits

FIGURE D.1. Cookie Manager's User Interface



Notes: Consent interface when a user clicks “cookie settings”. 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 treatments, meaning they can always choose to reject or select all cookies upon coming to the “settings” menu.

TABLE D.1. 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 D.2. 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	917	74.74				
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 banners appear every 10 minutes of browsing), and 60 minutes (where cookie banners 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 D.3. 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 D.4. 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

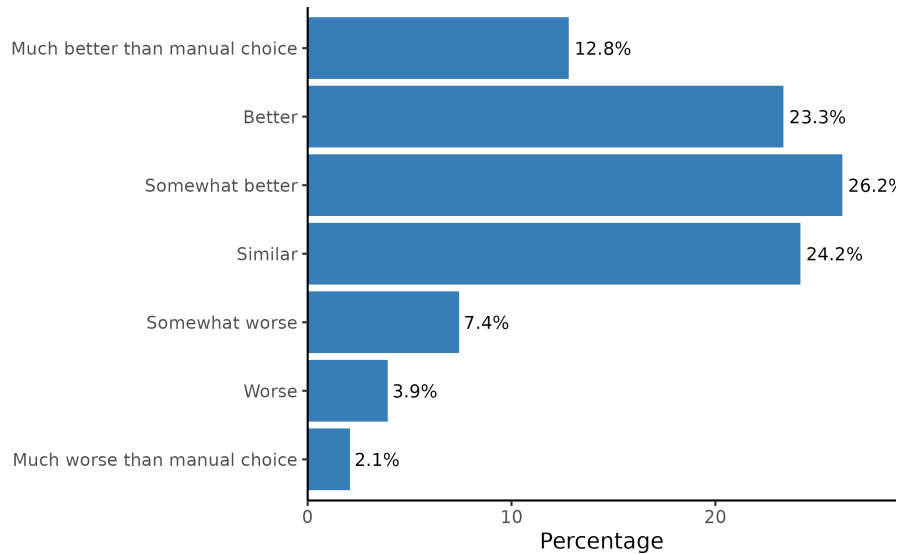
Notes: Banner frequency is randomized at the user level. We therefore exclude domain rank during the organic browsing, but include the number of banners exposed at the survey stage for covariate balance checks. * $p < 0.05$, ** $p < 0.01$; *** $p < 0.001$.

TABLE D.5. 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.

FIGURE D.2. User Preferences for Global Privacy Control




Notes: The figure plots the distribution of answers to the following question in the endline survey: Consider a tool that allows you to specify how you would like to answer cookie consent questions online. This tool will then automatically hide all cookie banners and answer them in the 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 D.6. Selective Cookie Choice by Experimental Condition

	Survey	Organic
	Accept Some (1)	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. * $p < 0.05$, ** $p < 0.01$; *** $p < 0.001$.

FIGURE D.3. Treatment Effects by Survey Domain (Acc-Set vs Baseline Banner)



Notes: The figure displays the treatment effects (point estimates and 95% confidence intervals) of the Acc-Set condition relative to the baseline 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 gray) denotes a different outcome.

TABLE D.7. Heterogeneity of Dark Pattern Effect by Prior Visit

	Accept All (1)	Reject All (2)	Close Window (3)
Has Prior Visit × Rej-Acc-Set	0.028 (0.020)	-0.009 (0.015)	-0.002 (0.015)
Has Prior Visit × Acc-Set	0.007 (0.018)	0.038** (0.014)	-0.030 (0.016)
Has Prior Visit × Acc-GrRej-GrSet	0.011 (0.018)	-0.005 (0.013)	0.001 (0.015)
Has Prior Visit × Acc-Rej-Set	0.017 (0.019)	-0.008 (0.014)	-0.008 (0.016)
Has Prior Visit × Rej-Set	0.011 (0.022)	-0.044* (0.018)	0.022 (0.020)
Has Prior Visit	0.015 (0.015)	-0.029* (0.012)	0.006 (0.012)
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. * $p < 0.05$, ** $p < 0.01$; *** $p < 0.001$.

TABLE D.8. Heterogeneity of Dark Pattern Effect by Domain Popularity

	Accept All (1)	Reject All (2)	Close Window (3)
Domain Rank (Log 10) × Rej-Acc-Set	-0.002 (0.005)	-0.001 (0.004)	0.001 (0.004)
Domain Rank (Log 10) × Acc-Set	-0.007 (0.005)	0.004 (0.004)	0.005 (0.004)
Domain Rank (Log 10) × Acc-GrRej-GrSet	0.000 (0.005)	-0.003 (0.004)	0.002 (0.004)
Domain Rank (Log 10) × Acc-Rej-Set	-0.002 (0.005)	0.002 (0.004)	0.000 (0.004)
Domain Rank (Log 10) × Rej-Set	0.005 (0.006)	-0.015** (0.005)	0.013* (0.006)
Domain Rank (Log 10)	-0.010* (0.004)	0.007* (0.003)	0.003 (0.003)
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 D.9. 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. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

TABLE D.10. 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 D.11. Choice Variation Decomposition (Outcome: “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 D.12. Consumer Surplus Under Counterfactual Policies (Utility Scale Results)

Counterfactual	Average	Type 1	Type 2	Type 3
		Acceptors	Rejectors	Discerners
U.S. Status Quo	1.87	3.29	-2.79	2.35
EU Norm	2.47	3.23	0.97	2.24
Optimal Banner Interface	2.52	3.57	-0.14	2.49
Optimal Banner Interface (with correct beliefs)	2.78	3.79	0.64	2.55
Global Accept	5.69	8.47	0.4	4.79
Global Reject	4.9	5.81	5.36	3.58
Global Privacy Control	6.56	8.47	5.36	4.79

Notes: The values represent unscaled consumer surplus *per choice* under various counterfactual policies in the utility scale. “Pooled estimate” refers to the estimate across all subjects, and the other columns correspond to subset-specific estimates. “U.S. status quo” refers to an accept-settings interface, combined with an accept-all default when consumers close window; “CS maximizing” refers to a baseline interface with an accept-all default when consumers close window; “EU norm” refers to a baseline interface with a reject-all default when consumers close window. “Global Accept” and “Global Reject” force each individual to either always accept or always reject all cookies. The last row, “Global Privacy Control,” allows each individual to make their preferred global choice.