Bocconi

Consent and Choice: Experimental Evidence on User Behavior at Consent Banners

Working Paper N. 289
October 2026

Cristiana Firullo
Cristobal Cheyre
Francesco Decarolis
Alessandro Acquisti

BAFFI
CENTRE ON ECONOMICS,
FINANCE AND REGULATION

Working Paper Series

Consent and Choice: Experimental Evidence on User Behavior at Consent Banners*

Cristiana Firullo[1]
Cornell

Francesco Decarolis
Bocconi and CEPR

Cristobal Cheyre
Cornell

Alessandro Acquisti
MIT

Abstract

Cookie banners appear on websites when visitors first arrive, informing them about a site's use of cookies and tracking technologies, and often asking them whether they consent to their data being collected and shared with third parties. We use a panel of 789 U.S. Chrome users, whose online behaviors were tracked by a browser extension they kept installed for up to three months, to investigate how users interact with cookie banners in the wild, over time, across all the websites they visit. We find that users engage with only one in three banners they encounter, and do so mostly on websites that force a response; these websites register three times more banner engagement than the ones where banners can be ignored, but lose more than twice as many visitors after their encounter with the banner. Users' explicit choices (accepting or rejecting tracking) exhibit across-user heterogeneity, and are significantly influenced by whether a one-click reject option is offered. We set a single signal from the choices users make during the first third of the days we observe them, and find that this signal is consistent with about two thirds of their later encounters. We calculate that users would spend 58 percent less time on banners if websites could not force a response.

JEL codes: D18, D83, L51, L86.

Keywords: privacy, consent banners, dark patterns, GDPR.

*The Internet Behavior Experiment was supported by the Alfred P. Sloan Foundation, the Carnegie Mellon University Block Center for Technology and Society, the Carnegie Mellon University CyLab Security and Privacy Institute and the Program on the Economics & Privacy at George Mason University, the University of Pennsylvania Center for Technology, Innovation and Competition and the Warren Center for Network & Data Sciences, the NET Institute, and the Pittsburgh Supercomputing Center. Cristobal Cheyre acknowledges support from the Cornell Center for Social Sciences. Cristobal Cheyre and Alessandro Acquisti acknowledge support from the National Science Foundation (Awards 2237327/2237328/2237329). Francesco Decarolis and Cristiana Firullo gratefully acknowledge financial support from Amazon Europe Core S.à r.l and Clifford Chance for the analytical contributions regarding the legal and regulatory aspects of the study. None of the funders had any role in the study design, data collection, analysis, or decision to publish.

[1]Email: cf528@cornell.edu.

1. Introduction

Consent banners are the current way websites inform users and ask them for consent with their tracking and data collection strategies. The origins of their implementation can be found in the required compliance with Article 5(3) of the EU ePrivacy Directive which, as amended in 2009, requires consent before a website places a cookie (or similar technology) that is not strictly necessary for the service the user requested (Article 29 Data Protection Working Party, 2012; European Data Protection Board, 2024). In 2018, the General Data Protection Regulation (GDPR) supplied the standard that consent must meet.[2] Its stricter standard led to a sharp increase in the number of websites showing a banner and in the use of consent management platforms (Degeling et al., 2019; Hils et al., 2020). Banners have since spread beyond the EU. In the United States, no federal rule requires consent to cookies, but several states have introduced notice and opt-out requirements that websites often implement through banners,[3] and websites also adopt banners for self-regulatory reasons. The European requirement of prior consent nevertheless remains the most demanding standard and the reference point for the consent tools deployed worldwide (Bradford, 2020; Sanchez-Rola et al., 2019): most websites follow the same consent practice for visitors inside and outside the EU (van Eijk et al., 2019).

Banner design strategies and users' consent choices are consequential for both websites and their visitors: they can affect websites' ability to monetize user data, as well as the privacy trade-offs their visitors will face. In theory, user engagement with consent banners should reflect their preferences over the expected trade-offs of data sharing. In practice, behavioral research has highlighted how online privacy choices can be affected by non-normative factors (Acquisti et al., 2015, 2016), such as interface design features nudging users toward accepting or rejecting tracking (Utz et al., 2019; Nouwens et al., 2020). While some effort has been devoted to estimating the cost of time spent engaging with cookie banners (European Commission, 2025), and the impact of such engagement on consumer welfare (Farronato et al., 2026), no study to date has empirically captured users' natural interactions with the consent banners they encounter in the wild, over time, during

[2] Regulation (EU) 2016/679, Arts. 4(11), 7 and 94(2); Case C-673/17, *Planet49*, EU:C:2019:801; European Data Protection Board (2020). See Directive 2002/58/EC, Art. 5(3), as amended by Directive 2009/136/EC.

[3] E.g., Cal. Civ. Code §§ 1798.120 and 1798.135; Cal. Code Regs. tit. 11, § 7025 (Global Privacy Control).

their organic browsing, and thus the real world effect of websites' banners on user behavior.

We leverage a panel of 789 U.S.-based Chrome users who participated in the Internet Behavior Experiment (IBE) to investigate their interactions with website banners throughout several weeks. IBE participants install a research extension that logs, among other things, the pages they visit, the consent banners they meet, and the clicks they make on those banners, over a period of up to 90 days (Cheyre et al., 2024). Participants are randomly assigned to one of three experimental conditions: Opt-In (ads are shown, and are likely to be behaviorally targeted, because participants are opted-in to tracking by advertising networks); Anti-tracking (ads are shown, but are significantly less likely to be behaviorally targeted, as the study software opts participants out of ad personalization across major platforms and advertising networks); and Ad-blocking (ads are minimized, as they are blocked to the greatest extent possible by current ad-blocking technology). For every participant, we record banners organically encountered while browsing, together with the banners' text and menus of options. Furthermore, we record what participants do on the websites where they encounter a consent banner: accept tracking, reject it, ignore the banner and keep browsing the website, or leave the website altogether. We complement behavioral data with self-reported survey answers capturing participants' preferences and attitudes, including those regarding privacy and tracking. First, we ask whether users engage with consent banners at all, and whether engagement is explained by the characteristics of the user (such as demographics or privacy preferences), or by those of the environment (such as the category of website and the menu of options offered by the banner), or by the user's familiarity with the website. Second, we study what users choose when they do engage with a banner: whether they accept or reject tracking, and which of the above characteristics can explain that choice. Finally, we ask how consistently a single signal, set once and applied to every website, would reproduce the choices the same user makes, and how much of the time users spend on banners depends on websites that force a response.

Unlike prior related work, which presented participants with banners designed by the researchers (Nouwens et al., 2020; Habib et al., 2022; Bielova et al., 2024; Papenmeier et al., 2025; Farronato et al., 2026), observed study participants on a single website (Utz et al., 2019; Hils et al., 2020), or captured engagement data from the websites themselves

(Usercentrics, 2024; Jha et al., 2025), we study user-banner interactions as they naturally occur in the panel, and we do so over time. The same participant is observed for several weeks, during which they may visit multiple sites with different banners (and the same websites again and again), allowing us to measure how interactions dynamically evolve. We find that participants engage with only one in three banners, and that whether a user engages with a banner depends more strongly on contextual factors (such as the type of website and the menu of options the banner offers) than on user characteristics or preferences. When the website cannot be browsed until the user makes a choice in the banner, for instance, engagement is three times higher. However, forcing user choice is not costless for websites: websites that force that choice lose more than twice as many visitors after the banner as those that do not. Participants are more likely to engage with banners on websites they actively browse (i.e., they navigate multiple pages from the same website), such as laboratory portals or brokerage websites (labcorp.com, fidelity.com), and are more likely to ignore banners on websites they pass through, such as recipes or celebrity news websites (allrecipes.com, people.com). Next, we compare the engagement of the same participant at banners that offer a one-click reject and at banners that do not, and find that banners providing a one-click reject are engaged with more than twice as often as banners without one (47 against 20 percent of encounters); participants reject tracking on 26 percent of the choices they make at these banners, while they never reject at banners that do not provide a one-click option. Furthermore, participants are slightly more likely (but significantly so) to leave the website after encountering a banner without a one-click reject option. We also find that banners whose text mentions personalization are engaged with more often, but accepted at the same rate once engaged with, compared with those that do not. The menu of options shown in the banner affects participants' choices, but within banners with identical options, what participants choose depends on user characteristics: 56 percent of the variation in accepting tracking lies between participants, and one standard deviation of stated privacy concern raises rejecting by 4.3 percentage points, on a base of 16 percent of choices (but it does not predict whether the participant engages with the banner in the first place); participants in the top third of stated concern reject tracking more than twice as often as those in the bottom third. We also find that engagement with banners and banner choices do not differ significantly across the three experimental conditions.

We then examine the stability of consent decisions within-user, across-sites: whether a given participant, when faced with consent banners across multiple websites, tends to repeat the same choice (e.g., always “Accept All"), or whether instead their choices exhibit context-dependent variability by changing across sites. A signal set from the first third of the days we observe a participant is consistent with 90 percent of that participant’s later *explicit* choices (*accept* vs *reject*), and with 65 percent of their later encounters when *ignoring* the banner is counted as an outcome the signal can match. Consistency differs between users who typically accept or reject. The signal is consistent with 94 percent of the later choices of participants who only accepted in that period (three in four participants), and with 77 percent of those of users who both accepted and rejected, against 64 percent for a signal that always accepts and 56 percent for a signal that always rejects. Finally, we study whether self-reported privacy preferences captured at the start of the study are consistent with or deviate from banner choices. We use a validated privacy concerns scale to classify participants into three groups by their self-reported privacy concerns (low, medium and high), and compare the low and high groups with their banner choices. We consider two possible interpretations of a participant’s decision to ignore (that is, not engage with) a banner. Assuming that a participant who ignores a banner believes that, by doing so, they are *not* consenting to tracking, high self-reported privacy concerns are consistent with 75 percent of banner encounters, and low privacy concerns are consistent with 27 percent of banner encounters. Whereas, assuming that a participant who ignores the banner does so believing and accepting that they will be tracked, then low self-reported privacy concerns are consistent with 97 percent of banner encounters, and high self-reported concerns are consistent with 7 percent of banner encounters. Restricting the comparison to explicit choices, participants with high concerns actually *accept* 77 percent of the banners they engage with, those with medium concerns 82 percent, and those with low concerns 91 percent.

In short, our results suggest that observed consent rates are a property of the choice architecture as much as of users’ traits. When visitors can ignore consent requests, most do. Whether they engage with a banner is strongly affected by the website’s design, while what they choose (when they do engage) often reflects stable personal preferences. Such findings speak to three groups of studies. First, a large literature shows that the design of consent interfaces affects consent choices: defaults, dark patterns (designs that steer the

user toward acceptance) and the cost of refusal change acceptance rates (Utz et al., 2019; Nouwens et al., 2020; Habib et al., 2022; Egberts and Krämer, 2026; Sjøflot and Opsahl, 2026). Some designs affect later choices too: in a single session, participants who first saw two banners with no reject button on the first layer refused or customized less often on a third, neutral banner, 10 against 19 percent (Bielova et al., 2024). Among websites using the IAB Transparency and Consent Framework, a cross-industry standard designed to comply with privacy regulations, about half of the banners that allow refusal preselect options, and 10 percent of websites store a positive consent before the user makes any choice (Matte et al., 2020). Bielova (2022) surveys the studies on dark patterns. On a mock website with opt-out banners modeled on the California Consumer Privacy Act (CCPA), Mazumdar et al. (2023) find that 80 percent of 224 U.S. participants ignore the banner and only one opts out, although 32 percent later report having refused the sale of their data. Most of these studies show participants banners designed by the researchers, on one website or on mock websites, and observe each participant during a single visit or session. On a single website, most visitors do not interact with the banner at all (Utz et al., 2019). Two studies follow the same users across the websites they browse, but replace each website's banner with one designed by the researchers: Nouwens et al. (2020) follow 40 participants for up to three weeks, and Farronato et al. (2026) follow participants for one week and find that, under their baseline design, 22 percent of the time participants close the banner (using the 'X' button) without choosing. Consent management platforms record clicks on the websites' own banners, but cannot link a user across websites or record a banner that is ignored (Jha et al., 2025). In an online experiment in which a banner with an accept and a reject button blocked each of twelve template websites until participants chose one, almost four in ten participants changed their choice across websites, while the others always accepted or always rejected cookies (Papenmeier et al., 2025); as participants could not ignore the banner, the experiment does not capture user behavior in the wild, where users can ignore banners. In our study, we capture users' interactions with banners that websites themselves display to participants in their organic browsing, over several weeks (up to three months per participant). Hence we can measure how much of consent behavior is explained by the website's interface rather than by the person, how often the same user ignores these banners across the websites they visit, and how user engagement behavior (including ignoring the banner altogether, or abandoning the website) changes over time. In so doing, we find that *ignore* is the most common user

behavior. Among the participants who make at least five explicit choices during the study period, 61 percent always make the same choice (i.e., either always accept or always reject).

Second, a recent stream of literature studies consent set once in the browser. We can think of centralized consent as a way of shifting user choices from websites to browsers, through a machine-readable signal that websites will read rather than collect. In the European Union, the proposal of the European Commission known as the "Digital Omnibus" (November 2025) would require websites to respect such a signal and browsers to offer a way to set it, and would ask standardization organizations to draft standards for interpreting it, while the text does not specify whether the signal is one setting applied to every website or can differ from one website to the next. In the United States, California already requires websites to honor the Global Privacy Control as an opt-out of the sale or sharing of personal data, and the California Opt Me Out Act, signed in October 2025, requires web browsers to offer a setting that sends such an opt-out signal from January 2027.[4] The Behavioural Insights Team (2024) tests browser-level settings in a simulated browser, and Farronato et al. (2026) randomize the design of consent banners shown to U.S. users by a browser extension and estimate a model of their choices as they browse the Internet for one week. Valuing the time users spend on banners at their wage, they estimate that a single setting in the browser would raise consumer surplus by more than any site-by-site banner design, mainly because it removes the time cost of repeated decisions. Our study differs from these ones in design and in results. Farronato et al. (2026) replace each website's banner with one of six banners designed by the researchers, which appear also on websites without a native banner, at most once every ten or sixty minutes. Participants are told that the banner will appear and that their choices will be enforced, and they browse organically for one week. Their participants can accept, reject, customize or close the banner, but ignoring the banner and leaving the website are not recorded actions. In some of these studies, UI design was effectively preventing the user from ignoring the banner (a behavior that our findings suggest to be the most common response to banner encounters). Instead, we observe the banners that websites themselves display, with their own menus and layouts and whether or not they block page access until a choice is made, during up to three months of browsing in a study that is not presented as being about banners. We record whether participants accept, reject,

[4]Assembly Bill 566 (2025), amending the California Consumer Privacy Act.

ignore the banner and keep browsing, or leave the website.[5] We find that a signal set from the first third of the days we observe each user is consistent with 90 percent of their later explicit choices, and with 65 percent of all their later encounters (when ignoring the banner is counted as an outcome the signal can match).

Third, a stream of literature measures how much tracking is worth to advertisers and publishers, with estimates ranging widely by method (Goldfarb and Tucker, 2011; Marotta et al., 2019; Johnson et al., 2020). Aridor et al. (2023) measure, from a data intermediary's records, how GDPR opt-outs changed the observability and value of consumers. In a field experiment in which Chrome users were randomly assigned to browse without third-party cookies, publisher advertising revenue fell by 29 percent (Gu et al., 2026). This stream of research has also emphasized the impact of data protection regulation on the extent to which platforms can track users, with reductions of up to 80 percent across different EU locations (Prasad et al., 2026). We complement these studies by measuring how users interact with consent banners: how often they engage with them, what they choose, and how many of their later choices a single signal would reproduce. The consent decision is itself informative: in a simulated recommender system, when consent rates differ across demographic groups, a user who withholds consent can reveal more about their attributes than a user who shares their cookies (Miehling et al., 2023). Furthermore, our data allow us to consider possible economic implications of proposals to centralize consent at the browser level. Based on our analysis of user behaviors with consent banners, we estimate that a user who visits as many new websites as our median participant, and for whom half of the websites with a banner force a response, would spend about 42 minutes per year on cookie banners, most of them on websites where ignoring the banner is not an option. If websites were prevented from forcing users to make a choice, this time would fall by 58 percent, to 18 minutes per year. These findings are relevant for the Digital Omnibus proposal of the European Commission (November 2025), which would move the consent question from the banner to a machine-readable signal set in the browser, which websites would have to honor, and which does not specify whether the signal applies to every website or can differ from one website to the next. While our data cannot predict

[5] In Farronato et al. (2026), participants who close the banner with the 'X' button without choosing are recorded as closing it. Closing a banner is a click on the banner, while ignoring it, as we define it, means continuing to browse without clicking on the banner at all. Closing therefore captures only part of the encounters in which a user does not make a choice, and cannot measure how often banners are ignored.

how users would choose under such a setting (a question asked once by the browser is a different choice architecture from the banners we observe, and websites, advertisers and platforms would respond to it in ways that our data do not capture), our findings suggest that the time users spend on consent depends more on whether websites force a response than on how often consent is asked, and that a single signal would reproduce most, but not all, of the choices users make website by website.

The rest of the paper is organized as follows. Section 2 describes the data, the banner encounter sample, and the empirical approach. Section 3 presents the results in two parts: whether people engage with a banner and what they choose. Section 4 studies the consistency of a single signal and discusses possible trade-offs associated with centralized consent mechanisms. Section 5 concludes.

2. Data and Empirical Approach

The Internet Behavior Experiment recruits U.S. Chrome users and invites them to install a research software that includes a browser extension. The extension records the pages participants visit, their navigation events, the consent banners that appear, the cookies set, among other data described in Cheyre et al. (2024). Participants complete a survey at onboarding, monthly, and at exit. In the Opt-In condition, the study software opts participants into various forms of online tracking and behavioral targeting. In the Anti-tracking condition, the software opts participants out of the tracking of the most common third-party services and removes tracking cookies. In the Ad-blocking condition, the software keeps tracking enabled but blocks the advertising. Participants remain in the study for up to three months. At the time of writing, the analysis sample consists of 789 participants who have been recruited into the study between May 2025 and September 2026: 276 browsing under Opt-In, 248 under Anti-tracking, and 265 under Ad-blocking.[6] The sample includes the participants who were paid for completing the onboarding and who were active on at least 14 days. The extension recognizes a banner when the markup of the page matches one of eight consent management platforms (CookieHub, CookieInformation, CookieScript, CookieYes, Cookiebot, Didomi, Iubenda, OneTrust) which jointly cover at least 66 percent of the web traffic in Europe (Nouwens et al., 2025).

[6]The experiment is ongoing and targets a total sample of 1,200 participants. This manuscript will be updated with final estimates.

For every banner, the extension records the text, the options it offers, and the option the participant selects: accept, reject, customize, or none. Repeated records of the same ignored banner within a session are collapsed into one encounter.[7] The extension may miss some banners, for instance those of a platform outside the eight or those on a page the participant leaves before the banner appears, so the share of websites where we record a banner is a lower bound.

Stated preferences are captured through the onboarding survey, before participants are assigned to an experimental condition. Captured measures include, among others, privacy concerns (measured as the average agreement with the eight items of the validated short version of the Internet Users' Information Privacy Concerns scale (Malhotra et al., 2004; Groß, 2021)), whether tracking is useful, whether targeted advertising benefits consumers, and whether people would soon realize they do not miss targeted ads if ads stopped being targeted. In the analysis, we divide participants into three groups of similar size by their privacy concern: low, medium and high. The groups are the terciles of the concern average in the analysis sample: low concern is an average of 5.5 or less, medium concern an average above 5.5 and up to 6.125, and high concern an average above 6.125 (297, 229 and 263 participants).[8] Familiarity with a website is the number of prior visits by a participant to that website during their tenure in the panel. How widely visited a website is, measured as the number of participants who visit it during the panel, is used to control for the likelihood that a user may have already known a website before joining the panel. Website categories come from a hand-checked LLM-classified taxonomy of the registered domain (Appendix A.4), for example shopping (e.g., cvs.com, ikea.com), news and lifestyle (e.g., people.com, allrecipes.com), productivity, meaning work and account tools (e.g., zoom.us, calendly.com), health (e.g., labcorp.com, aetna.com), finance (e.g., fidelity.com, cash.app) and streaming (e.g., siriusxm.com, hbomax.com).

[7] Banners that offer no choice at all, such as notice banners with a close button as the only option, are excluded from the analysis: 74 banners on 68 websites, for a total of 391 encounters out of 16,801. We also exclude encounters on paid-survey and reward websites (138 domains, led by three reward platforms whose visits reflect paid work rather than browsing, and whose banners are almost never engaged with). Including those websites lowers the share of banners engaged with, 27.7 against 30.8 percent (Appendix A.7).

[8] We use terciles rather than a fixed level of agreement because 86 percent of participants have a concern average of 5 or more, so a group of participants who disagree with the statements would be too small to compare. We also repeat the analysis using estimated concerns as a continuous measure.

We classify each encounter based on the action taken by our participants. On a first level, at every encounter, a participant can either engage or not engage with the banner. We define as *engagement* the encounters in which users make an explicit choice, such as 'Accept', 'Reject', even though banners may also display other options such as 'Customize' settings or 'Dismiss'.[9] When instead the participant does not select an option (i.e., does not engage), we look at the first different page they open in the same browser tab. If they click a link on the page, we classify the encounter as *ignoring* the banner while staying on the page, since they used the page while not engaging with the banner. If they go back or type another address and do not return to the website within thirty minutes, we classify the encounter as *leaving* the website. Therefore, we study the user choice at each encounter in two steps: first, whether to engage with the banner or not, where not engaging includes both ignoring the banner and leaving the website; second, among the banners the participant engages with, whether they accept or reject. We keep the two levels apart and we further explore whether they are explained by different processes and characteristics, for example the website's type, the participant's familiarity with the website, the banner's menu, the participant's stated attitudes and demographics, and the experimental condition. One important example is given by websites that force a choice on banners, as their page cannot be used until the participant clicks on one available option. In that case, if our extension does not record an explicit choice, then users must have left the website.[10]

3. Results

Our unit of observation is a banner encounter: one participant on one website within one browsing session, where a browsing session is a stream of page visits, on any website, with no more than thirty minutes between one visit and the next. The sample currently includes 16,410 banner encounters, and Appendix Table A1, Panel B, describes the menus of the banners met in each condition. In a typical week, over 6.3 active days, the median

[9] If a participant opens the settings (e.g., 'Customize' settings), we count the encounter as engagement only if the participant thereafter clicks the accept or reject button of the banner within thirty minutes. Choices made inside the settings panel are not recorded, hence, lacking an explicit accept or reject choice, these encounters are classified as "ignoring" the banner.

[10] We identify those banners that force a choice if (1) the web page cannot be browsed before a user makes a choice on the consent banner and (2) the options available on the banner do not include a 'Dismiss' option (e.g., the button 'X'), as in that case the user could just access the page without making an explicit choice (our definition of engagement). Appendix A.3 describes how we measure it.

participant views 805 pages on 48 distinct websites, and encounters 1.3 banners (most websites do not show a banner again once it has been engaged with; Figure 4).[11]

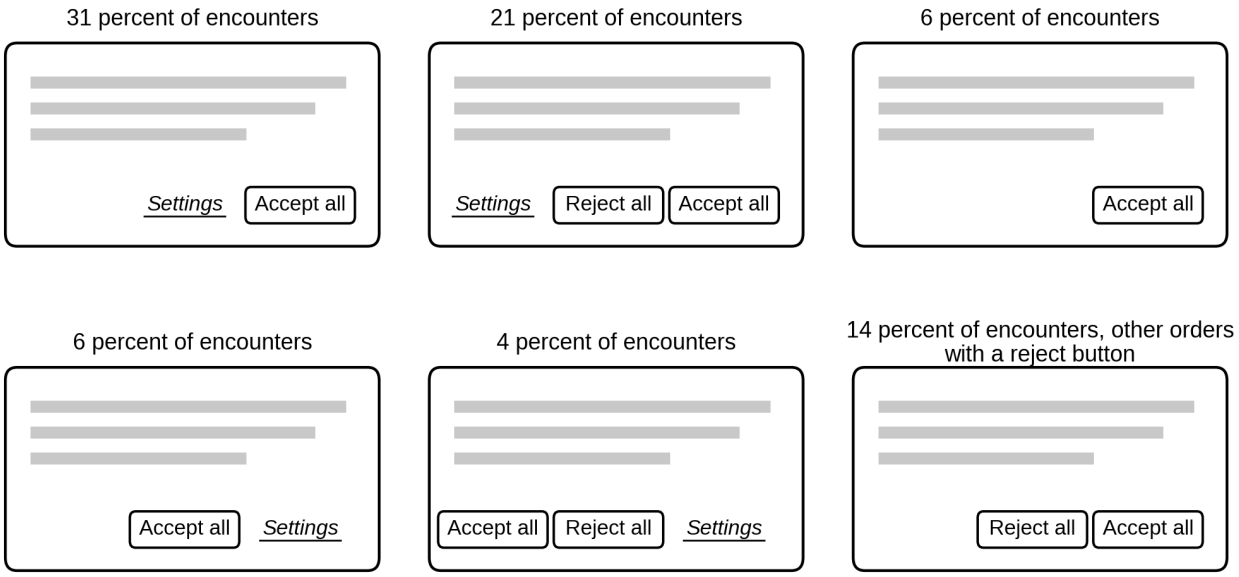

FIGURE 1. The menus of the banners participants met. Construction in Appendix A.2.

Over the whole panel, the median participant meets five banners for every hundred distinct websites visited. Figure 1 shows the menus participants met, by the buttons offered and their order on the page. Whether a banner offers a one-click reject or a customize option is read from the recorded banner text with fixed rules: 40 percent of encounters offer a one-click reject and 75 percent a customize option. Banners participants meet do not differ across conditions: a one-click reject appears on 38 to 43 percent of the banners met in the three conditions ($p = 0.23$) and a customize option on 73 to 75 percent ($p = 0.63$, Appendix Table A1, Panel B). Figure 2 shows the decision on whether to engage or not with the banner, covariate by covariate, while Figures 5 and 6 show the decision on whether to accept or reject cookies, respectively, conditional on having engaged with the banner. Whenever we compare choices across websites or banners, we hold the other covariates fixed, with standard errors clustered by participant.

[11] A page is a pageview, with consecutive loads of the same URL within two minutes counted once, and a website is a registered domain.

3.1. Whether participants engage

On average, users engage with about 31 percent of banners, ignore 62 percent, and leave the website about 7 percent of the time. Whether a banner is engaged with is explained more by the characteristics of the website and of the banner than by those of the participant: 84 percent of the variation in engaging is the same participant behaving differently across banners (Appendix Figure A1). The characteristics of the environment (website, type of banner, etc.) explain 17 percent of this variation, and the banner's menu, whether the website forces a choice, the website's type and whether the participant has visited it before account for most of what they explain.

3.1.1. Characteristics of the user

Our results suggest that demographics are barely related to whether participants engage with banners. The only demographic trait which correlates with engagement at consent banners is age: older participants engage slightly more often, by about 1.4 percentage points for ten more years of age, while gender, ethnicity, education, income, employment and familiarity with technology make no significant difference. We find that the attitudes participants state at onboarding about privacy, tracking and advertising are not related to engagement, either. We observe that the engagement rate of more concerned participants is not significantly different from the one of less concerned ones. Similarly, beliefs about tracking and advertising do not correlate with different engagement rates (none of the three differences is significant, see Figure 2, panel A). [12] Only 16 percent of the variation in engaging is explained by differences between participants, where only about 0.5 percent can be attributed to demographics and attitudes towards privacy, tracking and advertising (see Appendix Figure A1). The remaining 84 percent is explained by differences between those banners the same participant meets.

[12] The four questions on privacy, advertising and tracking attitudes do not help predict whether a participant engages with a banner: the prediction is no better than tossing a coin (Appendix A.5).

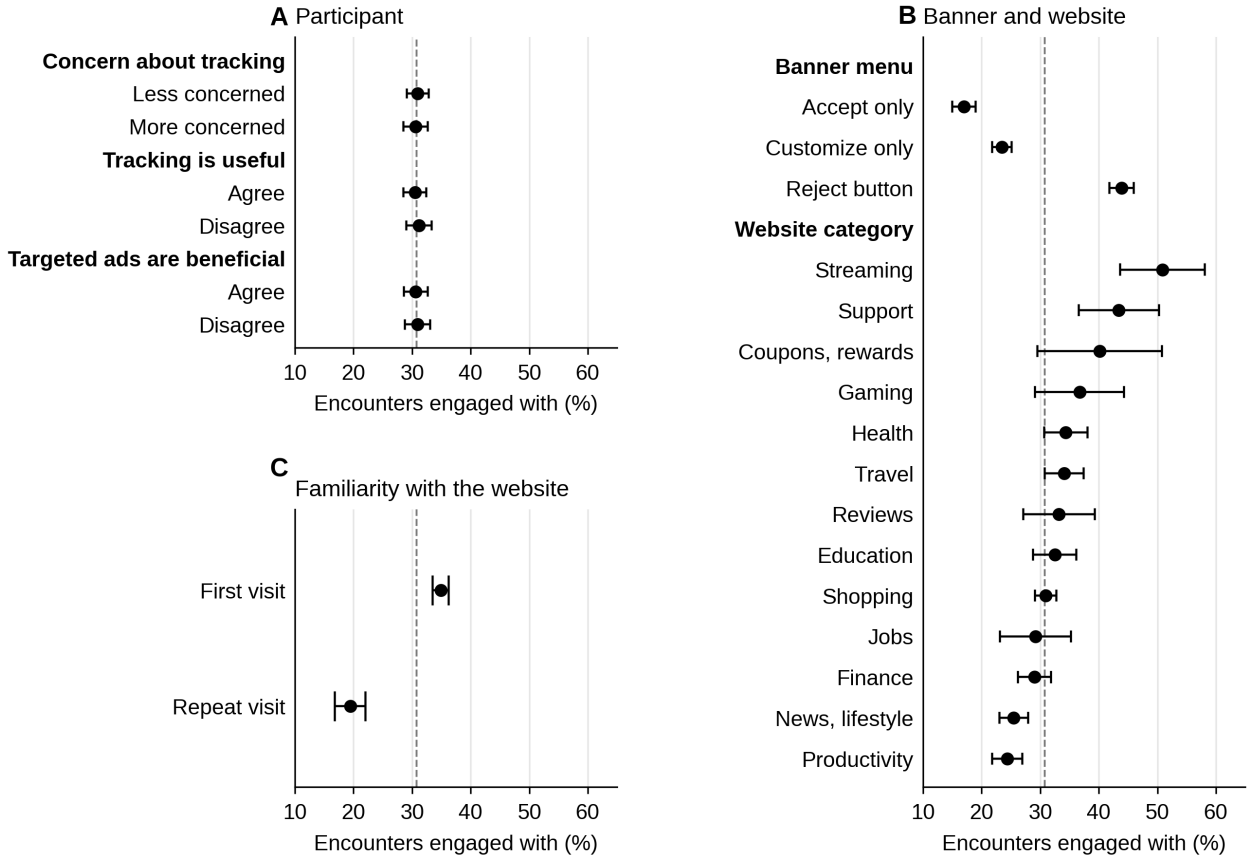

FIGURE 2. Share of encounters engaged with, by covariate. Construction in Appendix A.2.

3.1.2. Characteristics of the environment

Controlling for the menu of options offered, website popularity, the participant's familiarity with it, the experimental condition, the month and the participant's characteristics, we find that the share of banners engaged with varies significantly by website category. In particular, banner engagement is the lowest on productivity and on news and lifestyle websites (24 and 25 percent of encounters) and highest on streaming and on support and account websites (51 and 43 percent; Figure 2, panel B). We consider whether banners offer the 'Accept' or 'Customize' buttons only, and whether they offer a 'Reject' button. Users engage with banners 47 percent of the time when they offer a rejection option, against 20 percent where they do not (Figure 2, panel B),[13] and the gap remains when we compare the same participant on the same kind of website across templates (Section 3.2 shows what the

[13]Counting each participant and website once, with their first encounter, 37 percent of pairs are engaged with, against 31 percent of encounters, and the differences between banners are similar: 52 against 24 percent with and without a one-click reject (47 against 20 percent of encounters), and 80 against 32 percent where the website forces a choice and where it does not block (80 against 26).

additional choices are). Across all encounters, participants leave the website slightly less often when the banner offers a reject button (6 against 8 percent; Figure 3). The difference in engagement between banners with and without such a button is the same across the three experimental conditions, even though rejecting has no consequence in any of them. The customize option appears on three banners in four, and we record an acceptance or a rejection after 5 percent of the clicks on it, since the choices made inside the settings panel are not recorded. Figure 3 compares three features of the banner, with the website category and the month held fixed.

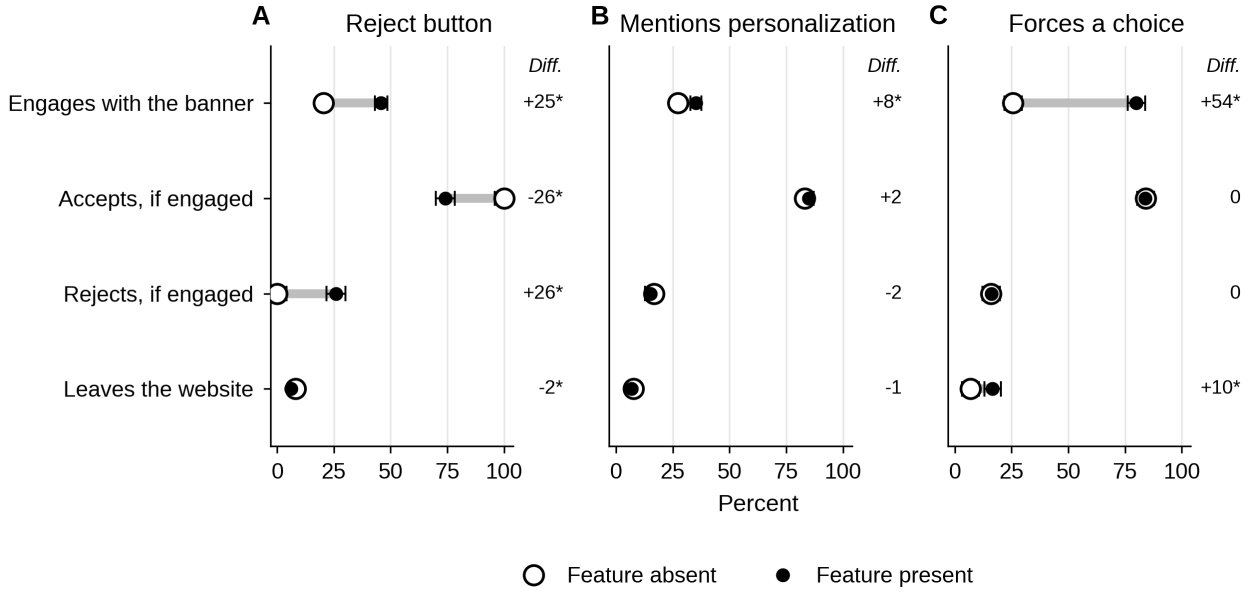

FIGURE 3. Engaging, accepting, rejecting and leaving, by feature of the banner. Construction in Appendix A.2. A website forces a choice when its page cannot be used until the participant makes a choice in the banner (Appendix A.3).

Banners whose text mentions personalization or improving the user's experience (43 percent of encounters) are engaged with more often, 35 against 27 percent with the category and the month held fixed. Once engaged with, they are accepted and rejected at the same rates as other banners (85 against 83 percent, and 15 against 17). A banner that forces a choice is engaged with in 80 percent of encounters, against 26 percent where it can be ignored, and once engaged with it is accepted at the same rate as any other, 84 percent in both cases. Among the three features, only the reject button changes what participants choose. The last row of Figure 3 shows that participants are 2 percentage points more likely to leave the website if banners do not offer a reject button, and more than twice as likely if banners force a choice, while not significantly more likely if banners

mention personalization.

Experimental conditions do not influence engagement with consent banners. Engagement is 29.6 percent of encounters under Opt-In, 31.9 percent under Anti-tracking and 31.4 percent under Ad-blocking ($p = 0.77$). With the website category and the month held fixed, engagement is 1.2 percentage points higher under Anti-tracking than under Opt-In, with a 95 percent interval from -3.3 to 5.6, and 1.4 points higher under Ad-blocking, from -2.6 to 5.4.

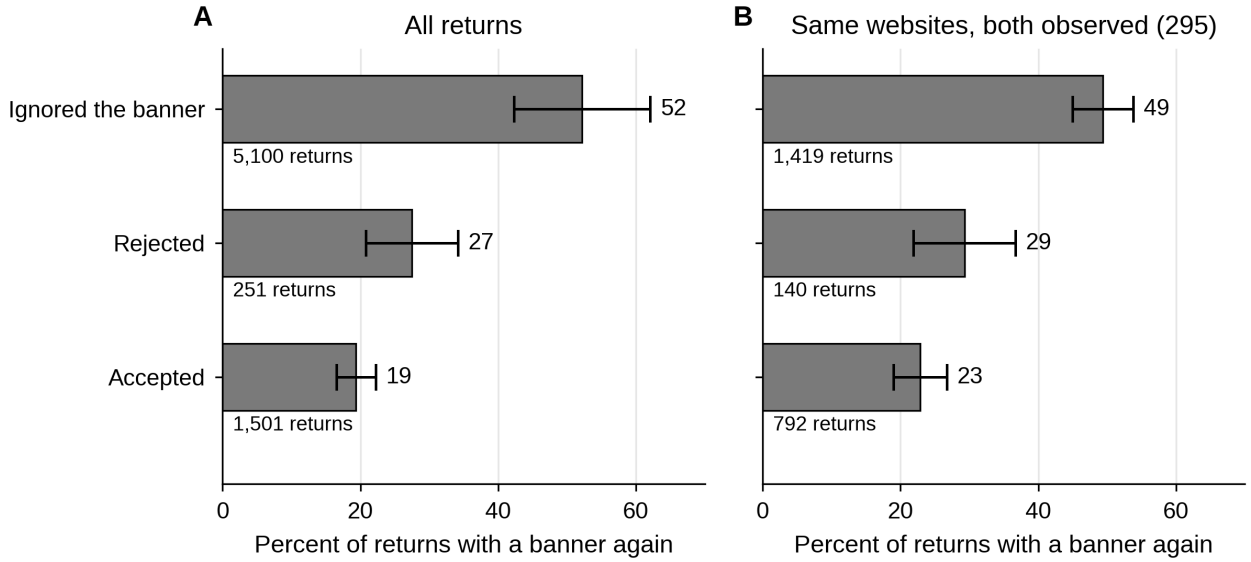

FIGURE 4. Whether a banner is recorded again on the next visit to the same website. Construction in Appendix A.2.

Finally, whether a banner is shown again when a participant returns to a website depends on what they did at the previous encounter on that website. We observe that a banner is shown again on 52 percent of the returns to websites where participants previously ignored the banner, against 27 percent where they rejected and 19 percent where they accepted. Figure 4, panel A, shows these shares. Panel B keeps only the 295 websites where some returns follow an ignored banner and others follow an explicit choice: there a banner is shown again on 49 percent of returns after an ignored banner, 29 percent after a rejection and 23 percent after an acceptance, a difference of 26 percentage points within website ($p < 0.001$), the same in each of the three experimental conditions. The banners a participant meets on repeat visits are therefore mostly banners they have

ignored before. The difference between a rejection and an acceptance is not significant within website.

3.1.3. Previous visits to the website

We then study whether having visited a website before influences the probability of engaging with its consent banner. We find that participants are more likely to engage with banners on websites they visit for the first time than with banners on websites they return to: 35 percent of the first encounters with a website within the panel are engaged with, against 19 percent of repeat encounters (Figure 2, panel C). Of the variation in engaging within participants that the covariates explain, the menu accounts for two fifths, whether the website blocks for a third, and the website category and familiarity for about a twelfth each (Appendix Figure A1).

3.2. What people choose

Among those banners participants engage with, 84 percent accept and 16 percent reject. The share of acceptances is close to the 78 percent that the Behavioural Insights Team (2024) finds considering only explicit choices under a neutral website-level banner in a simulated browser experiment. However, in our data, it describes only the third of banners that participants actually engage with. While the decision to engage with a banner is not predicted by participant's characteristics, the *choice* made at cookie banners is explained by those characteristics: 56 percent of the variation in accepting among choices lies between participants (Appendix Figure A1), and their stated attitudes account for most of the part that the characteristics we observe explain, which is 10 percent of it.

3.2.1. Characteristics of the user

Participants who agree that tracking is useful accept 88 percent of the banners they engage with, and those who disagree 79 percent ($p = 0.003$); the more concerned accept 82 percent of theirs, and the less concerned 88 percent ($p = 0.03$). The belief that targeted advertising benefits consumers, instead, is not associated with significantly different acceptance rates (Figure 5, panel A).[14] Since stated attitudes are not related to whether a participant engages

[14]The attitudes predict whether a participant accepts better than tossing a coin, but far from accurately (Appendix A.5). Of the variation in accepting between participants that the covariates explain, the belief that tracking is useful accounts for 35 percent, concern for 34 percent, the belief about advertising for 12 percent and demographics for 14 percent (Appendix Figure A1).

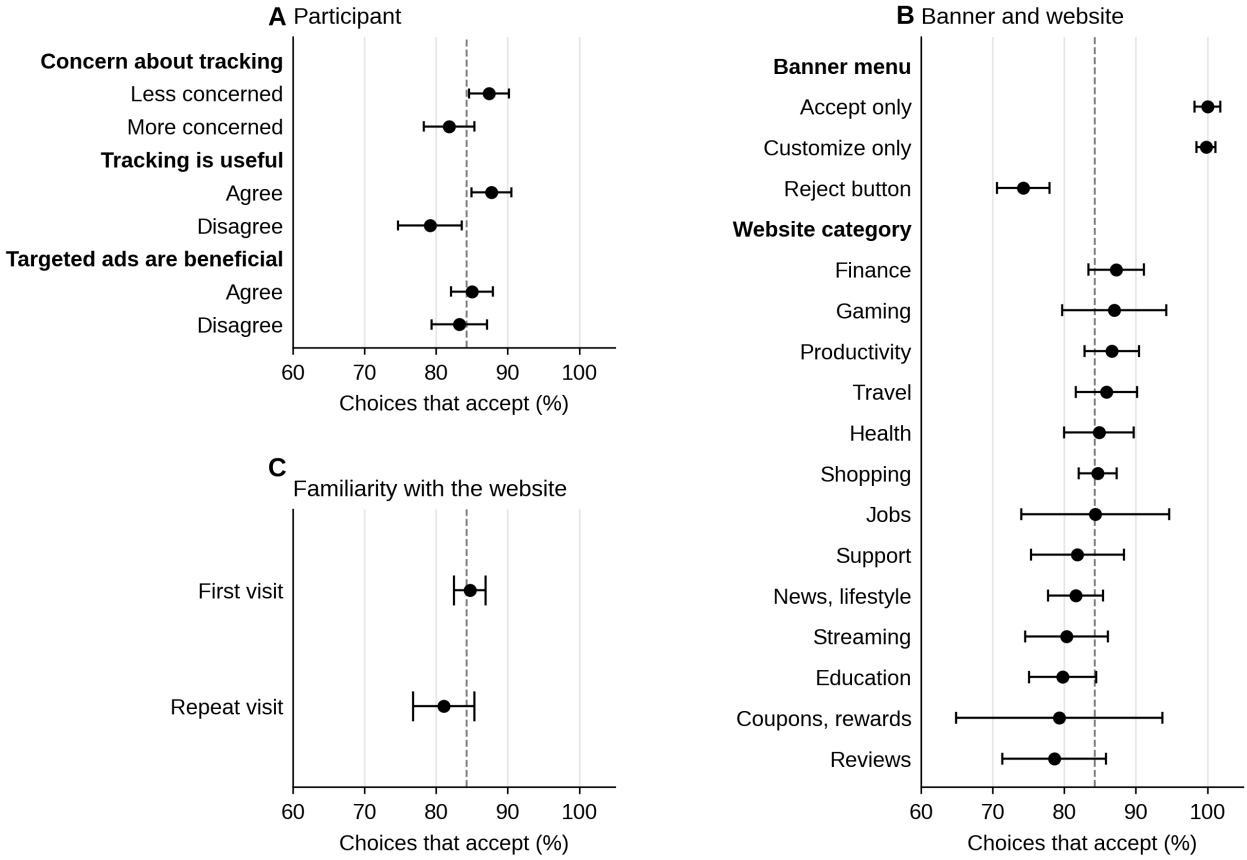

FIGURE 5. Share of choices that accept, by covariate. Construction in Appendix A.2.

with a banner, a model that pools together the decision of engaging with the nature of the choice would find a weak link between stated attitudes and banner clicks, as two thirds of the banners are ignored.

3.2.2. Characteristics of the environment

Among the characteristics of the encounter, the banner's menu of options is the only one significantly associated with what participants choose once they engage. When a one-click reject button is available, participants accept tracking in 74 percent of their explicit choices and reject it in the remaining 26 percent. In contrast, when no reject button is offered, every explicit choice we record is an acceptance (Figures 5 and 6, panel B).[15] However, this comparison tells only part of the story. Offering a reject button is also

[15]On these banners, refusing requires opening the settings. We record the click that opens them (271 clicks on banners without a one-click reject) but not the choices made inside the settings panel, so these encounters count as ignored. If all of them ended with a refusal, rejections would be at most about 12 percent of explicit choices on these banners.

associated with substantially higher engagement: participants make an explicit choice in 47 percent of encounters when the button is available, compared with only 20 percent when it is not. As a result, despite the lower acceptance rate conditional on engagement, the share of all encounters that result in acceptance is actually higher when a reject button is offered: 35 percent compared with 20 percent ($p < 0.001$). The increase in engagement more than offsets the greater likelihood of rejection among those who make a choice. These findings illustrate why acceptance rates cannot be interpreted independently of the interface through which consent is requested. The menu influences not only whether participants accept or reject tracking, but also whether they make an explicit choice at all. Consistent with this, within participants, the menu accounts for more than four fifths of the variation in acceptance explained by the observed covariates (Appendix Figure A1).

The absence of a reject button is also associated with how participants respond to consent banners, even among those who generally prefer to accept tracking. To examine this, we classify participants based on their choices on banners offering a one-click reject option. Among participants who make at least two such choices, 113 reject more often than they accept (we call this group "Rejectors"), while 350 accept more often than they reject ("Acceptors"). Both groups are more likely to ignore banners when a reject button is unavailable, despite their different preferences. Among Rejectors, engagement falls from 49 percent of encounters when a reject button is offered to just 12 percent when it is not ($p < 0.001$). Correspondingly, the share of encounters in which they ignore the banner and continue browsing rises from 44 to 78 percent ($p < 0.001$). Importantly, Acceptors also ignore banners more frequently when a reject button is absent, although the increase is 11 percentage points smaller than among Rejectors (Appendix Figure A2). The absence of a one-click reject option also affects what Rejectors choose. They accept tracking more often (12 versus 7 percent of encounters, $p < 0.001$) and are slightly more likely to leave the website (10 versus 7 percent, $p = 0.01$), although their likelihood of leaving does not differ significantly from that of other participants ($p = 0.28$). These patterns suggest that the association between the reject button and engagement cannot be explained solely by participants seeking to refuse tracking. Even those who generally accept tracking engage less often when the option to reject is unavailable.

As observed for engagement, different website categories show different acceptance and rejection rates among explicit choices. Across the thirteen categories of Figure 2, the share of choices that are acceptances ranges from 79 to 80 percent on coupons and rewards, reviews and education websites to 88 percent on finance websites, and from 80 to 88 percent among the categories with at least 200 choices, while the percentage of rejections spans from 22 percent for coupons and rewards websites and about 20 percent for reviews, education and streaming websites, to 12 percent for finance websites and 13 percent for gaming websites (Figure 6, panel B). We find similar patterns for familiarity: 85 percent of the choices given on a first visit to a website within the panel are acceptances and 81 percent of those given on a repeat visit, a difference of 4 points, against 35 and 19 percent of encounters engaged with (Figure 5, panel C). Together, category and familiarity account for less than a tenth of the variation in accepting within participants that the covariates explain (Appendix Figure A1).[16]

[16]A reject brings the banner back more often than an accept (27 against 19 percent of the returns to the same website, a difference that is not significant within website), and participants who rejected mostly reject again. Keeping only the first choice of each participant on each website, so that no website enters twice, acceptances are 84.5 percent of first-visit choices and 82.8 percent of repeat-visit choices ($p = 0.57$).

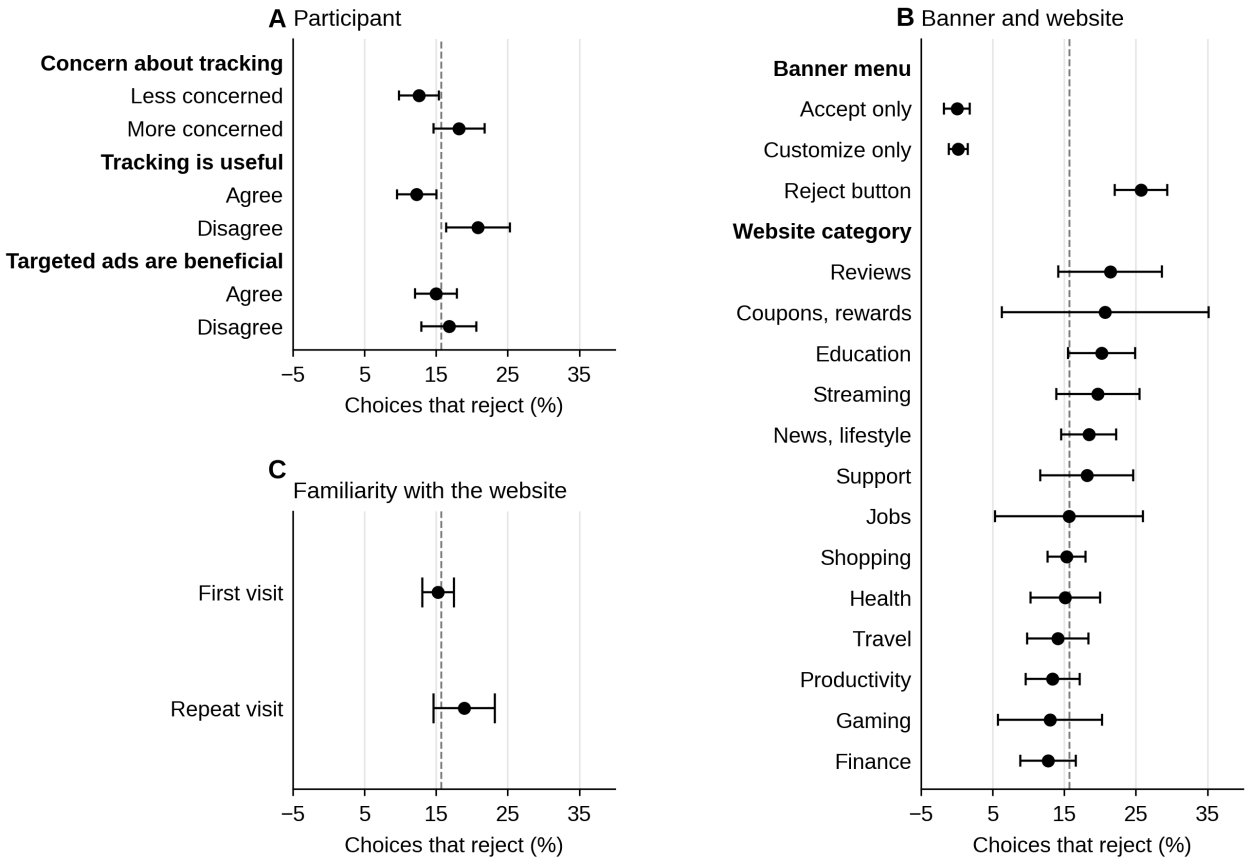

FIGURE 6. Share of choices that reject, by covariate. Construction in Appendix A.2.

As noted, whether a website forces users to make a choice strongly influences engagement with consent banners. We now examine how this feature interacts with the menu of options offered by the banner. Figure 7 compares participants' responses across three types of websites: those that do not block access to the page, those that block access but allow users to dismiss the banner with an 'X' button, and those that require an explicit choice before the page can be accessed. When the banner can be ignored, participants mostly continue browsing without making a choice. They do so in 70 percent of encounters without a reject button, compared with 52 percent when one is available. When a reject button is offered, both acceptance and rejection are more frequent: acceptance increases from 22 to 32 percent of encounters, and rejection from zero to 11 percent. The pattern differs when websites force a choice. Without a reject button, participants accept tracking in 83 percent of encounters, while the remaining 16 percent leave the website. When a reject button is offered, acceptance falls to 60 percent and rejection rises to 25 percent, but the share of participants leaving remains unchanged at 16 percent. Thus, when websites

force a choice, the availability of a reject button is associated with a different distribution of consent choices, rather than a difference in the likelihood of leaving the website.

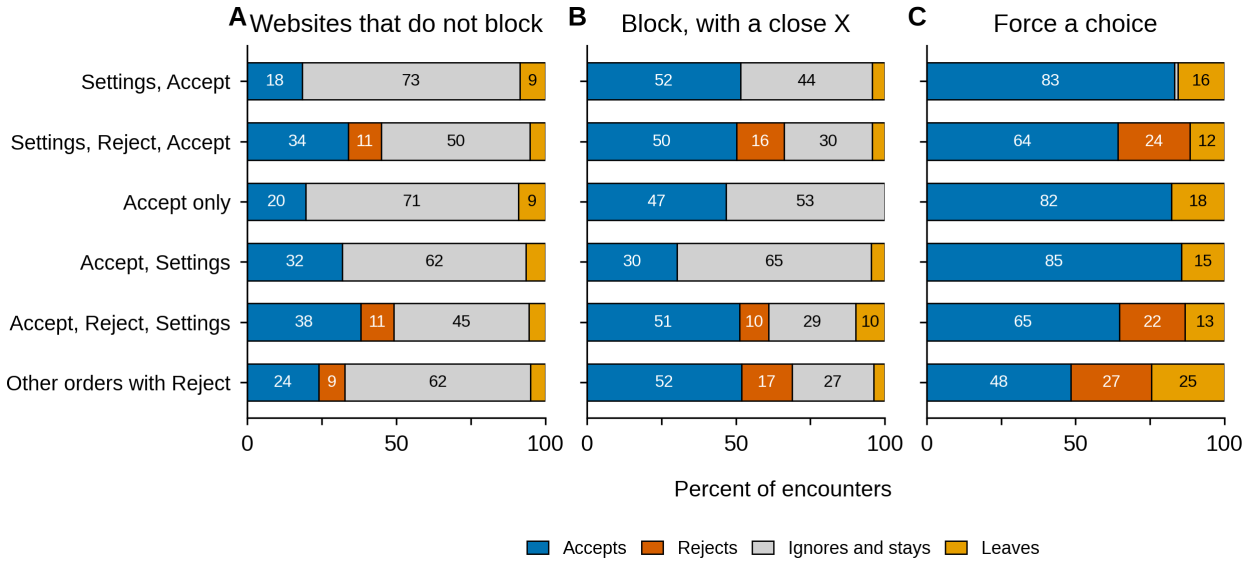

FIGURE 7. What participants do at each menu, on websites that do not block, that block with a close X, and that force a choice. Construction in Appendix A.2.

These results shed light on an important fact: where websites force users to make a choice, engagement is much higher, even in cases where the user would likely have ignored the consent banner. Since websites choose whether to force a choice, this comparison is descriptive. We discuss how this can impact welfare calculations in Section 4.2.1.

3.2.3. Previous visits to the website

We then study whether having visited a website before influences what participants choose once they engage with its banner. As noted above, the share of acceptances at first and later encounters with a website (85 and 81 percent) differs within sampling error, and later encounters account for only 643 of the explicit choices, since a banner is rarely shown again after a choice (Figure 5, panel C, and Figure 4).

3.2.4. Impact of experimental treatments

Finally, we study whether the condition participants are assigned to is related to their choices at banners. Two expectations are plausible. On the one hand, the condition changes what participants see while browsing, and this could change how they answer

banners. For instance, a concerned participant in the Opt-In condition sees more targeted ads and might reject more often, while a participant in the Ad-blocking condition sees little to no advertising and they may have less reason to reject. On the other hand, a banner asks about one website at the moment of arrival, before its ads are shown, so choices may follow the participant's attitudes and the banner's menu rather than what they experienced on other websites. We find that choices do not differ across conditions. Acceptance among choices is 86 percent under Opt-In, 84 percent under Anti-tracking and 82 percent under Ad-blocking, differences within sampling error, and the increase in rejecting for one standard deviation of stated concern is 5.8 percentage points under Opt-In, 2.5 under Anti-tracking and 3.8 under Ad-blocking, three estimates that cannot be distinguished from one another (joint $p = 0.41$).[17]

4. Generalizing Consent Choices Across Websites

4.1. Within-participant variation in consent choices

We use the study data to investigate within-participant variation in consent choices across time and websites, and then project the possible impact of propagating a single consent choice signal to all websites a person visits. Figure 8, panel A, shows that, on average, participants engage with only one in three banners: in 27 percent of later encounters they make the same choice as their signal, and in 4 percent the opposite one, while they do not engage with the remaining 69 percent (ignoring the banner or leaving the website). In Panel B, we restrict our attention to those 520 participants who made at least one explicit choice (i.e., accept or reject) in the first third of their days in the study, and we distinguish participants based on the explicit choices made during the first third of their days in the study. We define the "own signal" as the choice a participant makes most often at banners during the first third of the days we observe them. Participants' own signal is consistent with 90 percent of their explicit choices (which account for about 30 percent of their actions). A signal that accepts for every participant is consistent with 84 percent, and one that rejects for every participant with 27 percent. The participants who only accepted in the first third of their days (392) keep accepting: their signal is consistent with 94 percent of their later choices, the same as accepting for everyone.

[17] In a linear model with the full set of controls, clustered by participant, the two interactions with the condition are -3.3 (2.5) and -2.0 (2.9) percentage points.

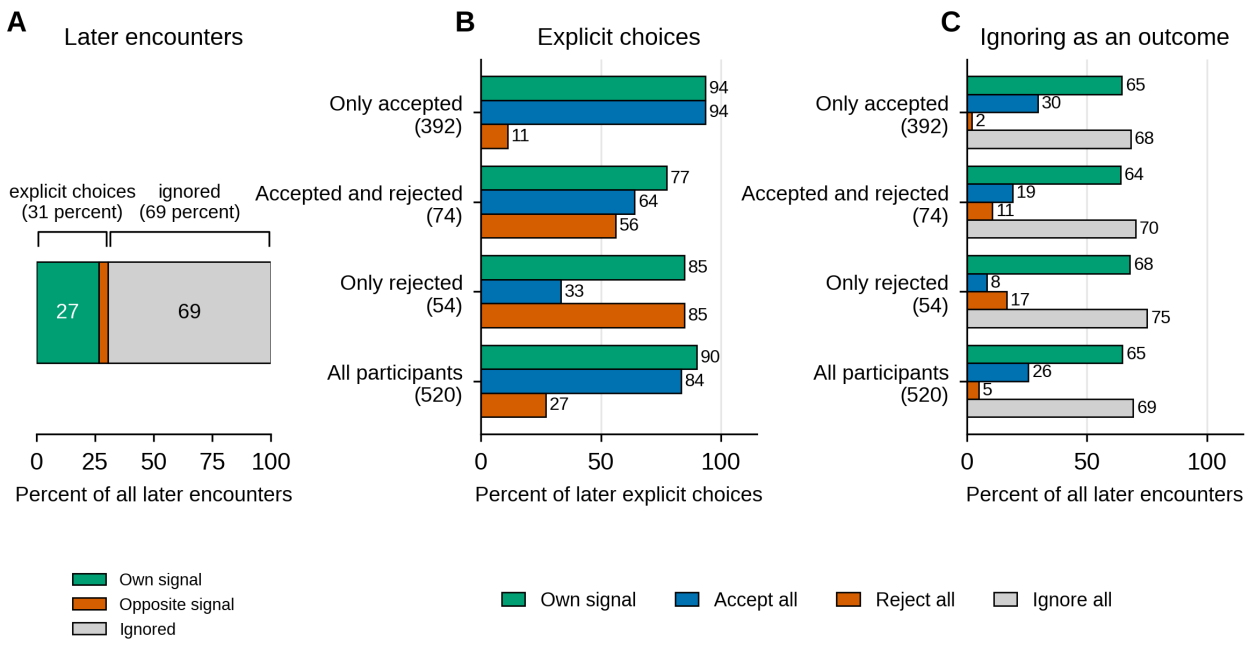

FIGURE 8. Later encounters of each participant and consistency of a signal set in the first third of the days we observe them, on explicit choices and with ignoring counted as an outcome. Construction in Appendix A.2.

The signal is consistent with 85 percent of the later choices of the participants who only rejected (54), and with 77 percent for those who both accepted and rejected (74), against 64 percent for a signal that always accepts and 56 percent for one that always rejects. The shares of the two signals can add up to more than 100 because they are computed on different choices: under a $Reject$ signal, acceptances on banners without a one-click reject are left out, since the participant could not reject in one click (Appendix A.2). Panel C counts ignoring as a third possible reaction, both in the signal and in the later encounters. For most participants (63 percent), *ignoring* is the most frequent action in the first third of their days (namely, their "own signal"), and it is consistent with 74 percent of their later encounters. For those who mainly accepted (33 percent), their own signal is consistent with 43 percent of their later encounters, and for those who mainly rejected (4 percent) with 18 percent. Over all participants, the own signal is consistent with 65 percent of later encounters, against 69 percent if every participant were assigned to ignore, 26 percent if assigned to accept and 5 percent if assigned to reject.

Finally, we analyze the relationship between the privacy concern participants stated at onboarding and what they do at banners later in the study. We use the three concern

groups of Section 2 and the encounters of the 520 participants after the first third of their days. The three groups behave in a similar way (Figure 9, panel A): whatever their concern, participants ignore about 70 percent of later banners and accept about a quarter of them. Participants with high concern reject 7 percent of their later encounters, against 3 percent for those with low concern, and among the banners they engage with they accept 77 percent, against 82 percent for those with medium concern and 91 percent for those with low concern.

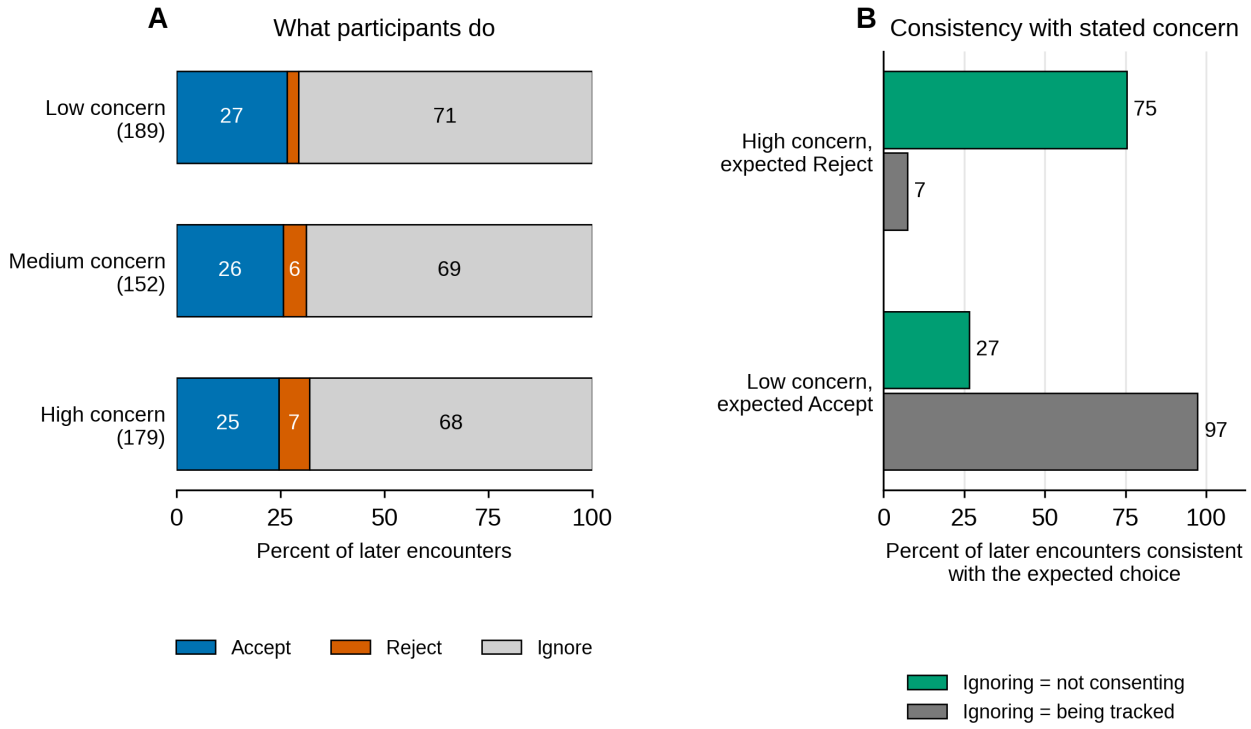

FIGURE 9. Later encounters by stated privacy concern, and consistency of stated concern with later encounters under two readings of ignoring the banner. Construction in Appendix A.2.

Whether this behavior is consistent with stated concern depends on how participants interpret ignoring a banner. We take *Reject* as the choice closest to high concern and *Accept* as the choice closest to low concern. If participants who ignore a banner believe that they are not consenting to tracking, the behavior of participants with high concern is consistent with their stated concern in 75 percent of their later encounters, against 27 percent for those with low concern. If instead they believe that ignoring a banner means being tracked, the shares are 7 percent for participants with high concern and 97 percent for those with low concern (Figure 9, panel B). In other words, most of the consistency comes from the banners participants ignore, and which group appears consistent depends

on how ignoring is interpreted. On the banners they engage with, participants with high concern accept three in four.

As discussed in Section 3, explicit choices are about a third of the encounters, while the other two thirds are banners the participant ignored. When a single signal is chosen and propagated, it must also affect those encounters where the user ignored the banner. We can think of three possible explanations: either the user ignored the banner thinking that, by doing so, their data would not be tracked (as if they chose $Reject$), or that their data would be tracked (as if they chose $Accept$), or they had no preference over their data being tracked or not. The consistency of the signal depends on which reading holds: if ignoring means the same choice as the signal, the signal set in the first third of their days is consistent with about 96 percent of later encounters; if it means the opposite choice, with 27 percent. If the user has no preference, the overrides are confined to the encounters the user engaged with. When a banner offers no one-click reject, participants who usually reject ignore it on 78 percent of encounters, against 44 percent when they can reject (Section 3.2), so at least part of ignoring reflects a refusal that the banner does not allow in one click.

4.2. Implications of centralized consent mechanisms

We now discuss possible economic implications of propagating a single signal for the economic actors involved. We use our results to discuss scenarios in which, rather than eliciting consent on every website encountered, a user can make an initial choice (accept or reject third-party cookies), and that signal will be read by all websites they will later visit.[18]

4.2.1. Users

Under one possible implementation of centralized consent, users would not be prompted to make an explicit choice on every website they visit, as their browser-level choice would be propagated to all websites they visit. This implies that users no longer spend a

[18]This is one reading of a machine-readable signal described in the proposed Article 88b of the GDPR, introduced by the Digital Omnibus proposal from the European Commission (November 2025). This is also the simplest benchmark for a question the data can answer, namely how well one persistent choice would reproduce the decisions the same person otherwise makes website by website.

certain amount of time making choices in consent banners. Under this implementation, centralizing consent may have other implications for users beyond the time they spend on banners, such as a loss of personalization, or a mismatch between a single choice and the contextual choices users would make on different websites (Section 4.1; different implementations of centralized consent mechanisms could allow users to both set default choices for most sites and adapt them for specific sites). While we refrain from making conclusions about what user behavior will be under a different choice architecture (one where users can set a consent signal on their browsers, and websites respect them), we can estimate how much time participants in the study spent making banner choices, and use such calculations to inform our discussion regarding potential time savings of centralized consent mechanisms. The time users spend interacting with banners depends on various factors, including the number of banner prompts received, the actual response rate to banners, the amount of time needed to make a choice, and the availability of the option to ignore the banner. We measure the number of new websites participants visit, the share of banners they engage with and the time a choice takes, and we ask how the time a user spends on banners would change under different shares of websites that show a banner and that force a choice. In the *status quo*, users are prompted with consent banners every time they visit a new website that uses consent management platforms.

The number of banners a person meets depends on how many websites they visit for the first time, since websites mostly do not show the banner again once it has been engaged with.[19] Our median participant visited 76 new websites in a month. Consider a user who visits as many new websites, and assume that two in three of them show a consent banner, so that the user encounters about 51 banners in a month.[20] As shown in Figure 3, the probability of engaging with a banner on a website that does not force a response is around 26 percent, whereas on websites that force a response a user who stays

[19]A banner that is ignored is more often shown again: in our dataset, on the websites where some returns follow an ignored banner and others follow an explicit choice, a banner was recorded on 49 percent of returns after an ignored banner, against 24 percent after an explicit choice. Participants rarely engaged with it later: of the participant and website pairs whose first banner was ignored, 16 percent met a banner again on that website on a later day and 2.2 percent later made an explicit choice there. Counting only new websites therefore understates the banners a user meets, and only slightly the banners they engage with.

[20]Among the 10,000 most popular websites in each European country, 67 percent show a consent banner (Nouwens et al., 2025). In our panel, the median participant meets a banner on 5 of every 100 websites visited, a lower bound since the extension recognizes eight consent management platforms. The minutes per year below therefore describe a user exposed to banners as often as on European websites, while the percentage reduction does not depend on this share.

must engage with the banner. Assuming half of websites would force a response, we could expect that the number of consent banners people engage with in a month is around 32, 26 of them on websites that force a response.[21] If we focus on the time spent interacting with consent mechanisms (and ignore other associated costs, such as consent fatigue or frustration with repeated consent requests), we estimate the median choice to take 6.6 seconds from the moment the page has finished loading, 6.6 seconds for an acceptance and 5.3 for a rejection.[22] Therefore, this user would spend about 212 seconds per month, which is about 42 minutes per year. Assuming that a website that stops forcing a response would see the engagement of the websites that do not force one today, the same user would engage with about 13 banners a month and spend about 18 minutes per year, 58 percent less. The reduction depends on how many websites force a response: it would be about 41 percent if a quarter of the websites with a banner forced a response, and about 22 percent if one in ten did.

4.2.2. Firms

Whether ad-funded websites lose or gain revenue depends on multiple factors, including: the precise way in which generalized consent mechanisms would be implemented; whether a substantial fraction of users would move from being tracked to not being tracked; and whether and how the online advertising ecosystem would adapt to a new equilibrium in which, for instance, a larger proportion of users opting out of tracking brings more prominence and funding to contextual advertising and first-party publishers' data. From previous studies, we know that an ad shown without tracking is estimated to be worth less, from 4 percent less in Marotta et al. (2019) to about half in Ravichandran and Korula (2019) and Johnson et al. (2020). Using data from Yahoo's ad auctions, Alcobendas et al. (2023) simulate that removing third-party cookies from Chrome would lower publishers' revenue by 30 percent. These estimates, however, come from settings in which tracking was reduced for some users, platforms, or websites, while advertisers could still track users elsewhere; hence, part of the loss those studies measure may be local effects (spending that moves to the users and websites that remain trackable), rather than general economic

[21] In the EU, 46.5 percent of consent banners are found to block the website until they are engaged with (Bouhoula et al., 2024), close to our assumption of 50 percent.

[22] Timed on the 1,527 choices for which the extension records the load-complete event; from the page commit, available for all 4,386 timed choices, the medians are 7.7 seconds overall, 7.7 for acceptances and 7.3 for rejections.

effects of broad curtailing of online tracking and ad targeting (Acquisti, 2024). In our data, 95 percent of encounters end without a rejection (with an acceptance, an ignored banner or the participant leaving the website), which counts as consent under U.S. practice, so most of the impressions that a change in how consent is collected would affect are currently tracked. Consent management platforms would also be affected. Websites obtain their banners from these platforms, which offer the templates and the menus that, as we find in Section 3.2, influence whether users engage with a banner and what they choose. In our data, eight platforms supply all the banners we record, and one of them, OneTrust, supplies 85 percent of encounters. If users expressed their consent once in the browser, the role of these platforms would move from designing the banner that asks for consent to reading and applying a choice that users make elsewhere, and our data do not tell us how the platforms and the menus they offer would change. Whether large platforms or publishers would ultimately gain depends on how generalized consent mechanisms will be designed and implemented. If generalized consent mechanisms only applied to websites, and not to data platforms, large platforms that collect first-party data on their own services may gain relative to publishers: a choice expressed in the browser applies to the tracking that passes through the browser, such as third-party cookies, while the data a platform collects from its logged-in users, or exchanges between servers, would largely fall outside it (Assembly Research, 2026). A rule that only reduced third-party tracking could therefore lower the value of the impressions sold by publishers more than that of the impressions sold by platforms. In a case brought by Spanish media companies, a Madrid court found in 2025 that Meta had processed the personal data of its users for behavioral advertising without a valid legal basis, which gave it an advantage over online media in the advertising market, and ordered it to pay 479 million euros to the publishers (Assembly Research, 2026).

5. Conclusion

We study the behavior of 789 U.S. Chrome users for several weeks and record the consent banners they encounter during their organic browsing. We find that participants engage with only one in three consent banners, and that whether a banner is engaged with is explained by the characteristics of the environment rather than by those of the user: the type of website, the menu offered by the banner, the familiarity the user has with the website, and whether the website can be browsed before a choice is made. Where the

banner forces a response, engagement is three times higher than where it can be ignored. What is chosen is instead explained by characteristics of the person, such as privacy preferences and demographics, with one exception: when the banner offers a one-click reject option, participants engage more often and, among their choices, reject more often. Finally, the same person often makes different choices on different websites: more than a third of the participants who engage with at least five banners both accept and reject, and a signal set from the first third of the days we observe each participant is consistent with 90 percent of their later explicit choices, and with 65 percent of all their later encounters when ignoring the banner is counted as an outcome the signal can match.

These findings have implications for proposals that move consent from the banner to a setting in the browser. First, since most banners are ignored, a central signal would also apply to the many encounters in which users currently make no choice. We do not know what ignoring a banner means for the user, who may be content to be tracked, may think that not clicking is enough to refuse, or may not care, and the consistency of any signal on those encounters depends on it. Second, a single signal set once would override some of the explicit choices participants make website by website. For instance, if we assume that participants with high stated concern would choose $Reject$, they are a third of participants, and they accept 77 percent of the banners they engage with, while they ignore two thirds of the banners they meet. Third, most of the time users currently spend on banners is spent on websites that prevent browsing unless a choice is made. Under the assumptions of Section 4.2.1, preventing websites from forcing a response would reduce that time by more than half, without propagating a signal the user may not want.

These implications have limits. We observe choices made at banners designed by websites, under current rules, by U.S. users who can ignore most of them. The consent banners we analyze are not randomized within websites or categories of websites, so these comparisons are descriptive rather than causal. How the same users would answer a single question asked by the browser depends on how that question is designed: in a simulated browser, preselecting the rejection of cookies raises the share of participants who decline from 22 to 39 percent (Behavioural Insights Team, 2024), and privacy choices are known to depend on the context in which they are made (Acquisti et al., 2015). For this reason, we do not use our results to predict the share of users who would send a $Reject$

signal. The same caution applies to firms. How advertising revenue and the relationship between publishers, intermediaries and users would change depends on how many users refuse and on how websites and platforms respond. A central signal could strengthen the position of the large platforms that already hold first-party data, or it could give content providers back control over their relationship with their users, and we do not observe either response. The single signal we study is the simplest version of a choice made once. A browser setting may also let users make exceptions for some websites or revise their choice, which most participants in the survey of the Behavioural Insights Team (2024) say they would want, and the Global Privacy Control used in California only signals an opt-out of the sale or sharing of personal data; our data do not tell us how users would use these options.

References

Alessandro Acquisti. The economics of privacy at a crossroads. In Avi Goldfarb and Catherine Tucker, editors, *The Economics of Privacy*. University of Chicago Press, Chicago, 2024. NBER conference volume.

Alessandro Acquisti, Laura Brandimarte, and George Loewenstein. Privacy and human behavior in the age of information. *Science*, 347(6221):509–514, 2015. doi: 10.1126/science.aaa1465.

Alessandro Acquisti, Curtis Taylor, and Liad Wagman. The economics of privacy. *Journal of Economic Literature*, 54(2):442–492, 2016.

Miguel Alcobendas, Shunto Kobayashi, Ke Shi, and Matthew Shum. The impact of privacy protection on online advertising markets. In *Proceedings of the 24th ACM Conference on Economics and Computation*, EC '23, page 62, New York, NY, USA, 2023. Association for Computing Machinery. ISBN 9798400701047. doi: 10.1145/3580507.3597712. URL https://doi.org/10.1145/3580507.3597712.

Guy Aridor, Yeon-Koo Che, and Tobias Salz. The effect of privacy regulation on the data industry: Empirical evidence from GDPR. *The RAND Journal of Economics*, 54(4):695–730, 2023. doi: 10.1111/1756-2171.12455.

Article 29 Data Protection Working Party. Opinion 04/2012 on cookie consent exemption. Technical Report WP194, European Commission, 2012.

Assembly Research. The economic costs of weakening privacy and data protection. Technical Report BEUC-X-2026-086, Report commissioned by BEUC, The European Consumer Organisation, September 2026. URL https://www.beuc.eu/sites/default/files/publications/BEUC-X-2026-086_The_economic_costs_of_weakening_privacy_and_data_protection.pdf.

Behavioural Insights Team. Evaluating browser-based cookie setting options to help the UK public optimise online privacy behaviours: Results report. Technical report, Department for Science, Innovation and Technology, 2024. URL https://www.bi.team/wp-content/uploads/2024/05/DSIT-Browser-based-cookie-systems-final-report.pdf.

Nataliia Bielova. Survey of academic studies measuring the effect of dark patterns on acceptance consent rate of users in consent banners. Technical report, Commission Nationale de l'Informatique et des Libertés (CNIL), LINC, Paris, December 2022. URL https://www.cnil.fr/sites/default/files/atoms/files/full_2022-12-02_v2.pdf.

Nataliia Bielova, Laura Litvine, Anysia Nguyen, Mariam Chammat, Vincent Toubiana, and Estelle Hary. The effect of design patterns on (present and future) cookie consent decisions. In *33rd USENIX Security Symposium (USENIX Security 24)*, pages 2813–2830. USENIX Association, 2024.

Ahmed Bouhoula, Karel Kubicek, Amit Zac, Carlos Cotrini, and David Basin. Automated large-scale analysis of cookie notice compliance. In *33rd USENIX Security Symposium*. USENIX Association, 2024.

Anu Bradford. *The Brussels Effect: How the European Union Rules the World*. Oxford University Press, New York, 2020.

Cristobal Cheyre, Li Jiang, Florian Schaub, Zijun Ding, Cristiana Firullo, Yucheng Li, and Alessandro Acquisti. An experimental infrastructure to investigate the impact of online tracking, targeting, and advertising on consumer behavior and consumer welfare. Working paper, 2024.

Martin Degeling, Christine Utz, Christopher Lentzsch, Henry Hosseini, Florian Schaub, and Thorsten Holz. We value your privacy ... now take some cookies: Measuring the GDPR's impact on web privacy. In *Network and Distributed System Security Symposium (NDSS)*, 2019. doi: 10.14722/ndss.2019.23378.

Alexander Egberts and Julia Krämer. Preference revelation under design frictions: Evidence from cookie consent banners. Technical report, SSRN, May 2026. Working paper, version posted 20 May 2026.

European Commission. Commission staff working document accompanying the digital omnibus proposals, 2025. SWD(2025) 836 final, Brussels, 19 November 2025.

European Data Protection Board. Guidelines 05/2020 on consent under Regulation 2016/679. Technical report, European Data Protection Board, 2020. Version 1.1.

European Data Protection Board. Guidelines 2/2023 on technical scope of art. 5(3) of ePrivacy directive. Technical report, European Data Protection Board, 2024. Version 2.0.

Chiara Farronato, Andrey Fradkin, and Tesary Lin. Designing consent: Choice architecture and consumer welfare in data sharing. *Management Science*, 2026. Forthcoming. Earlier version: NBER Working Paper 34025 (2025).

Avi Goldfarb and Catherine E. Tucker. Privacy regulation and online advertising. *Management Science*, 57(1):57–71, 2011.

Thomas Groß. Validity and reliability of the scale internet users' information privacy concerns (IUIPC). *Proceedings on Privacy Enhancing Technologies*, 2021(2):235–258, 2021.

Zhengrong Gu, Garrett A. Johnson, and Shunto J. Kobayashi. Can privacy technologies replace cookies? Ad revenue in a field experiment. *Proceedings of the National Academy of Sciences*, 123(19), 2026. doi: 10.1073/pnas.2603752123.

Hana Habib, Megan Li, Ellie Young, and Lorrie Faith Cranor. "okay, whatever": An evaluation of cookie consent interfaces. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems (CHI)*, 2022. doi: 10.1145/3491102.3501985.

Maximilian Hils, Daniel W. Woods, and Rainer Böhme. Measuring the emergence of consent management on the web. In *Proceedings of the ACM Internet Measurement Conference (IMC)*, pages 317–332, 2020. doi: 10.1145/3419394.3423647.

Nikhil Jha, Martino Trevisan, Marco Mellia, Daniel Fernandez, and Rodrigo Irarrazaval. Privacy policies and consent management platforms: Growth and users' interactions over time. *ACM Transactions on the Web*, 19(3):1–25, 2025. doi: 10.1145/3725737.

Garrett A. Johnson, Scott K. Shriver, and Shaoyin Du. Consumer privacy choice in online advertising: Who opts out and at what cost to industry? *Marketing Science*, 39(1):33–51, 2020.

Naresh K. Malhotra, Sung S. Kim, and James Agarwal. Internet users' information privacy concerns (IUIPC): The construct, the scale, and a causal model. *Information Systems Research*, 15(4):336–355, 2004.

Veronica Marotta, Vibhanshu Abhishek, and Alessandro Acquisti. Online tracking and publishers' revenues: An empirical analysis. *Workshop on the Economics of Information Security*, 2019. Working paper.

Célestin Matte, Nataliia Bielova, and Cristiana Santos. Do cookie banners respect my choice? Measuring legal compliance of banners from IAB Europe's Transparency and Consent Framework. In *2020 IEEE Symposium on Security and Privacy (SP)*, pages 791–809, 2020. doi: 10.1109/SP40000.2020.00076.

Torsha Mazumdar, Daniel Timko, and Muhammad Lutfor Rahman. Are current CCPA compliant banners conveying user's desired opt-out decisions? An empirical study of cookie consent banners. In *Cryptology and Network Security (CANS 2023)*, Lecture Notes in Computer Science, pages 186–207. Springer Nature Singapore, 2023. doi: 10.1007/978-981-99-7563-1_9.

Erik Miehling, Rahul Nair, Elizabeth Daly, Karthikeyan Natesan Ramamurthy, and Robert Redmond. Cookie consent has disparate impact on estimation accuracy. In *Advances in Neural Information Processing Systems*, volume 36, 2023.

Midas Nouwens, Ilaria Liccardi, Michael Veale, David Karger, and Lalana Kagal. Dark patterns after the GDPR: Scraping consent pop-ups and demonstrating their influence. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (CHI)*, 2020. doi: 10.1145/3313831.3376321.

Midas Nouwens, Janus Bager Kristensen, Kristjan Maalt, and Rolf Bagge. A cross-country analysis of GDPR cookie banners and flexible methods for scraping them. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. ACM, 2025. doi: 10.1145/3706598.3713648.

Frank Papenmeier, Josephine Halama, and Carl Reichert. Accepting cookies: Nudging, deceptive patterns and personal preference. *Computers in Human Behavior*, 168:108641, 2025. doi: 10.1016/j.chb.2025.108641.

Tanya Prasad, Rut Vora, Soo Yee Lim, Nguyen Phong Hoang, and Thomas Pasquier. RegTrack: Uncovering global disparities in third-party advertising and tracking, 2026. URL https://arxiv.org/abs/2603.02679.

Deepak Ravichandran and Nitish Korula. Effect of disabling third-party cookies on publisher revenue. Technical report, Google, 2019.

Iskander Sanchez-Rola, Matteo Dell'Amico, Platon Kotzias, Davide Balzarotti, Leyla Bilge, Pierre-Antoine Vervier, and Igor Santos. Can I opt out yet? GDPR and the global illusion of cookie control. In *Proceedings of the 2019 ACM Asia Conference on Computer and Communications Security (AsiaCCS)*, pages 340–351, 2019. doi: 10.1145/3321705.3329806.

Liv Hilde Sjøflot and Tobias A. Opsahl. Deceptive cookies: Consent by design – a mixed methods study, 2026. URL https://arxiv.org/abs/2605.15056.

Usercentrics. Optimizing consent data and user trust. Whitepaper, Usercentrics GmbH, 2024. Consent management platform data, 1.09 billion banner displays over 30 days in late 2022 and early 2023.

Christine Utz, Martin Degeling, Sascha Fahl, Florian Schaub, and Thorsten Holz. (un)informed consent: Studying GDPR consent notices in the field. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security (CCS)*, pages 973–990, 2019. doi: 10.1145/3319535.3354212.

Rob van Eijk, Hadi Asghari, Philipp Winter, and Arvind Narayanan. The impact of user location on cookie notices (inside and outside of the European Union), 2019. Workshop on Technology and Consumer Protection (ConPro '19); SSRN 3361360.

A. Appendix

A.1. Sample balance

TABLE A1. Summary statistics and balance across experimental conditions

	All	By experimental condition			p
		Opt-In	Anti-tracking	Ad-blocking	
A. Participants					
Age (years)	50.4	51.2	49.5	50.3	0.39
Female (%)	63.0	61.2	62.5	65.3	0.61
Non-white (%)	26.7	24.3	29.8	26.4	0.35
College degree (%)	61.6	63.8	57.3	63.4	0.24
Log household income	10.80	10.79	10.74	10.86	0.24
Employed full time (%)	41.1	39.9	38.7	44.5	0.36
Self-rated IT skill (1 to 5)	3.43	3.45	3.49	3.35	0.18
Privacy concern (standardized)	0.02	0.06	−0.01	0.02	0.65
Participants	789	276	248	265	
B. Banner encounters					
Banner offers a one-click reject (%)	40.3	38.3	43.1	40.2	0.23
Banner offers a customize option (%)	74.5	75.4	74.7	73.1	0.63

Notes: Panel A: one observation per participant (789 participants). Percent rows report shares. Privacy concern averages the eight items of the short Internet Users' Information Privacy Concerns scale, measured in the onboarding survey on 1 to 7 agreement scales, and is standardized to mean zero and unit variance among the 1,253 panel participants who answered the onboarding survey, so its mean in the analysis sample is not exactly zero. p is from a joint test that the three condition means are equal. Panel B: one observation per banner encounter, on the sample that excludes paid-survey and reward websites (Section 2); p from joint tests clustered by participant. Menu availability is read from the recorded banner text, never inferred from clicks. Encounter-level rows describe composition, which is downstream of assignment; every cross-condition estimate in the paper therefore has a participant-level companion.

A.2. Notes to the figures

Figure 1. Stylized banners drawn with the buttons each type offers, in the order in which they appear on the page, and the share of the 16,410 encounters of each type. The buttons

are read from the banner text the extension records (accept, reject, settings, close); the recorded data carry no styling, so every button is drawn alike. Types with a reject button in other orders (settings, accept, reject; reject, accept; and others) are pooled in the last mock-up. Banners whose buttons cannot be read, because the recorded text is truncated, the banner offers only a close button, or the labels are not recognized, are 18 percent of encounters and are not drawn.

Figure 2. Each row shows the share of encounters that end with an explicit choice at one level of a covariate, on the 16,410 encounters of 789 participants. Panel A shows the attitudes of the participant, panel B the banner menu and the website category, and panel C familiarity with the website; the three panels share one axis, and the dashed line is the share over all encounters. The share is the average prediction of a linear probability model of engaging on the covariate, the other covariates of the figure, how widely visited the website is, demographics and month fixed effects, with every encounter set to that level; bars are 95 percent intervals clustered by participant, and the horizontal axis of each panel is fitted to its intervals. Attitudes are split at the sample mean of the standardized onboarding item; familiarity separates the participant's first recorded visit to the website from later ones; categories with fewer than 200 encounters are omitted.

Figure 3. Each panel is one feature, absent (open circle) or present (black dot), joined by a grey segment; rows are four margins. The column on the right reports the difference in percentage points between present and absent. Each point is the average predicted outcome with every encounter set to that version of the feature, from a linear model with website-category and month fixed effects, with 95 percent intervals clustered by participant; an asterisk marks a difference significant at 5 percent. On websites that force a choice, a participant who uses the website without a recorded choice is counted with the websites that do not block, since at that visit the banner did not force a choice, and one whose banner page is not in the navigation log is left out (Section 2). The second and third rows condition on an explicit choice. The last row is the share of all encounters in which the participant leaves the website: the first different page opened in the same tab is reached by going back or typing an address, and no page of the website follows within thirty minutes. A participant who ignores the banner and opens no other page in the same tab within thirty minutes counts as staying, or as leaving on a website that forces a choice. A banner mentions personalization when its recorded text contains "personalized", "personalised"

or "personalization", or mentions improving, enhancing or tailoring the user's experience. A website forces a choice when its page cannot be used until the participant makes a choice in the banner (Appendix A.3). Websites that block with a close X, and websites whose blocking could not be established, are left out of Panel C. Websites choose their own templates and wording, so the comparison is descriptive. To check the rule for leaving on websites that force a choice, we look at the pages the participant opens in any tab in the thirty minutes after the banner. Of the 115 encounters on these websites with no recorded choice and no other page in the same tab, in 88 the participant visits only other websites, in 15 they use the same website in another tab, and in 12 they open no page at all. Counting the second group as staying and leaving out the third, participants leave the website after 14 percent of banners on websites that force a choice, against 7 percent on websites that do not block, where the rule is unchanged: twice as often, against 17 and 7 percent under the rule of Section 2.

Figure 4. The unit is a return: a day on which the participant meets a banner on a website is paired with the next day they visit that website, and the outcome is whether a banner is recorded then. Panel A: the share of returns with a banner again, by what the participant did at the previous visit (ignored the banner, rejected or accepted), with 95 percent intervals clustered by participant. Panel B: the same shares on the 295 websites where both groups are observed. The difference within website, from a linear model with website fixed effects and standard errors clustered by participant, is 26 percentage points ($p < 0.001$); between rejecting and accepting it is 3 percentage points ($p = 0.52$, 62 websites). 6,852 returns, 658 participants, 1,703 websites.

Figure 5. Each row shows the share of encounters engaged with that accept, at one level of a covariate, on the 5,050 encounters engaged with of 691 participants; bars are 95 percent intervals clustered by participant. Construction as in Figure 2, with the same website categories, each figure in decreasing order of its own share.

Figure 6. Each row shows the share of encounters engaged with that reject, at one level of a covariate, on the 5,050 encounters engaged with of 691 participants; bars are 95 percent intervals clustered by participant. Construction as in Figure 2, with the same website categories, each figure in decreasing order of its own share. Accepting and rejecting are estimated in separate models, so the two shares of a category need not sum to 100.

Figure 7. One bar per menu type, with the buttons read from the recorded banner text in page order (Figure 1); types with a reject button in other orders are pooled. Each bar splits the encounters by outcome: accepts, rejects, ignores the banner and stays, leaves the website. On websites that force a choice, a participant who uses the website without a recorded choice is counted with the websites that do not block, since at that visit the banner did not force a choice, and one whose banner page is not in the navigation log is left out (Section 2). Panel A: websites that do not block; Panel B: websites that block but whose banner can be closed with an X; Panel C: websites that force a choice, by the measures of Figure 3. Websites whose blocking or close X could not be established are left out; banners whose buttons cannot be read are not shown. For this reason, participants leave after 16 percent of the encounters shown in Panel C, against 17 percent of all encounters on websites that force a choice (Appendix A.3). Shares are raw, without controls.

Figure 8. The days we observe a participant run from their first to their last page in the navigation log. The signal of a participant is their most frequent explicit choice (accept or reject) in the first third of those days, with ties counted as accepting; it is compared with the explicit choices the participant makes afterwards (520 participants, 2,682 later choices). Consistency is the share of later explicit choices equal to the signal, among the choices that are equal to the signal or opposite to it with the choice of the signal available in one click; an acceptance on a banner without a one-click reject under a $Reject$ signal and ignored banners are left out. Panel A: the later encounters of these participants, split into the choice of the signal, the opposite choice (including acceptances on banners without a one-click reject under a $Reject$ signal), and ignored banners. Panel B: participants classified by their choices in the first third of their days (only accepts, accepts and rejects, only rejects); 'Accept all' and 'Reject all' apply the same signal to every participant. Panel C: the signal is the most frequent reaction in the first third of the days, accepting, rejecting or ignoring the banner, and consistency is the share of all later encounters in which the participant has the same reaction (the same 520 participants, 9,011 later encounters); 'Ignore all' applies ignoring to every participant.

Figure 9. Encounters after the first third of each participant's observed days, for the 520 participants of Section 4.1. Concern groups are the terciles of the average of the eight items of the short Internet Users' Information Privacy Concerns scale (Section 2): low (189 participants), medium (152) and high (179). Panel A: share of later encounters that end

with an acceptance, a rejection, or an ignored banner. Panel B: share of later encounters consistent with the choice closest to stated concern ($Reject$ for high concern, $Accept$ for low), counting an ignored banner as not consenting or as being tracked.

Figure A2. Participants are classified by their explicit choices on banners that offer a one-click reject (at least two choices between accepting and rejecting): 113 reject more often than they accept, 350 accept more often than they reject, among the participants who also meet banners without the button. Bars: percent of each group's encounters on banners with (light) and without (dark) a reject button that end with each outcome; ignoring and leaving follow the definitions of Section 2. Banners with and without the button appear on different websites.

Figure A3. Analysis sample: participants paid for completing the onboarding and active on at least 14 days (789 participants). Currently enrolled only: 687 participants. At least three encounters: 688 participants. Shares of encounters engaged with, overall and by banner feature and familiarity; shares of choices that accept, and that reject where a reject button is offered; consistency of the signal set in the first third of each participant's observed days, on explicit choices and with ignoring counted as an outcome (Section 4.1). Shares are raw, without controls. The largest difference from the analysis sample is in engagement on repeat visits, 17 percent among the participants currently enrolled against 15 percent in the analysis sample (raw shares, while the 35 and 19 percent of Section 3.1 are adjusted for the covariates of Figure 2); every other share differs by about one percentage point or less. Blocking is known for 99 percent of encounters.

A.3. Websites that block and websites that force a choice

A website blocks when its page cannot be used until the participant clicks on the banner, either on one of its options or on a close button. A website forces a choice when it blocks and its banner cannot be closed with an X without choosing. Both are measured once per website, from four sources taken in order of priority. First, for OneTrust banners, which are 85 percent of encounters, the snapshots of the pages saved during the panel: a banner blocks when its dialog is marked as modal in at least one snapshot. This is the only source that describes the banner at the time of the visit; on a check by hand, it identified 8 of 8 blocking banners, with 1 false positive in 28. Second, a browser audit, which loads the website, tries 36 points outside the banner and the scroll wheel, and counts the banner as

blocking if none of them reaches the page. Third, the configuration the website publishes on OneTrust, which states whether the banner blocks the page for a visitor from the United States. Fourth, for CookieYes and Didomi banners, the recorded selector of the banner, a pop-up or an overlay. The audit and the published configuration describe the website when we measured it, not at the time of the visit.

For the websites that these four sources leave unclassified, we use three further sources, in this order. First, a second browser audit, run in October 2026 with the same 36-point test on the homepage and then on the page where the participant saw the banner, together with the configuration published by the other consent management platforms (Cookiebot, CookieScript, CookieHub, Iubenda) and the OneTrust rules for each U.S. state; on 44 websites already classified by the earlier sources, it classifies 38 and agrees with the earlier classification in 35 of them. Second, the archived copies of the website in the Internet Archive within six months of the visit, from which we read the consent management platform used at that time and its configuration for U.S. visitors; on 40 websites already classified, it classifies 21 and agrees on whether the website blocks in 20 of them. Third, the navigation log: on a website still unclassified, an encounter in which a participant ignores the banner and then follows a link or submits a form on the page shows that the banner did not force a choice at that visit, and the website is counted as not blocking. Banners shown only to visitors from some U.S. states are classified by the rule for those states. These three sources reduce the websites that could not be classified from about 580 to 88.

Whether a blocking banner can be closed with an X is read from the configuration the website publishes on OneTrust, from the banner saved by the browser audit, from a click on a close button recorded during the panel, or from recorded banner text that ends with a close button. On a website classified as forcing a choice, an encounter in which the participant uses the page without a recorded choice is counted with the websites that do not block. Of the 4,747 websites of the encounter sample, 3,812 do not block, 418 block with a close X, 429 force a choice and 88 could not be classified. The 88 websites that could not be classified (1 percent of encounters) are kept as a separate category in the regressions and left out of the panels of Figure 7. Counting them all as forcing a choice raises the share of encounters in which participants leave after a banner on websites that force a choice from 17 to 20 percent, and counting them all as not blocking leaves it unchanged.

A.4. Website taxonomy

Every website is classified once, by its registered domain, so that all subdomains of a website share a category unless one is listed separately. The categories come from fixed rules (lists of domains and name patterns), from a language model for the domains the rules leave unclassified, and from a classification by hand of every website of the encounter sample that the first two steps leave unclassified, followed by an audit of the most visited websites of each category. A website whose subdomains fall in different categories is counted in each of them, so the websites in Table A2 add up to more than the 4,747 websites of the sample. Table A2 lists, for each category used in the paper, the number of encounters and websites in the encounter sample and the three websites with the most encounters.

TABLE A2. Website categories in the encounter sample

Category	Encounters	Websites	Websites with the most encounters
Shopping	4,874	1,794	cvs.com, ikea.com, acehardware.com
News and lifestyle	2,888	417	people.com, allrecipes.com, fandom.com
Health	1,515	562	aarp.org, aetna.com, labcorp.com
Finance	1,504	338	intuit.com, cash.app, synchrony.com
Productivity	1,496	436	zoom.us, speedtest.net, ups.com
Education	1,244	423	stackexchange.com, sciencedirect.com, stackoverflow.com
Travel	854	264	delta.com, expedia.com, flightaware.com
Gaming	343	109	ignitioncasino.eu, baseball-trivia.com, wheeloffortune.com
Streaming	338	108	fandango.com, iheart.com, siriusxm.com
Support and account pages	288	80	hp.com, aarp.org, siriusxm.com
Job listings	247	83	paylocity.com, cvshealth.com, jobbnorge.no
Reviews	225	20	trustpilot.com, rottentomatoes.com, metacritic.com
Coupons and rewards	211	59	mygiftcardsplus.com, studentbeans.com, hometesterclub.com
Government and non-profit	138	98	tylerportico.com, obama.org, adoptsheboygancounty.org
Utilities	126	30	t-mobile.com, consumercellular.com, pge.com
Social media	91	21	linktr.ee, connectnetwork.com, gettingout.com

Notes: The 16,410 encounters of the analysis sample; categories with fewer than 200 encounters are omitted from the figures. AI assistants, email, adult websites and search engines account for the remaining 28 encounters.

A.5. Prediction and variance decomposition of engaging, accepting and rejecting

To measure how well the stated attitudes predict participants' choices, we use the attitudes and choices of a random half of the participants to fit a model of engaging, and of accepting

among explicit choices, and use the model to predict the choices of the other half. We repeat this with five different random halves. We report the AUC, the probability that the model ranks a randomly chosen banner the participant engaged with above a randomly chosen banner they ignored: 0.50 corresponds to tossing a coin and 1 to a perfect prediction. The attitudes predict engaging with an AUC of 0.50, and accepting with an AUC of 0.67.

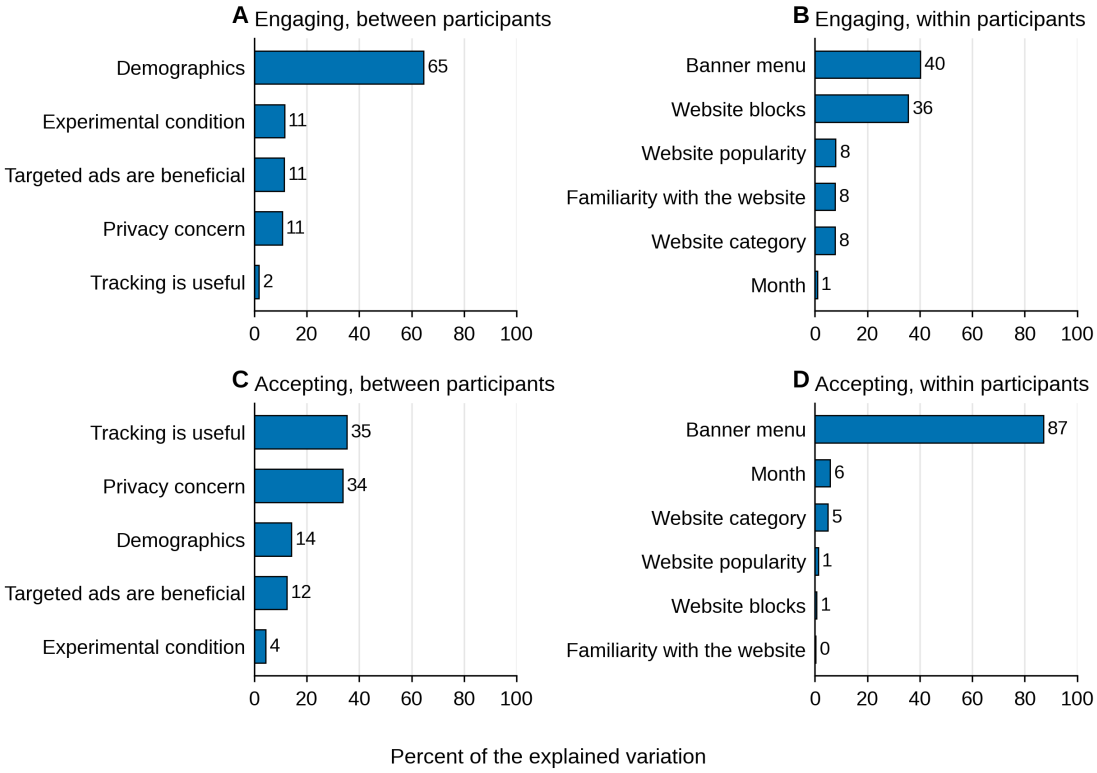

FIGURE A1. Variance decomposition of engaging and accepting. Total variance is split into between- and within-participant parts with participant fixed effects; each part is then decomposed exactly, by Shapley value, over the predictor blocks able to operate in it.

Figure A1 splits the variance of engaging and of accepting into a part between participants and a part within participants, and decomposes each part exactly, by Shapley value, over the blocks of predictors that can operate in it. The covariates explain 17 percent of the variation in engaging within participants and 10 percent of the variation in accepting between participants; the shares below refer to these explained parts. Engaging varies within participants: 84 percent of its variance is the same person meeting different banners, and of the within part the covariates explain, the banner's menu accounts for 40 percent, whether the website blocks until a choice is made for 36, and how widely visited the website is, the website category and familiarity with the website for 8 each.

Accepting varies between participants: 56 percent of its variance lies between people, and of the between part the covariates explain, the belief that tracking is useful accounts for 35 percent, stated concern for 34 and the advertising beliefs for 12. Rejecting is split in the same way, 56 percent between participants and 44 percent within them: between participants, the belief that tracking is useful accounts for 35 percent of the explained part and stated concern for 34; within participants, the banner's menu accounts for 87 percent, and whether the website blocks for none.

A.6. Banners without a reject button

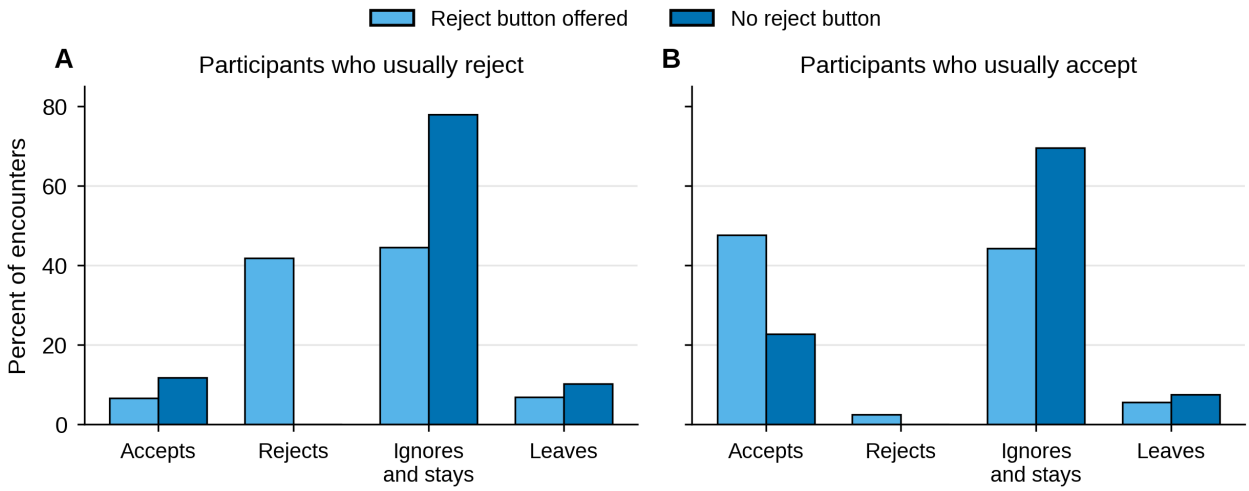

FIGURE A2. What participants do on banners with and without a one-click reject, by the choice they usually make where it is offered. Construction in Appendix A.2.

A.7. Robustness to the sample

Figure A3 reports the main shares on the analysis sample and on two alternative samples: the participants currently enrolled in the study, and the participants of the analysis sample with at least three banner encounters. We also repeat the main shares including the encounters on paid-survey and reward websites (19,786 encounters of 795 participants). Including those websites lowers the shares of banners engaged with, 27.7 against 30.8 percent overall and 38.8 against 46.7 percent when the banner offers a one-click reject, and leaves it almost unchanged when the banner does not offer one (19.8 against 19.6 percent), while the shares of explicit choices change little: participants accept 84.4 against 84.3 percent of the banners they engage with, and reject 24.2 against 25.9 percent of those with a one-click reject.

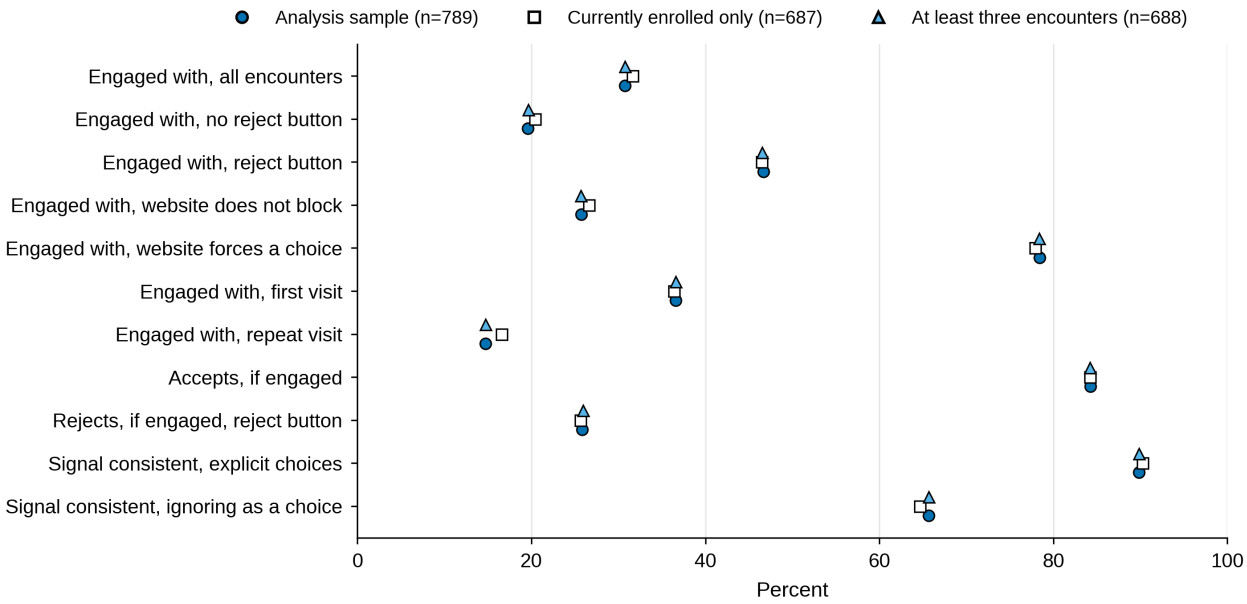

FIGURE A3. Main shares on the analysis sample and on two alternative samples. Construction in Appendix A.2.