



HAL
open science

The Effect of Design Patterns on (Present and Future) Cookie Consent Decisions

Nataliia Bielova, Laura Litvine, Anysia Nguyen, Mariam Chammat, Vincent Toubiana, Estelle Hary

► To cite this version:

Nataliia Bielova, Laura Litvine, Anysia Nguyen, Mariam Chammat, Vincent Toubiana, et al.. The Effect of Design Patterns on (Present and Future) Cookie Consent Decisions. USENIX Security 2024 - 33rd USENIX Security Symposium, Aug 2024, Philadelphia (PA), United States. pp.2813–2830. <hal-04913806>

HAL Id: hal-04913806

<https://hal.science/hal-04913806v1>

Submitted on 27 Jan 2025

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire HAL, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons CC BY 4.0 - Attribution - International License

The Effect of Design Patterns on (Present and Future) Cookie Consent Decisions

Nataliia Bielova*

Inria research centre at Université Côte d’Azur

Anysia Nguyen

Behavioural Insights Team (BIT)

Vincent Toubiana†

Commission Nationale de l’Informatique et des Libertés (CNIL)

Laura Litvine

Behavioural Insights Team (BIT)

Mariam Chammat

Interministerial Directorate for Public Transformation (DITP)

Estelle Hary‡

RMIT University

Abstract

Today most websites in the EU present users with a consent banner asking about the use of cookies or other tracking technologies. Data Protection Authorities (DPAs) need to ensure that users can express their true preferences when faced with these banners, while simultaneously satisfying the EU GDPR requirements. To address the needs of the French DPA, we conducted an online experiment among 3,947 participants in France exploring the impact of six different consent banner designs on the outcome of users’ consent decision. We also assessed participants’ knowledge and privacy preferences, as well as satisfaction with the banners. In contrast with previous results, we found that a “bright pattern” that highlights the decline option has a substantial effect on users’ decisions. We also find that two new designs based on behavioral levers have the strongest effect on the outcome of the consent decision, and participants’ satisfaction with the banners. Finally, our study provides novel evidence that the effect of design persists in a short time frame: designs can significantly affect users’ future choices, even when faced with neutral banners.

1 Introduction

Today almost any website in the EU presents the user with a consent banner asking about the use of cookies or other tracking technologies. While the EU ePrivacy Directive [29]¹ requires a legally valid consent before reading or writing of cookies and use of other tracking technologies, the EU General Data Protection Regulation (GDPR) [66] has set high-level requirements for such consent to be valid, leaving website owners a lot of margin to design consent banners.

*The work was primary carried out while Nataliia Bielova was a Senior Privacy Fellow at the LINC lab of the CNIL in 2021-2022.

†The views and opinions expressed in this paper do not necessarily reflect the views of the CNIL or any individual Commissioner.

‡The work was carried out while Estelle Harry was a designer at the LINC lab of the CNIL before May 2023.

¹ePrivacy Directive was last amended in 2009 (ePD) and is known among computer scientists and website developers as “cookie law”. The upgrade of the ePD into an ePrivacy Regulation is currently under discussion.

Although consent banners are a mean massively adopted by the industry to collect such legally-required consent, in practice, their usage raises numerous issues from both a regulatory and a behavioral stand-point.

From a regulatory perspective, EU Data Protection Authorities that are in charge of enforcing the GDPR and – for some of them – ePrivacy, help websites to comply by providing high-level guidelines and practical recommendations. The French Data Protection Authority (CNIL) has recently updated its guidelines [13] and recommendations [14] regarding the use of cookies and other tracking technologies, and consent banner interfaces. While providing more concrete recommendations, regulators still need to ensure that *users can reflect their true preferences in the banner interface*, thus making “freely given, specific, informed and unambiguous indication of the data subject’s wishes”, requested by the GDPR [66, Art.4]. However, no definite proof exists today that consent banners (even in their recommended format) allow users to express their true preferences.

From a behavioral perspective, multiple studies observe a high level of discrepancy between users’ preferences and their decisions. This can largely be attributed to the “choice architecture” - the environment and form in which choices are presented [65]. Recent research shows [43] that small alterations in the choice architecture can have substantial effects on decision making. In the digital environment, such intentional alterations are commonly referred to as “dark patterns” or “deceptive design”, a term that refers to any practices that direct, deceive, coerce, or manipulate users into making choices that are often not in their best interest [10, 37]. A recent study by the European Commission [22] demonstrated that deceptive design is particularly common on EU websites and mobile apps. Dark patterns have also been found in consent banners of numerous websites in recent research [53, 54, 62].

The concept of “dark patterns” has lately gathered interest, appearing recently in the Digital Services Act [26], the Digital Markets Act [25] and the Data Act proposal [24] published by the EU. Similarly, the European Data Protection Board (EDPB) has recently released guidance on dark patterns in

social media [27]. These advances have again increased and confirmed the demand from regulators – not just DPAs – for rigorous evidence on the impact of such practices on users.

The evaluation of the impact of dark patterns on users’ decisions when facing consent banners is a very active research effort [4, 39, 40, 52, 54, 68]. However, studies differ in their number of participants, context of the experiment, and targeted population, providing different results for the same type of consent banners². Additionally, Grassl et al. [39] found a consent rate for UK users of 94% when exposed to dark patterns and 53% when they are exposed to an alternative design, called “bright” patterns. This suggests that users who have been continuously exposed to dark patterns might have already developed the habit of accepting, and exposing them to “bright patterns” might change their behavior. This hypothesis, however, has not been rigorously tested so far.

While these studies offer important insights to regulators, consultations with experts from the CNIL, that participated to this project, have revealed that these studies may not be specific enough to be used in guidelines or to inform policy makers about the effect of banner design, because they may test slightly different designs and be performed on non-French population. This is especially the case when no consensus in prior studies has been found on the effect of dark (and bright) patterns. To be used in a legal document, the result of a study on cookie banners should not be subject to legal uncertainty.

In the context of this research, conducted as a *partnership between CNIL and external researchers*, the services of CNIL asked a French public agency to perform a large scale study on French population to evaluate the impact of cookie banners design on consent rates. This study aims to enlighten CNIL’s legally-binding decisions and recommendations related to cookie banners. Our study therefore aims, not only to contribute to the active literature on the impact of dark patterns on user decisions, but also to directly address the requirements and needs of the services of CNIL, which has fined several companies on the topic [15, 20]. Our overarching goal is to *study the factors that could prevent users from expressing their true privacy preference when faced with a consent banner, and to propose levers allowing them to express their true choice*. More specifically, we conducted a controlled online experiment with 3,947 participants in France to answer the following research questions:

- **RQ1:** What is the state of knowledge of users about cookies and tracking technologies, and their general preferences on data sharing online?
- **RQ2:** What is the impact of dark pattern design on the outcome of users’ consent decision in France?

²Studies range from 40 [54] to 2,000 [68] participants per banner, and were done with German and Austrian [33, 52, 68], Danish [2], UK [39], and US [40, 51, 54] participants. For example, in a neutral banner with identical “accept” and “decline” buttons, 72% of German participants decide to accept [68], while only 60% of US participants do so [40].

- **RQ3:** Can new designs based on behavioral levers allow users to make consent choices closer to their intention?
- **RQ4:** Does the effect of dark patterns and behavioral levers on user consent decision persist over time, even in the absence of these design patterns?

To answer these questions, and based on a close guidance from the services of CNIL, we have designed six variants of a consent banner (§3). We then conducted a between-subjects online experiment with 3,947 participants in France to test the impact of these design variants (§4). By analysing participants’ interactions with the banners and survey responses, we evaluate users’ knowledge and preferences for data sharing, the impact of banner design on users’ consent decisions, as well as user sentiment and satisfaction. Finally, we discover that banners design can significantly affect users choices in a short-term (§5). We then discuss our findings and their implications (§6).

2 Background and related work

2.1 EU legal requirements on consent and the power of Data Protection Authorities

GDPR and ePrivacy directive. The EU ePrivacy directive first published in 2002 was amended in 2009 [29]. The regulation of cookies, and subsequently of consent banners, comes from Article 5.3 of ePrivacy which specifies that “storing of information, or the gaining of access to information already stored, in the terminal equipment of a subscriber or user is only allowed on condition that the subscriber or user concerned has given his or her consent [...]”. The validity of consent is evaluated according to the Directive 95/46/EC, which has been superseded by the EU General Data Protection Regulation (GDPR) that came in force in May 2018 [66]. According to the GDPR, “‘consent’ of the data subject means any freely given, specific, informed and unambiguous indication of the data subject’s wishes” [66, Art.4]. Unlike the GDPR, ePrivacy is not subject to the one-stop-shop³ and hence every ePrivacy national regulator (in France, the Data Protection Authority) is competent to build a case where it finds local non-compliance with Art.5.3 of the ePrivacy directive.

Non-binding: recommendations, EDPB guidelines. In order to clarify what is expected for a consent to be valid, the EDPB has published guidelines on consent [30]. To provide more legal certainty, DPAs can publish guidelines on how to comply with a given regulation. These guidelines cannot proscribe practices that are non-compliant: instead, to limit legal uncertainty, guidelines list examples and practices of how to comply with the regulation. However, such guidelines are non-binding: it is possible to comply with a regulation

³It’s a mechanism which defines the lead authority as the authority in which the EU headquarters of the company are established.

without following them. In France, the CNIL recommendations [14] contain two examples of consent banners that are considered compliant. In practice, there are many types of consent banners that differ from these examples and yet are compliant with the regulation.

Binding: case law. Jurisprudence is built on previous cases that have been published, and therefore has a relatively narrow scope compared to recommendations that attempt to cover many cases. However, unlike recommendations, jurisprudence is legally-binding. For example, recent jurisprudence in France indicates that consent banners must offer an option to reject cookies on the first layer [16, 17].

2.2 Impact of design on consent decisions

Even though the regulation is aimed at empowering users to make choices that reflect their true preferences, multiple studies observe a high level of discrepancy between users' preferences and their decisions. Kulyk et al. [48] demonstrate that users click "accept" in consent banners despite having concerns over data collection. Utz et al. [68], Borberg et al. [7] and Habib et al. [40] explored reasons why users accept consent and found reasons such as decision fatigue, being in a rush, thinking that consent is required for the website functionality or simply to dismiss the consent interface.

Starting with the EU Commission report in 2016 [69], various research studies then aimed at evaluating the impact of consent banner design on users' decisions. Utz et al. [68] measured users' decisions on four types of banners, as well as variations of those banners incorporating "nudges" towards acceptance. They also found that 48% of users believe that clicking "decline" on a neutral banner would block or break the website. Nouwens et al. [54] have explicitly evaluated the effect of dark patterns in consent banners: they found that moving "decline all" button to the second layer of the banner increased the probability of acceptance by 23 percentage points. Machuletz et al. [52] analysed how visualisation and control over purposes, as well as dark patterns, can impact user decisions. Bermejo Fernandez et al. [4] found that color-based nudging bars that visualise the amount of accepted cookie purposes can significantly impact the participants' decisions to change the default cookie settings. Grassl et al. [39] were the first to measure the effect of "bright patterns" on users' consent decisions and found that placing the decline option on the second layer and pre-selecting a decline option substantially influenced users' consent decisions.

Bielova et al. [6] compared 17 guidelines of EU DPAs and the EDPB and 11 consent-focused user studies and found out that the guidelines and user studies rarely impact each other: their analysis yielded 11 gaps, 5 insights and only 3 consistencies [6]. Singh et al. [60] studied which consent banners users would prefer if they could choose and identified five key design factors which impacted participants' liking/disliking of consent notice UI designs: ease of use, amount

of information, customisability, decision-making time, and clarity/transparency.

3 Selecting design parameters

Recently, CNIL has made a chain of repressive actions against websites with consent banners that do not lead to a valid consent⁴. When looking to select design parameters to be included in our experiment, we looked beyond the current state of knowledge at the needs of regulatory authorities building such cases. Indeed, the primary goal of this study was to ensure that our learnings could be actioned by DPAs. We designed the study specifically to inform policy-making in France: the types of dark banners match some banners that appear on French websites, and the study was conducted with French participants. The outputs of this study could therefore enlighten the French regulators and policy makers when defining future guidelines on cookie banners or to inform future discussions on the design of cookie banners and its significant impact on consent rate.

Recommended consent design. The CNIL's latest guidelines [13] require that "the expression of the user's refusal must [...] be able to be translated by an action presenting the same degree of simplicity as the one allowing to express his consent". While further interpreting the ePrivacy Directive [29] in light of the notion of consent under the GDPR [66, Art.4], CNIL strongly recommends that the mechanism allowing to express a refusal to consent to read and/or write operations⁵ should be accessible *on the same screen and with the same ease* as the mechanism allowing to express consent [14, par. 31]. Additionally, CNIL provides a concrete example of a recommended consent banner, where "accept all" and "decline all" buttons are presented in visually identical ways [14, par. 34]. We therefore adopt the design recommended by the CNIL as a control banner to compare users' decisions in other banners against it. The LINC, CNIL's innovation Lab, has also recently analyzed the related works on the impact of cookie banner design [50] and identified that prior works [4, 33, 40, 52] do not compare the tested designs against a control banner – we hence do not compare our methods and results to these works.

3.1 Selecting dark patterns

A number of exchanges with employees of the CNIL allowed us to identify further design parameters that may prevent users from expressing their true preferences, and provide robust evidence of the impact of these dark patterns on users' behavior in France (RQ2).

⁴Over the last two years, several companies have for example been sanctioned by the CNIL for making it more difficult to refuse cookies than accepting them [16–19].

⁵That is, to read, set or send cookies or use other tracking technologies in the browser.

3.1.1 Unequal path to decline

Our first consent design with dark pattern, called “unequal path to decline”, allows users to decline only on the second layer of the banner, accessible under “customize my choices” option. Beyond the fact that such design is not recommended by the CNIL, it has been considered a violation of the regulation since “several clicks are required to refuse all cookies, against a single one to accept them.” [17].

No large-scale study reached conclusions. Even though ten studies have measured the effect of dark patterns in consent banners between 2020 and 2022 [50], only two of them [39, 54] have evaluated the effect of a consent banner design with “unequal path to decline” with respect to the control banner and reached different conclusions⁶. Nouwens et al. [54] have tested banners with 40 US students and reported a statistically significant difference in the outcome of consent decisions (45% refusal rate for control banner vs 23% for “unequal path” banner, $p < 0.001$). Grassl et al. [39], based on behavior of 228 UK users, have reported that no statistically significant difference was observed among two banners. Therefore, no study up to now have reached a confirmation on the effect of such banners on a large population, and no study was conducted with French users.

Dark patterns. First, this consent design presents a configuration barrier that directs users towards acceptance of consent [36]: it contains an *Obstruction* dark pattern, because it hides the option to decline behind the “customize my choices” link [62], and according to the EDPB guidelines [27], such design contains *Longer than necessary*⁷ pattern that requires more steps for the users to choose a more privacy-friendly option. It also includes *Interface Interference*, *Manipulating Visual Choice Architecture* and a low-level pattern of *False Hierarchy* [38] since the option to accept is highly visible, while the option to decline is hidden behind a less visible “customize my choices” link.

Interest of the services of the CNIL. Since there are only two peer-review comprehensive academic studies that evaluate this particular design of consent banner, and that none was large-scale and done exclusively with the French population, CNIL employees have expressed interest in including this banner design.

3.1.2 Visual aid: weighted-accept

Even when placing both accept and decline buttons together on the banner’s first layer, websites still implement other dark

⁶We exclude the study of Utz et al. [68] because even though the authors included the banner with no decline option on the first layer, their banner did not provide the second layer to decline all trackers, and therefore all users who interacted with this banner have clicked “accept all”, and thus is incomparable to our banner design. Bouma-Sims et al. [8] in 2023 included a similar design (options-link banner in Fig.18), however did not report on comparing it to the banner with equally presented “accept” and “decline”.

⁷This dark pattern is called “Adding Steps” in the latest ontology of Gray et al. [38].

patterns, for example by highlighting the accept button, making it more visually prominent.

No significant impact according to previous research.

Three academic studies have analyzed users’ reactions upon interacting with such a banner design. Utz et al. [68] found that around 79% of German users click “accept all” on this type of design, while 72% users do so on the control banner with equally shown accept and decline buttons, however authors did not provide any testing for statistical significance of this finding. Grassl et al. [39] experimented with 228 UK users and did not find any statistically significant difference in consent decisions between users interacting with this type of banner versus a neutral banner. Berens et al. [3, Table 8] also did not find any statistically significant difference between such types of banners. Therefore, no studies found support for difference in users behavior.

Dark patterns. This design pattern presents a configuration barrier that directs users towards acceptance of consent [36]. It therefore includes a dark pattern called *Manipulating Visual Choice Architecture* [38], where privacy-unfriendly user choices are perceived more salient and prioritized. This design implements a dark pattern called *False Hierarchy* [37] giving users a more visual precedence over “accept all” button than “decline all” button. From the recent EDPB guidelines [27], such design contains a dark pattern called *Look over here*, where “action or information is put in competition with another element” and thus nudges users using visual style.

Interest of the services of the CNIL. While CNIL recommends [14] that accept and decline options should be accessible on the same screen and with the same ease, the regulator does not have further requirements on concrete design and colors of the buttons. As a result, CNIL’s employees expressed interest in testing this type of design in the experiment - the results of the study could help CNIL to update its recommendations in the future.

3.2 Selecting behavioral levers

CNIL employees have also expressed interest, by collaborating with behavioral and social scientists (co-authors of this paper), in experimenting with banner designs that are not explicitly mentioned in their recommendations yet, but could help users to express their true privacy preferences (RQ3). To do that, we reviewed the prior works on privacy [1, 12, 40, 64], and the broader literature on decision-making [23, 46, 47, 59].

The review first aimed to identify the most likely barriers to users making informed choices, which were grouped into two key barriers:

1. Users still lack knowledge on cookies. The literature [44, 56, 61] indeed showed that despite the presence of information on some cookie banners, internet users have limited knowledge about cookies and their purposes. This is exacerbated by the fact that they are not helped by cookie banners that are often unclear (including legal jargon, long

sentences, etc.), as well as information that is often partial or even biased, framed to highlight benefits of cookies without mentioning potential costs to the user [58].

2. Users lack motivation to actively make informed consent choices. This lack of motivation is linked to a desire for speed above all, which leads to the development of mental shortcuts, and cost-benefits calculations (e.g. instant access to the site vs. data protection) which in the short-run often favour the sharing of information.

As a second step, we used both the behavioural literature and design principles to identify levers that could help lift these barriers. As a result, we proposed 18 potential design solutions such as simplifying the text, using heuristics, or changing the choice architecture and the display of buttons to make refusing easier or more desirable (a full list can be found in the Supplemental Materials [5, Appendix C]). To select behavioral levers for our experiment, we organised a workshop with 7 participants, including four CNIL agents (a lawyer with expertise on sanctions, a design expert, a law and technology expert, and a computer scientist), and three additional experts outside the CNIL (a behavioral economist, an online experimentation expert, and an expert in cognitive neuroscience).

During the workshop, all participants ranked 18 behavioural levers based on two criteria: potential impact and implementation feasibility. To evaluate a lever's *potential impact*, we asked participants to reflect on the following questions: (1) What proportion of users would this reach? Is it adapted to a specific population of users or all of them? (2) What is the expected effect of the design on refusal rate? (3) Could this backfire and lead to negative outcomes such as prevent people from making an informed choice?

To evaluate the *implementation feasibility* of the 18 levers and design changes, we asked the participants : (1) Could the CNIL potentially recommend this type of design? (2) Are there any judicial considerations and implications that don't allow this type of banner? (3) Could this design be easily scalable to all consent banners if it improved informed consent? With these questions in mind, each participant had to give a score of potential impact and implementation feasibility, from 1 (low) to 3 (high) for each lever.

Using the results of this voting system (presented in Supplemental Materials [5, Appendix C]), we selected three behavioral levers to be included in the banners: (1) Highlight consequences of their choice; (2) Pre-select "Refuse all"; (3) Use visual aids that can be used as heuristics. We further discuss each of the levers in the rest of this section.

3.2.1 Highlight consequences of choice

Behavioral Lever. The first lever seeks to emphasize the negative consequences of acceptance and to remind users that they are consenting to the use of trackers. This design aims to re-balance the little importance Web users give to future

risks, sometimes overshadowed by a need for speed. This lever also aims to create an element of surprise, which would provoke reflection and allow to interrupt the acquired habit of mechanically clicking on "accept all".

Interest of the services of the CNIL. This lever has been particularly supported by CNIL agents as it makes visible, in the act of selecting an option, the consequences of the choices. CNIL guidelines on consent [13, par.24], CNIL recommendations [14, par.8], and GDPR require that for consent to be valid, users should be informed of the *consequences of their choices*.

Previous work suggests users may be impacted by the text on banners' buttons. Habib et al. [40, Section 4.4.1] were the first to compare generic text labels on buttons ("accept" and "decline") with specific ("accept all" and "necessary-only") labels and found no significant impact on users consent decisions. Berens et al. [3, Table 9] also found no statistically significant differences between the different text labels placed on decline button. Finally, Ma and Birrell [51, Table 1] tested five combinations of specific texts placed on accept and decline buttons (using positive/negative slant that refers to the utility of cookies and positive/negative framing about benefits or harms of the choice) - the authors found a statistically significant impact of some of the button text on users' consent decisions. As a result, we decided to test such behavioral lever by emphasizing the consequences of choice within the text of the accept and decline buttons and specifically included the consequences of being tracked in the button's text labels, which has not been evaluated in previous works.

3.2.2 Visual aid: weighted-decline

Behavioral Lever. The second lever functions as a "mirror" to the *weighted-accept* dark pattern described above, this time visually emphasizing rejection rather than acceptance. This lever aims to tap into users' tendency to favor pre-selected options [46], but this time seeking to counterbalance acceptance habits developed over time. This "mirror" implementation allows us to compare the effect of the same design feature (the pre-selection of an option), but used for a different purpose.

No significant impact according to previous works. Grassl et al. [39] first evaluated the impact on the users' consent decisions of *bright patterns*, where the design nudges are reversed towards the privacy-friendly option. A test on 255 UK users of a design where the decline option is more visually prominent showed no support for the hypothesis that such design has a substantial effect on the user consent decision. Berens et al. [3, Table 8] have further studied the impact of visual representation of "highlighted-reject" with 100 users per banner, and, similarly to Grassl et al., found no statistically significant difference between "highlighted-reject" and control banner. We include this behavioral lever in our study with a large sample of French users to re-evaluate the impact of such design on users decisions.

Table 1: List of design parameters, possible implementations for each, and the corresponding dark patterns or behavioral levers.

Design Parameter	Possible Implementations	Dark patterns (DP) / Behavioral levers (BL)
<i>Path to decline</i>	<i>equal</i> (accept and decline on 1st layer)	-
	<i>unequal</i> (decline is only accessible on the 2nd layer)	DP: Obstruction [62], Interface Interference [37], Longer than necessary [32]
<i>Text within button options</i>	<i>generic</i> (“Accept all” and “Decline all”)	-
	<i>consequences-of-choice</i> (“Accept to be tracked” and “Continue without being tracked”)	BL: Element of surprise, negative consequences of acceptance
<i>Visual aid in button layout</i>	<i>neutral</i> (accept and decline buttons are shown with identical neutral color, equal shape and size)	-
	<i>weighted-accept</i> (accept button is more visually prominent than decline)	DP: Aesthetic Manipulation [37], False hierarchy [37], Hidden in Plain Sight [32]
	<i>weighted-decline</i> (decline is more visually prominent - opposite to weighted-accept)	BL: Tendency to favor pre-selected options
	<i>traffic-light</i> (highlight buttons in the colors of traffic light and use other familiar pictograms)	BL: Familiar visual metaphors and pictograms

3.2.3 Visual aid: traffic lights

Behavioral Lever. The third lever selected during the workshop is based on visual aids. It uses “traffic light” colours as well as other familiar visual metaphors (shield, attention sign) to link refusal, in users’ minds, to a recommended, safer option, which is equivalent to a higher data protection. This design relies on the finding that when making choices, especially when they are made in an environment with a lot of information and in a context where users are goal-oriented and have little time (all characteristics of online choices), users tend to use mental shortcuts (heuristics), to get to a decision quicker and preserve their mental bandwidth [67].

This is why users can tend to click on highlighted options (see “weighted decline” in Section 3.2.2), but is also a trait of user decision-making that could be used to try and direct behaviour towards the refusal of cookies, which the user might then implicitly perceive as the “safer” or “green” option.

Previous work has shown support for this type of heuristics in another context. Egelman et al. [28] showed that coloured bars illustrating the complexity of a password encouraged users to create more complex ones. Nevertheless, such design has never been tested previously in the context of consent banners.

3.3 Designing consent banners

Following the methodology set out in Habib et al. [40, Table 3]⁸, we have identified three key design parameters of consent banner interfaces, which we vary to implement the dark patterns and behavioral levers: the nature of the path to decline (equal or unequal), the text within the button options, and the

⁸Differently from Habib et al., we map implementations of each parameter to the corresponding dark patterns and behavioral levers that we have identified.

Table 2: Overview of the six consent banners design variants and their values for the design parameters in our experiment. Design choices that differ from *Control* banner, are bolded.

Banner name	<i>Path to decline</i>	<i>Text within button options</i>	<i>Visual aid</i>
<i>Control</i>	equal	generic	neutral
<i>No decline</i>	unequal	generic	neutral
<i>Highlighted accept</i>	equal	generic	weighted-accept
<i>Highlighted decline</i>	equal	generic	weighted-decline
<i>Consequences</i>	equal	consequences of choice	neutral
<i>Tricolor</i>	equal	generic	traffic light

presence (or not) of visual aid in the button layout. The *control banner* is without dark patterns or behavioral levers, and shows both “Accept All” and “Decline All” buttons in an identical style and size, following CNIL recommendations [14]. A summary of the chosen design parameters, as well as their implementations with dark patterns and behavioral levers is shown in Table 1. Table 2 provides an overview of the six variants of consent banners created: one *control* banner, 2 *dark patterns* banners, and 3 *behavioral levers* banners (all banners are provided in Appendix B).

In order to isolate the effect of the design from that of the text, the main text in the banners remained fixed on all six banners. The text used was developed on the basis of the current CNIL recommendations [14] and made as short and simple as possible, following key insights from efficient communications. The text has been further reviewed by the three experts of the CNIL to be as clear as possible, compliant

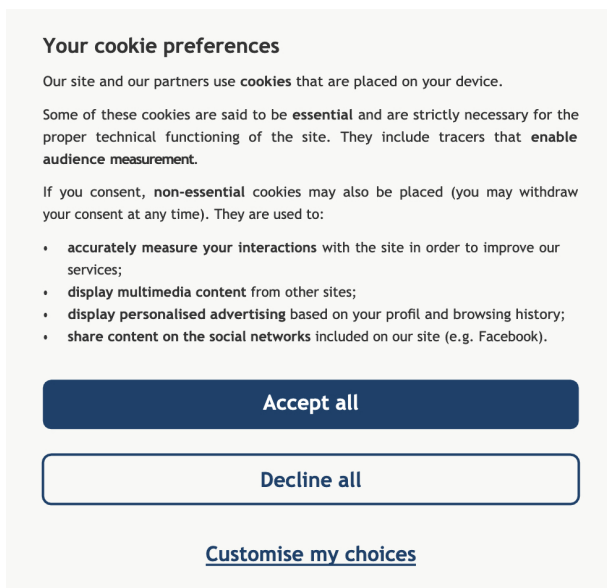


Figure 1: *Highlighted accept* banner included in the online experiment (English translation).

with the current recommendations and not to present partial or biased information⁹.

Defining minimal contrast levels. To choose the concrete design for accept and decline buttons, we were guided by the following reasoning. The European Data Protection Board (EDPB) has recently published the result of the joint Task Force on consent banners, where the practices of “deceptive button colours” and “deceptive button contrast” were examined [31]. The board concludes that one practice could be manifestly misleading for users, when the decline button is such that “the contrast between the text and the button background is so minimal that the text is unreadable to virtually any user.”. The requirement for accessible contrast level between the text and the button background can be identified using W3C Web Content Accessibility Guidelines (WCAG) [70].

Accept and decline with equal contrast levels. However, during our discussions with CNIL employees, we have identified a specific case often used on the French websites, when the contrast ratio between the text and the button background is the same for accept and decline buttons, however their colors are “reversed”, thus making the accept button more visibly prominent than the decline button. We have used this idea while designing *Highlighted accept* and *Highlighted decline* banners. For illustration, Figure 1 shows the English translation of *Highlighted accept* banner designed for this experiment.

⁹We have designed the second layer of the banner following CNIL recommendations [14] - this layer has also been reviewed by two CNIL employees and can be found in Appendix B.

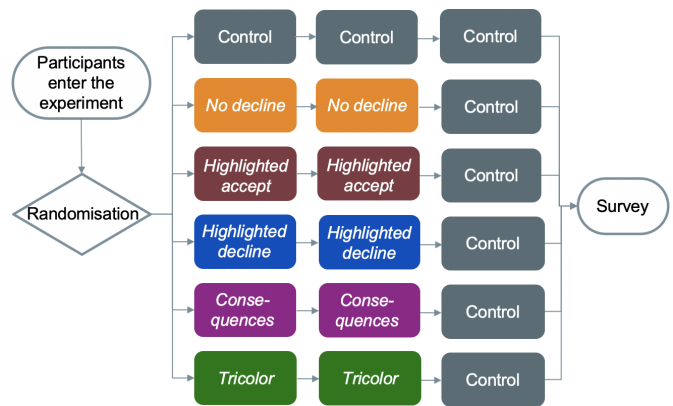


Figure 2: User journey through the online experiment.

4 User study design

To investigate the impact of the selected dark patterns and behavioral levers on user decision-making, we conducted a large-scale online between-subjects user study with the total of 3,947 participants, so that on average 657 participants have interacted with each banner (see Figure 2).

4.1 Experimental protocol

User journey through the online experiment. To provide the settings as realistic as possible, and building on previous consent banner studies that used e-commerce websites [40, 68], we asked participants to visit three fictitious e-commerce websites that sold respectively vinyl decks, a coffee machine, and a bluetooth speaker. These websites were built based on e-commerce templates, and had been tested during previous experiments. To mitigate biases linked to the simulated environment, we told participants that these were “real live websites” and asked them to find specific information on each website (e.g. the colour of the product, the battery life)¹⁰. To replicate as closely as possible the “goal-oriented” mindset of an individual navigating a website, we incentivized correct answers to the task by entering participants in a draw for a 10€ gift voucher. Upon opening the webpage, a consent banner popped up, probing a choice from them.

Figure 2 shows the user journey through the online experiment¹¹. Participants were randomly assigned to one of the six groups corresponding to the six banner variants described in Table 2. On the first two websites they visited, participants saw the banner that corresponded to their treatment group. On the third website, they all saw the *Control* banner. This last step allowed us to test whether the effect of the designs have persisted in a short-term (RQ4). Once the three websites

¹⁰Screenshots of the website and the task can be found in Supplemental Materials [5, Appendix B].

¹¹We designed the experiments so that participants would be able to read and interact with it both on desktop and mobile.

were reviewed, participants had to fill in a survey to evaluate their knowledge about cookies, their opinion on data sharing, and their satisfaction with the cookie banners they saw.

Outcomes of interest. First, we record participants' choice to accept, decline or personalize on the first treatment banner. For the purposes of this study, we create a binary variable, called *the rate of refusal/personalization*, which is equal to 1 if a user chooses to refuse or personalize cookies, and 0 otherwise. Second, we record the time spent to make a decision on the first consent banner, measured in seconds. Third, we record participants' choice to accept, decline or personalize on the third control banner. We similarly compute a rate of refusal/personalization on the third banner.

The survey (provided in Appendix A) at the end of the study allowed us to collect additional outcomes about the level of knowledge and understanding of cookies and attitudes towards data sharing online (RQ1). Participants were finally asked to revise the first consent banner they had interacted with, to express their overall sentiments towards this banner, and to say whether they would like to reconsider the choice they had initially made (which was reminded to them).

4.1.1 Ethical considerations

Though the authors' organizations do not have a formalized ethical review process like IRB, our organizations and the CNIL separately reviewed the design and did not identify any high-risk ethical issues. No other personal data than demographics was collected via the panel provider, and the three fictitious websites did not include any cookies or tracking technologies that require user consent. Via the panel provider, participants gave an informed consent for the data collected.

Participants were not informed ahead of time of the true intent of this experiment, but thought that they were taking part in a market research study, as it is common on the survey panel they have subscribed to. Indeed, the experimental protocol described above relies on collecting data on online behavior - which informing participants that we were researching the impact of consent banner design would have likely biased.

Upon entering the last section of the survey, participants were explicitly told the intention of the study and how the data provided by them would be used. Similarly, participants were only told at the end of the study that it was run in collaboration with a DPA. Participants were, upon finishing the experiment, provided with an email address they could reach out to if they had any question or concern about the study.

This study design was deemed acceptable given that the risks incurred by participants were evaluated as minor (especially as they correspond to behaviour adopted online every day), and given that participants were fully informed at the end of the study.

4.2 Participant Recruitment

Recruitment protocol. This large scale randomised controlled experiment was conducted online from September 22 to October 5, 2022 on a platform called Predictiv [55], created and managed by the Behavioural Insights Team to run online experiments. Participants were recruited via Cint, a panel supplier the co-authors of this paper use often.

Description of the sample. Participants over 18 years old and living in France were eligible to participate. To ensure representativeness, we used quotas on age, gender, and education based on official statistics published by the French National Institute of Statistics and Economic studies (INSEE) [45]. Quotas indicate the maximum percentages, not minimum percentages of participants being able to fill in the survey. Although the point of quotas is to control the proportion of specific groups that complete the survey, reaching the target quotas is not always guaranteed as it depends on the availability of participants in the marketplace. In order to keep participants who were paying attention, we included as a very first question an attention check question simply asking participants to click on two specific answers. If they didn't click on the two stated answers, they were considered as people not reading the question and therefore not paying attention.

4,026 participants passed the attention test and completed the experiment¹². However, data on demographics was missing for 79 individuals, leaving us with 3,947 participants for multivariate analysis using covariates. Supplemental Materials [5, Appendix D] provide user demographics. The experiment lasted on average 6 minutes 55 seconds and participants were remunerated \$0.68, which is close to the average remuneration for surveys on this type of marketing survey aggregators [63]¹³. They also had the opportunity to enter a draw for a 10€ gift voucher if they correctly answered questions about the websites reviewed. As described above, this was mainly used to ensure realistic interactions with the websites¹⁴.

¹²We tested attention using questions, at the start of the study, requiring individuals to read a relatively long question and then provide specific answers given in the instructions. Those failing this test were excluded. We limited non-interaction with the websites: (a) participants could only progress if they made a choice in the cookie banner; (b) for each website, participants had to respond to simple questions by searching for information within the website.

¹³Our study was run to mimic a usual marketing study, and the remuneration rate was provided by our panel aggregator. The estimate takes into account the market (what do other surveys and online experiments pay for studies of similar length), length of completion, whether the sample is hard to reach, and length of time the study runs. Similarly to remarks in other studies in the field [41, Section 3.3], and according to our panel provider, these amounts are in line with the industry standards. We recognize that this compensation is relatively lower than what was given by previous academic studies conducted on crowdsourcing platforms, which remunerated participants .18 to .33 USD/minute, compared to our .11 USD/minute [39, 40].

¹⁴Before launching the experiment, we ran a pilot study with 40 participants to make sure that reviewing three websites would not lead to too much attrition. Seeing that attrition remained at expected levels through the pilot study, we chose to maintain the experimental design as described above, thus allowing us to explore the effect of designs over time.

4.3 Analytical strategy

Power calculations. Power calculations on the primary outcome (the rate of refusal or personalization) were used to calculate the minimum detectable effect for pair-wise comparisons between banners. We used the STATA function *power twoproportions* and corrected for multiple comparisons by using the Benjamini-Hochberg method. Assuming a Type I error of $p < 0.05$, 80% power, a 14% baseline rate of refusal/personalization¹⁵, a total of 6 banner variants and a sample size of 4000 participants, the study was powered to detect effects greater than 6.43 percentage point, or a "treated" rate of refusal/personalization of 20.43%.

Data Analysis. We performed three types of statistical analysis. To test for the banners' effect on the rate of refusal/personalization on the first and third exposition, we used multivariate logistic regressions on (among other covariates) indicator variables for random treatment assignment. We used the logistic regression model shown in Eq. 1, where i is a participant; Y_i^1 is a binary variable indicating whether users have refused/personalized cookies on the first banner they saw; T_i is a vector of binary variables indicating which treatment arm participants belong to; Ψ_i is a vector of covariates including age, gender, education level and employment status.

$$Y_i^1 \sim \text{Bernoulli}(p_i); \text{logit}(p_i) = \alpha + \beta' T_i + \Gamma' \Psi_i \quad (1)$$

To test for the impact on time spent on banners, we used multivariate ordinary least square regressions of time spent on, again, treatment assignment¹⁶ using Eq. 2, where i is a participant; Y_i^1 is a continuous variable indicating how much time participants spent on the first banner they saw (in seconds); T_i is a vector of binary variables indicating which treatment arm participants belong to; Ψ_i is a vector of covariates including age, gender, education level and employment status; ε_i^1 is the error term.

$$Y_i^1 = \alpha + \beta' T_i + \Gamma' \Psi_i + \varepsilon_i^1 \quad (2)$$

To analyse survey answers, we displayed frequencies and used univariate regressions of a series of outcomes on indicators for treatment assignment, using each time the highest value as excluded category¹⁷. We used logistic regressions if survey responses were binary, and OLS if responses were continuous. In our reporting of findings, we highlight statistical differences by putting a star (*) next to values that are statistically different (with $p < 0.05$) to the highest value, which is systematically used as reference category in all regressions.

¹⁵This baseline proportion was taken from results of a study on a similar Control banner by Utz et al. [68].

¹⁶For these analyses, the variable for "time spent" was winsorized at the top (1%) to deal with outliers, which realistically resulted from participants leaving their computer and coming back to the experiment later.

¹⁷This strategy was chosen so that interpretations could always highlight the highest value and whether it significantly differs from other treatment conditions.

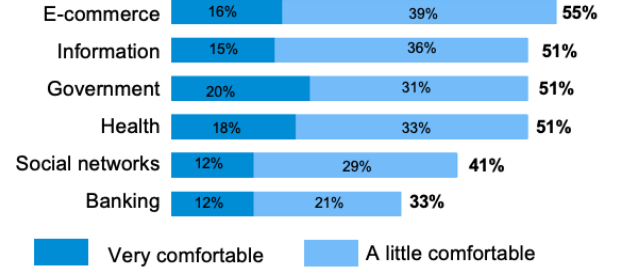


Figure 3: Comfort level of the participants with the idea of accepting all cookies, by type of website visited (N=4,026).

5 Results

5.1 Survey: users knowledge and preferences

User survey after interaction with three websites allowed us to address **RQ1** about users' knowledge about cookies and tracking technologies, and their general preferences on data sharing online¹⁸.

General views on data sharing. The results of our survey establish a *baseline for users' preferences* in our experiment: More than half of participants (52%) reported being comfortable sharing their data on the internet, even without knowing exactly what it is being used for or how it is being used and for what purpose. Conversely, less than one third (31%) say they are not comfortable sharing their data in any way. Therefore, at least 31% of internet users should refuse the use of cookies, regardless of banner design (if they understand the association between giving consent and data sharing). Designs that allow users to make a choice reflecting their preferences may encourage between 31% and 48% of users to refuse or personalize their consent. We have found that comfort with data sharing does not vary by age or education level of participants.

Types of data users accept to share. When asked about the specific data they are comfortable sharing, three-quarters (74%) of participants say they are somewhat or very comfortable sharing data about products purchased in the past (which is collected through cookies and other tracking technologies). However, only 46% say they are comfortable with sharing their browsing history (with only 15% saying they are very comfortable), a piece of data which is also collected by trackers. These results, taken together, show that users' views on data sharing are nuanced, and can depend on the type of data collected.

Impact of the type of website on comfort with accepting cookies. Kulyk et al. [49] have found that users' decisions about leaving a website are largely impacted by the trustworthiness, familiarity, and perceived importance of its contents to the user. Borberg et al. [7] found, however, that users' decision to consent to data sharing did not seem dependent on the

¹⁸Results for all survey questions can be found in the Supplemental Materials [5, Appendix E]

type of website. Our survey suggests the opposite: we find that participants report being more comfortable with accepting all cookies on e-commerce, government, health or information sites and less comfortable on social networks or banking sites (see Figure 3.).

Data collected via cookies. Although many participants in our study know that browsing history and past purchases are data collected via cookies (62% and 58% respectively), others believe that data about geolocation, social network, postal address or phone number are also collected via cookies (57%, 36%, 30% and 29% respectively).

Purposes of data collection via cookies. Similarly, participants do not fully understand the purposes for which cookies and other trackers collect data. While 61% of participants know that cookies are used to personalize ads, only 48% of participants know that the data collected can also be sold to other companies, despite this being common practice today. Even fewer participants are aware that the data collected can also be used to adapt prices and promotions presented (29%), information that could nonetheless factor into consumers' decisions. One third of participants (36%) think that the use of cookies can have beneficial functions on their user experience and only 17% know that cookies can be used to enforce the security of websites. Finally, while 74% know that cookies track what users do online, misconceptions persist – 15% of participants believe that cookies can collect data while listening to what users are saying.

5.2 Impact of design on consent decisions

We analyze the decisions of 3,947 users upon interacting with our six consent banners to address research questions **RQ2** and **RQ3** on the impact of banner design on users' decisions. We use data on time spent on banners, as well as user sentiment reported in the survey from Section 5.1, to enrich our interpretation of the data¹⁹.

Figure 4 shows participants' consent decisions upon interacting with the consent banner on the first website they were asked to visit, by treatment condition (or in other words, by consent banner seen)²⁰. Faced with a banner that presents options in a neutral way (*Control*), 17% of participants refuse or personalize the use of cookies. This is quite far from the expected 31% of users who said they were uncomfortable sharing their data under any circumstances (see Section 5.1), and reinforces the consensus view that recommending neutral banners might not be sufficient²¹.

¹⁹The results for all regression variables can be found in Supplemental Materials [5, Appendix F]

²⁰The figures above the bars represent the sum of the refusal and personalization rates. These do not correspond precisely to the sum of the figures displayed in the bars of the graph due to rounding.

²¹We also looked at these main results separately for mobile/desktop. Supplemental Materials [5, Appendix G] shows that rejection rates are overall lower on mobile, which is in line with previous literature (see for example [68]). The effects of the banners are, however, qualitatively the same across

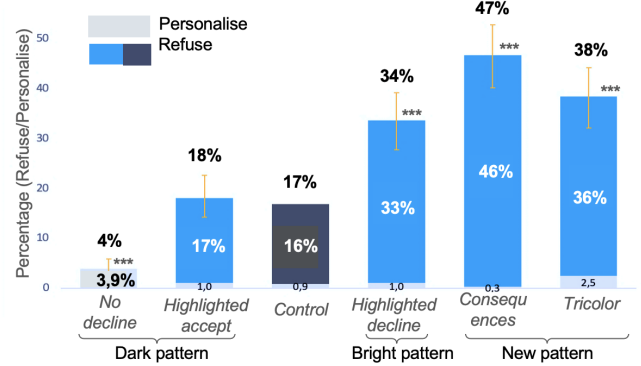


Figure 4: Participants' reported consent decisions (N=3,947) in their interactions with the consent banners during the visit to the first website: three stars (***) indicate $p < 0.001$, and orange bars correspond to the 95% confidence interval. Primary analysis using logistic regression, controlling for age, gender, education, and employment status of participants.

5.2.1 Impact of dark patterns on consent decisions

No decline banner. The refusal/personalization rate drops from 17% on a neutral banner to 4% ([95% CI, 2.5%-5.9%], $p = 0.000$) when users need to visit the second layer of the banner to decline consent. This statistically significant difference contradicts the previous results of Grassl et al. [39] that found no impact of such design on users' decisions. Moreover, our result complements the initial result of Nouwens et al. [54], who evaluated it with 40 US students and found that refusal rate dropped by 50% between the two banners (see § 3.1.1). With 657 French participants per banner, our study provides a new robust insight to the research community that was not demonstrated at large-scale in the past. This result highlights the potential detrimental effects of this design. Our result provides support for ongoing efforts by regulatory authorities to combat the use of *No decline* banner which make it harder for users to act in accordance with their preferences.

Highlighted accept banner. This banner, that visually emphasizes acceptance, has no statistically significant effect on users' consent decision compared to a control banner (18% [95% CI, 14.2%-22.6%], $p = 0.566$), confirming the results of Grassl et al. [39] and Berens et al. [3], but this time in a large-scale study and with French participants (see Section 3.1.2). This lack of difference between the control banner and the *Highlighted accept* banner suggests that individuals may have already developed the habit of accepting: highlighting or not highlighting this option does very little to sway their already habituated behaviour as their mental shortcuts are already in place. Indeed, when we asked participants why they accepted the use of cookies, 50% said they did so out of habit.

both interfaces.

5.2.2 Impact of behavioral levers on consent decisions

Highlighted decline banner The bright pattern *Highlighted decline* banner has a substantial effect on the outcome of the consent decision (34% [95% CI, 27.7%-39.2%], $p = 0.000$), doubling the refusal/personalization rate compared to the control banner (from 17% in control banner to 34% in highlighted decline). This new result is opposite to the previous findings of Grassl et al. [39] and Berens et al. [3], who have found no support for the hypothesis that such design impacts users' decisions (see Section 3.2.2). This result could be explained by users' tendency to accept the pre-selected option, which requires less cognitive effort and which users tend to think has been selected for them according to their interests [46]. It is also possible that highlighting rejection could create shock, or dissonance among users who are used to seeing acceptance highlighted as the recommended choice. This dissonance could then encourage them to question the habit of accepting and to rethink their choice.

Consequences banner. This banner, that highlights the negative consequences of data sharing through strong evocative words ("tracking"), has the biggest effect on the outcome of the consent decision (47% [95% CI, 40.2%-52.9%], $p = 0.000$) and increases the refusal rate by as much as three times (from 16% for neutral banner to 46% for consequences banner). In contrast to the findings of Habib et al. [40] and Berens et al. [3], who found no statistically significant impact of button text labels on user decisions (see Section 3.2.1), our finding demonstrates a clear impact of button text labels. This result supports the findings of Ma and Birrell [51], who found that the negative slant (similar to our text labels)²² have a significant impact on users compared to the positive slant. However Ma and Birrell didn't compare it to the neutral labels, and the sample of users who interacted with the banners was very small (up to 67 users per banner) while our experiment contained 657 participants per banner.

When surveyed, 71% of participants who saw this banner said they declined consent "in order not to be tracked", which was the highest number of participants to decline consent for this reason amongst all groups. This suggests that the negative consequences of tracking were more prominent in their minds when making the choice, thus re-balancing the importance that users place on future risks, which are sometimes overshadowed by a desire for speed.

Tricolor banner. Similarly, the *Tricolor banner* also sought to highlight the negative impacts of data sharing, by using familiar visual metaphors (traffic lights, shield icons). The use of these metaphors also had a large and significant effect (38% [95% CI, 32.1%-44.2%], $p = 0.000$): those who saw this banner were 2.3 times more likely to click refuse / personalize

²²The authors used a negative slant attempt to nudge users: "[such banners] present user with a choice between accepting cookies – which poses privacy risks by allowing a website to access and sell personal information – and denying cookies to prevent this harm". Our text labels highlighting the acceptance or refusal of being tracked are similar to such slants.

than those who saw the control banner. Such banners have never been tested in the previous works and our result provides a new insight to the research community (see Section 3.2.3). **Summary.** We conclude that all the three behavioral levers introduced in the banners have substantial effect on the outcome of consent decision. The refusal/personalization rates are close to what we would have seen if participants had made choices aligned with their initially reported preferences (see Section 5.1): between 31% and 48% of participants said they wanted to refuse to share their data or do so depending on the how exactly their data will be used – these numbers that are close to the 34% to 47% of users choosing to refuse or personalize after seeing the banners with behavioral levers.

5.3 Time spent and user satisfaction

We have analyzed the time spent on the banners, as well as user preferences over banners by measuring the level of satisfaction users report with their choices, the level of understanding and simplicity of navigation on the banner, the feeling of protection, and the general satisfaction with the banner (Appendix A). We used these criteria to evaluate how design banners would be perceived by users, and how they would make them feel as they entered a website. Table 3 shows the time spent to make decisions as well as users' satisfaction with the banner they interacted with and the choice they made²³.

Time spent interacting with the banner. Participants spent an average of 4.1 seconds on a *Control* banner, consistent with prior research [68]. Duration remained steady with dark patterns, but for *Consequences* and *Tricolor* users took more time, possibly due to surprise or reflection. "Time spent" on banners can serve as a proxy for effort and thinking in decision-making contexts, as longer duration often indicate increased cognitive engagement, complex decision-making, and information processing. The strongest dark pattern significantly prolonged consent decline time (up to 17 seconds), highlighting the importance of placing both options on the same layer.

User sentiment towards consent banner patterns. According to survey responses, participants who saw the *Consequences* and *Tricolor* banners were more likely to be satisfied with the banner they saw and with their choice, potentially illustrating a "wish for privacy" granted (see Table 3). Conversely, those who saw the *No reject* dark patterns banner were less likely to be satisfied, confirming that the more pernicious dark patterns do not allow users to make a choice that reflects their data protection preferences.

Finally, although the two new pattern banners, *Consequences* and *Tricolor*, appear to be the most successful at

²³One star indicates that this value is significantly different from the highest value at least to the 5% level. In the first row, the control banner and the two dark pattern banners are the banners that have a significantly different rate of refusal/personalization from the banner "Consequence". As these analyses were done with univariate regression, all 4,026 participants are included.

Table 3: Participants’ reported time spent and overall satisfaction with the banner they interacted with and the choice they made.

	<i>Control</i>	<i>No decline</i>	<i>Highlighted accept</i>	<i>Highlighted decline</i>	<i>Consequences</i>	<i>Tricolor</i>
Time spend on first banner (seconds)	4.1*	4.0*	3.8*	4.2*	5.7	5.7
Simplicity level (ease of navigation and understanding)	81%*	79%*	84%	86%	86%	86%
Satisfaction with the banner (wish to see more banners like this one)	78%*	72%*	82%	81%	82%	84%
Satisfied with the choice	54%*	44%*	55%*	58%*	64%	64%

impacting refusal/personalization rates, those who saw the banner emphasizing being tracked (*Consequences*) were less likely to think that the choices were presented with their interests in mind. This could ultimately affect the trust in the information provided, or lead individuals to avoid this information altogether if it is perceived as too "negative", thus making this banner design less desirable overall [34].

User satisfaction with the choice made. After reminding participants of the choice they had made when faced with the first consent banner, we asked them to reflect whether they were satisfied with their choice. Our analyses reveal that people who declined tended to be more satisfied by their choice: 78% of those who declined were satisfied, while only 48% of those who had accepted were satisfied. When faced with a neutral banner, only 54% of participants said they were satisfied with their choice. The proportion of satisfied users decreases by 10 points for the most extreme dark pattern *No reject* banner (44%), confirming the need to discourage this type of banner. Conversely, this proportion increases by 10 points (64%) for the banners where the refusal/personalization rate was the highest (*Consequences* and *Tricolor*).

In addition to being less satisfied on average, more individuals who saw the control banner or those with dark patterns also indicated that they did not care about their choice, potentially suggesting a sense of resignation, which has already been highlighted in qualitative studies [7]. This resignation could suggest decision fatigue, whereby users would prefer not to question their habits and may tend to avoid the information presented. It might therefore be important to encourage more unusual banners, containing motifs that encourage reflection.

5.4 The short-term time effect of design

We explored whether dark patterns and behavioral levers have a short-term effect on future user consent decisions (RQ4). To test this hypothesis, we asked participants to interact with three different consent banners on three websites (Figure 2). This protocol allowed us to explore, by studying the choice made on the third control banner, whether the short-term effect of the dark patterns or behavioral levers persists.

To select the number of exposures seen before measuring the refusal/personalization rate, we balanced the need to

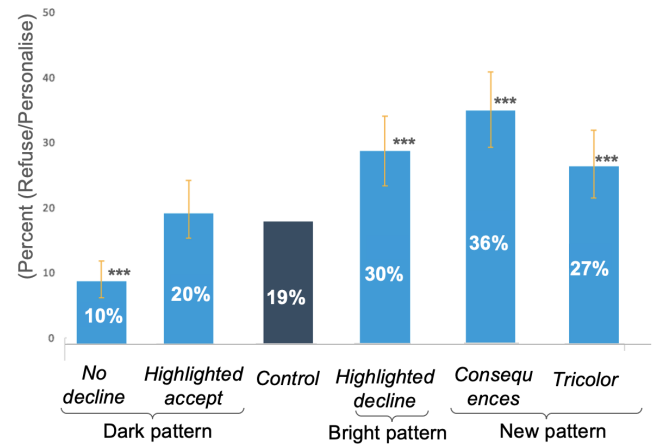


Figure 5: Participants’ reported consent decisions (N=3,947) on upon visiting the third website with *Control* banner.

measure several repeated choices, with the risk of creating fatigue in participants and increasing the dropout rate of the experiment. Previous work suggests that one to two exposures to a dialogue window are sufficient to create a habitual response [9]. We therefore opted for two exposures, before testing the survival of the effect on a third one. The frequency of exposure to consent banners when surfing the internet, however, means this type of choice setup can also realistically be found in the real world. The effects measured here will, nevertheless, likely remain measures of the upper-bound of survival effects given how shortly after they are measured.

Figure 5 shows the refusal/personalization rates upon third exposure to the *Control* banner²⁴. Unsurprisingly, the refusal/customization rate on the *Control* banner does not vary between the first and third exposure (17% for the first and 19% for the third) and this rate also does not vary on the *Highlighted accept* banner. After seeing two *No reject* banners, the rate of refusal/personalization is significantly lower than that of participants who saw *Control* banners only: in this group only 10% ([95% CI, 7.2%-13.8%], $p=0.000$) of participants refused/personalized. This suggests that the short-term effect

²⁴These are the results of a logistic regression, controlling for age, gender, level of education and employment status of individuals.

of *No reject* dark pattern can survive, likely because the repeated exposure has created a habit around always accepting all cookies, reinforced by the fact that accepting all has always allowed access to the content wanted.

We have observed a survival effect after two consequent exposures to the banners with behavioral levers. Among participants who first saw two *Highlighted decline* banners, rate of refusal/personalization was significantly higher (30% [95% CI, 24.3%-35.0%], $p=0.000$), compared to participants who only saw *Control* banners. This was also the case among participants who first saw the *Consequences* banners (36% [95% CI, 30.1%-41.8%], $p=0.000$) or *Tricolor* banners (27% [95% CI, 22.5%-32.9%], $p=0.000$). Taken together, these results provide the first quantitative evidence that banners designs significantly affect users choices in a short-term.

6 Discussion

6.1 Limitations

Our study has a few limitations. Our participants were compensated to perform a controlled experiment which, despite our best efforts, remains distinct to the usual setting in which internet users are confronted with consent banners. Indeed, the usual exposure to consent banners comes in a setting where participants have goal-oriented behaviors (looking for an information, trying to purchase a product etc.) and where attentional limitations might modulate the effects we observe.

Another noteworthy aspect is that our experiment involved interactions with only three specific e-commerce websites. While this approach enables controlled testing, it may limit the external validity of our findings. While using e-commerce prototype websites is commonly used in research literature on consent banners [8,40,68], it doesn't capture the full spectrum of online user experiences. Nevertheless, our survey shows that users' willingness to accept cookies on an e-commerce website (55%, see Figure 3) is similar to the overall attitude to data sharing (52%, Section 5.1) thus demonstrating that e-commerce websites are suitable, in the context of this experiment, for the internal validity of the study.

Despite our efforts to try and replicate the conditions of real online behaviors, our results might not generalize to all types of banners or contexts. Moreover, participants in a controlled study may be inclined to answer what they think the researchers want to hear (participant's bias), especially on sentiment and satisfaction questions. We therefore potentially overestimate true satisfaction and preferences. This, however, does not affect the internal validity of our results.

We have also compared various results with prior works on the impact of design on consent decisions (Section 5.2). Few works with greater external validity covered *Nodcline* and *Highlighted accept*, however often on a very small user base [54] and outside of France [68], while recent research shows that user choices differ across countries [8,41].

6.2 Insights on knowledge and preferences

Impact of the type of websites. The effect of website type on consent decisions is relatively understudied [7]. Our findings show that data sharing preferences vary depending on the sites visited (Section 5.1), and thus potentially on the purposes for which the collected data will be used. Future studies should explore how comfort levels vary across site types to inform regulatory decisions and guidelines.

Given widely spread misconceptions about cookies and tracking technologies, it is important that, at a minimum, consent banners provide clear information to users. However, the nature, design and frequency of consent banners limit the ability of individuals to engage in a thoughtful consideration of what they are or are not willing to share at the time the question is asked. This raises the question of the type of message that could be conveyed, or the type of awareness-raising action that could be carried out beyond simply changing the design of the cookie banners.

Support for alternative solutions to cookie banners. There exist browser extensions or solutions built-in Web browsers that can ask users to set their preferences once, and then automatically refuse or accept consent based on those preferences. For example, recent Global Privacy Control (GPC) [71] or its extension Advanced Data Protection Control (ADPC) could offer such potential solutions (even though they still face a number of challenges) [42]. Asking users to make only one choice per browser could, indeed, allow the time to better inform and encourage users to pay attention to their consent choices. When asked whether they would be interested in solutions like these, a majority of respondents to our survey (62%) said they would be willing to use this type of solution, and only 9% said they would definitely not be interested.

6.3 Design implications

Impact of dark patterns on consent decisions Our results for *No decline* banner (Section 5.2.1) confirm the impact of users' preference for speed and for paths with the least amount of friction. It's therefore essential for regulators across the EU to discourage and sanction designs that require more effort to refuse than to accept, following the example CNIL's recommendation [14] that is now supported with quantifiable result thanks to our findings. We, however, did not observe a significant effect of *Highlighted accept*, likely due to already existing habits of user to accept. We therefore believe it is important to reflect on the research methods that is based using the *Control* banner as baseline. Similarly to previous research on online decisions, we found that users express different preferences when asked, or when they interact with *Control* banners: 31% of users do not agree with data sharing online when asked in the survey (Section 5.1), while only 17% decline when facing a neutral banner interface (Section 5.2), that could have been caused by the long-term effect of banner

design. We therefore recommend to evaluate the design of consent banners not only against users' decisions in *Control* banners (that are likely impacted by the banners seen in the past) but also against users' true privacy preferences.

Need to further quantify the long-term effect of design patterns. We found that users' choices in consent banners are impacted by the designs of banners they interacted with in the past, even over a short time-frame (Section 5.4). Given that users have been continuously exposed to dark patterns in consent banners manipulating them towards acceptance in the last years [53, 54, 62], our finding can explain the "stickiness" of accept as an option. Our result opens many windows of opportunities for future research to investigate and quantify the habituation effects of design patterns in the long-term. Regulatory authorities could also benefit from our result: in the same way than bad habits can be reinforced, that good habits could be introduced, only by encouraging repeated exposure to "better design" that can be encouraged by the regulators.

Need for legally-required built-in solutions. Only half of the participants (54%) in our survey say they are satisfied with their choices after having seen a control banner. This satisfaction rate is even lower after having seen banners containing dark patterns (see Table 3). Banners containing behavioral levers generate a higher number of satisfied users. Nevertheless, it is important to note that almost a third of the participants said that, even after having seen the consent banners with behavioral levers, they were dissatisfied or not very concerned by the choice they made. This suggests a need to continue working on solutions that allow for a more informed choice or require less effort on a daily basis, such as browser built-in solutions. In the US, the California Consumer Protection Act (CCPA) requires online businesses to offer consumers methods to submit requests to opt-out of the sale of their personal information, that could be implemented in the user's browser via Global Privacy Control(GPC) [35]. In the EU, the update of ePrivacy directive [29] into a regulation could provide a similar legal requirement in the future, however this new law has not been finalised [21].

Recommendations and call for action. A strategy to amend the situation where users do not select their desired data sharing option in the neutral banner interface could include various actors in the ecosystem. *Regulators* could target most popular platforms in their enforcement actions: the more popular websites present explicit decline button, the faster the effect of "stickiness" of accept will weaken. Additionally, *developers and practitioners* of companies, who put privacy and ethics at the core of their business values, could be motivated to provide best practices and exemplary banners. For example, the French TarteAuCitron banner provider [11], which is used by many French administration websites whose websites are not monetized by ads, could be motivated to provide them a banner with behavioral lever design, where users consent choices are similar to their willingness to share data according

to our results in Section 5.2.2. Nevertheless, as highlighted in previous works [36, 57], various actors in this ecosystem often have conflicting goals, making it difficult to motivate all practitioners. Finally, we encourage *researchers* to collaborate more closely with the regulators and practitioners to study their motivations and interests in improving consent banner design, and understand which other types of designs can be important to study in the future work.

Acknowledgments

This work has been partially supported by the ANR JCJC project PrivaWeb (ANR-18-CE39-0008), by the ANR 22-PECY-0002 IPOPOP project of the Cybersecurity PEPR, and by the Interministerial directorate for public transformation (DITP) as part of a framework agreement on which the Behavioural Insights Team was a supplier (tender no 001-2018-DITP).

7 Conclusions

At the request of the French DPA, we conducted a 3,947-participant online experiment in France exploring the impact of six distinct consent banner designs on the users' consent decisions. We also measured participants' knowledge of cookies and privacy preferences, as well as satisfaction with the banners included in the experiment. In line with previous studies, we confirm that dark patterns that make rejecting cookies unrealistically hard strongly impact consent and satisfaction, and that dark patterns that highlight the accept option do not impact choices compared to a neutral banner. This suggests that users might have developed a habit of accepting, despite this not being in line with their true privacy preferences. Contrary to previous results, we however find that a bright pattern that highlights the decline option has a substantial effect on users decisions (up to doubling the rate of refusal). We also tested two new designs based on insights from behavioral sciences. The banners including behavioral levers have the strongest effect on consent decisions. Moreover, participants were the most satisfied when interacting with these banners. Finally, we provide the first evidence, to our knowledge, that the effect of banners designs persists over a short time frame, and can significantly affect future consent choices, even when faced with neutral banners.

Taken together, these results strongly reinforce the need for regulation around the most extreme dark patterns (such as the ones where the option to reject is on a second page), but also offer some hope : by including them in their guidelines and encouraging exposure to bright patterns or to banners including behavioral levers, data protection authorities have an opportunity to break the habit to systematically accept cookies, and offer a realistic path for users to express their true privacy preferences.

References

- [1] Alessandro Acquisti. Privacy in electronic commerce and the economics of immediate gratification. In *Proceedings of the 5th ACM Conference on Electronic Commerce*, EC '04, page 21–29, New York, NY, USA, 2004. Association for Computing Machinery.
- [2] Jan Michael Bauer, Regitze Bergstrøm, and Rune Foss-Madsen. Are you sure, you want a cookie?: The effects of choice architecture on users' decisions about sharing private online data. *Computers in Human Behavior*, 120, July 2021.
- [3] Benjamin Maximilian Berens, Heike Dietmann, Chiara Krisam, Oksana Kulyk, and Melanie Volkamer. Cookie disclaimers: Impact of design and users' attitude. In *Proceedings of the 17th International Conference on Availability, Reliability and Security*, ARES '22, New York, NY, USA, 2022. Association for Computing Machinery.
- [4] Carlos Bermejo Fernandez, Dimitris Chatzopoulos, Dimitrios Papadopoulos, and Pan Hui. This website uses nudging: Mturk workers' behaviour on cookie consent notices. *Proc. ACM Hum.-Comput. Interact.*, 5(CSCW2), oct 2021.
- [5] Nataliia Bielova, Mariam Chammat, Vincent Toubiana, Estelle Hary, Anyisia Nguyen, and Laura Litvine. The effect of design patterns on (present and future) cookie consent decisions: Supplemental materials. Technical report, 2023. <https://inria.hal.science/hal-04235032>.
- [6] Nataliia Bielova, Cristiana Santos, and Colin M. Gray. Two worlds apart! closing the gap between regulating eu consent and user studies. *Harvard Journal of Law and Technology (JOLT)*, 37, 2024. Accepted for publication.
- [7] Ida Borberg, René Hougaard, Willard Rafnsson, and Oksana Kulyk. "so i sold my soul": Effects of dark patterns in cookie notices on end-user behavior and perceptions. In *Usable Security and Privacy (USEC) Symposium*, 2022.
- [8] Elijah Bouma-Sims, Megan Li, Yanzi Lin, Adia Sakura-Lemessy, Alexandra Nisenoff, Ellie Young, Eleanor Birrell, Lorrie Faith Cranor, and Hana Habib. A us-uk usability evaluation of consent management platform cookie consent interface design on desktop and mobile. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, CHI'23. Association for Computing Machinery, 2023.
- [9] Cristian Bravo-Lillo, Lorrie Faith Cranor, Saranga Komanduri, Stuart E. Schechter, and Manya Sleeper. Harder to ignore? revisiting pop-up fatigue and approaches to prevent it. In Lorrie Faith Cranor, Lujo Bauer, and Robert Biddle, editors, *Tenth Symposium on Usable Privacy and Security, SOUPS 2014, Menlo Park, CA, USA, July 9-11, 2014*, pages 105–111. USENIX Association, 2014.
- [10] Harry Brignull. Dark patterns: User interfaces designed to trick people, 2018. <http://darkpatterns.org/>.
- [11] Amauri Champeaux. [tarteaucitron.io](https://tarteaucitron.io/en/). <https://tarteaucitron.io/en/>, accessed on 5 September 2023.
- [12] Hichang Cho, Jae-Shin Lee, and Siyoung Chung. Optimistic bias about online privacy risks: Testing the moderating effects of perceived controllability and prior experience. *Computers in Human Behavior*, 26(5):987–995, 2010. Advancing Educational Research on Computer-supported Collaborative Learning (CSCL) through the use of gStudy CSCL Tools.
- [13] CNIL. Lignes directrices « cookies et autres traceurs », 2020. <https://www.cnil.fr/en/decisions/lignes-directrices-recommandations-CNIL>.
- [14] CNIL. Recommandation « cookies et autres traceurs », 2020. <https://www.cnil.fr/sites/default/files/atoms/files/recommandation-cookies-et-autres-traceurs.pdf>.
- [15] CNIL. 2021, a record year for the CNIL's enforcement action, 2021. <https://www.cnil.fr/en/2021-record-year-cnils-enforcement-action>.
- [16] CNIL. Cookies: Facebook ireland limited fined 60 million euros, 2022. <https://www.cnil.fr/en/cookies-facebook-ireland-limited-fined-60-million-euros>.
- [17] CNIL. Cookies: Google fined 150 million euros, 2022. <https://www.cnil.fr/en/cookies-google-fined-150-million-euros>.
- [18] CNIL. Cookies: Microsoft ireland operations limited fined 60 million euros, 2023. <https://www.cnil.fr/en/cookies-microsoft-ireland-operations-limited-fined-60-million-euros>.
- [19] CNIL. Cookies: the CNIL fines Tiktok 5 million euros, 2023. <https://www.cnil.fr/en/cookies-cnil-fines-tiktok-5-million-euros>.
- [20] CNIL. Sanctions et mesures correctrices : la CNIL présente le bilan 2022 de son action répressive, 2023. <https://www.cnil.fr/fr/sanctions-et-mesures-correctrices-la-cnil-presente-le-bilan-2022-de-son-action-repressive>.
- [21] European Commission. Proposal for an ePrivacy Regulation. <https://digital-strategy.ec.europa.eu/en/policies/eprivacy-regulation>, visited on 5 September 2023.
- [22] European Commission. *Behavioural study on unfair commercial practices in the digital environment : dark patterns and manipulative personalisation : final report*. Publications Office of the European Union, May 2022.
- [23] Lynne M. Coventry, Debora Jeske, John M. Blythe, James Turland, and Pam Briggs. Personality and social framing in privacy decision-making: A study on cookie acceptance. *Frontiers in Psychology*, 7, 2016.
- [24] Proposal for a Regulation of the European Parliament and of the Council on harmonised rules on fair access to and use of data (Data Act), 2022.
- [25] Digital Markets Act - Regulation (EU) 2022/1925 of the European Parliament and of the Council of 14 september 2022 on contestable and fair markets in the digital sector and amending Directives (EU) 2019/1937 and (EU) 2020/1828 (digital markets act) (text with eea relevance), 2022.
- [26] Regulation (EU) 2022/2065 of the European Parliament and of the Council of 19 October 2022 on a Single Market For Digital Services and amending Directive 2000/31/EC (Digital Services Act), October 2022.
- [27] Guidelines 03/2022 on deceptive design patterns in social media platform interfaces: how to recognise and avoid them.
- [28] Serge Egelman, Andreas Sotirakopoulos, Ildar Muslukhov, Konstantin Beznosov, and Cormac Herley. Does my password go up to eleven? the impact of password meters on password selection. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '13, page 2379–2388, New York, NY, USA, 2013. Association for Computing Machinery.
- [29] Directive 2009/136/EC of the European Parliament and of the Council of 25 November 2009. <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=celex%3A32009L0136>, accessed on 2019.10.31.
- [30] European Data Protection Board. Guidelines 05/2020 on consent under Regulation 2016/679. https://edpb.europa.eu/sites/edpb/files/files/file1/edpb_guidelines_202005_consent_en.pdf.
- [31] European Data Protection Board. Report of the work undertaken by the cookie banner taskforce. https://edpb.europa.eu/our-work-tools/our-documents/report/report-work-undertaken-cookie-banner-taskforce_en.
- [32] European Data Protection Board. Guidelines 3/2022 on dark patterns in social media platform interfaces: How to recognise and avoid them. Technical Report Version 2.0, February 2023. https://edpb.europa.eu/our-work-tools/our-documents/guidelines/guidelines-032022-deceptive-design-patterns-social-media_en.

- [33] Julia Giese and Martin Stabauer. Factors that influence cookie acceptance: Characteristics of cookie notices that users perceive to affect their decisions. In *HCI in Business, Government and Organizations: 9th International Conference, HCIBGO 2022, Held as Part of the 24th HCI International Conference, HCII 2022, Virtual Event, June 26 – July 1, 2022, Proceedings*, page 272–285, Berlin, Heidelberg, 2022. Springer-Verlag.
- [34] Russell Golman, David Hagmann, and George Loewenstein. Information avoidance. *Journal of Economic Literature*, 55(1):96–135, March 2017.
- [35] Global Privacy Control (GPC) specification, 2023.
- [36] Colin Gray, Cristiana Santos, Nataliia Bielova, Michael Toth, and Damien Clifford. Dark patterns and the legal requirements of consent banners: An interaction criticism perspective. In *ACM CHI Conference on Human Factors in Computing Systems (ACM CHI)*, 2021. <https://arxiv.org/abs/2009.10194>.
- [37] Colin M Gray, Yubo Kou, Bryan Battles, Joseph Hoggatt, and Austin L Toombs. The Dark (Patterns) Side of UX Design. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, CHI '18, pages 534:1–534:14, New York, NY, USA, 2018. dl.acm.org.
- [38] Colin M Gray, Cristiana Santos, and Nataliia Bielova. Towards a preliminary ontology of dark patterns knowledge. In *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems (CHI EA '23)*, 2023.
- [39] Paul Graßl, Hanna Schraffenberger, Frederik Zuiderveen Borgesius, and Moniek Buijzen. Dark and bright patterns in cookie consent requests. *Journal of Digital Social Research*, 3(1):1–38, Feb. 2021.
- [40] Hana Habib, Megan Li, Ellie Young, and Lorrie Cranor. “Okay, whatever”: An evaluation of cookie consent interfaces. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, CHI '22, New York, NY, USA, 2022. Association for Computing Machinery.
- [41] Franziska Herbert, Steffen Becker, Leonie Schaewitz, Jonas Hielscher, Marvin Kowalewski, Angela Sasse, Yasemin Acar, and Markus Dürmuth. A world full of privacy and security (mis)conceptions? findings of a representative survey in 12 countries. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, CHI '23, New York, NY, USA, 2023. Association for Computing Machinery.
- [42] Soheil Human, Harshvardhan J. Pandit, Victor Morel, Cristiana Santos, Martin Degeling, Arianna Rossi, Wilhelmina Botes, Vitor Jesus, and Irene Kamara. Data protection and consenting communication mechanisms: Current open proposals and challenges. In *2022 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*, pages 231–239, 2022.
- [43] Dennis Hummel and Alexander Maedche. How effective is nudging? a quantitative review on the effect sizes and limits of empirical nudging studies. *Journal of Behavioral and Experimental Economics*, 80:47–58, 2019.
- [44] Les Français et la réglementation en matière de cookies, 2019. <https://www.ifop.com/publication/les-francais-et-la-reglementation-en-matiere-de-cookies/>.
- [45] The national institute of statistics and economic studies collects, analyses and disseminates information on the french economy and society. <https://www.insee.fr/en/accueil>.
- [46] Jon M. Jachimowicz, Shannon Duncan, Elke U. Weber, and Eric J. Johnson. When and why defaults influence decisions: a meta-analysis of default effects. *Behavioural Public Policy*, 3(2):159–186, 2019.
- [47] Daniel Kahneman, Jack L. Knetsch, and Richard H. Thaler. Anomalies: The endowment effect, loss aversion, and status quo bias. *Journal of Economic Perspectives*, 5(1):193–206, March 1991.
- [48] Oksana Kulyk, Nina Gerber, Annika Hilt, and Melanie Volkamer. Has the GDPR hype affected users’ reaction to cookie disclaimers? *Journal of Cybersecurity*, 2020.
- [49] Oksana Kulyk, Annika Hilt, Nina Gerber, and Melanie Volkamer. “This Website Uses Cookies”: Users’ perceptions and reactions to the cookie disclaimer. In *3rd European Workshop on Usable Security (EuroUSEC)*, London, England, April 23, 2018. Internet Societa, 2018.
- [50] Laboratoire de l’Innovation Numerique de la CNIL. A survey of user studies as evidence for dark patterns in consent banners, 2023. https://backoffice.cnil.fr/sites/default/files/atoms/files/full_2022-12-02_v2.pdf, accessed on 7 September 2023.
- [51] Eryn Ma and Eleanor Birrell. Prospective consent: The effect of framing on cookie consent decisions. In *Extended Abstracts of the 2022 CHI Conference on Human Factors in Computing Systems*, CHI EA '22, New York, NY, USA, 2022. Association for Computing Machinery.
- [52] Dominique Machuletz and Rainer Böhme. Multiple purposes, multiple problems: A user study of consent dialogs after GDPR. *Proceedings on Privacy Enhancing Technologies*, 2020(2):481–498, 2020.
- [53] 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 *IEEE Symposium on Security and Privacy (IEEE S&P)*, 2020. <https://hal.inria.fr/hal-03117294>.
- [54] 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 *CHI*, 2020.
- [55] Predictiv platform. <https://www.bi.team/bi-ventures/predictiv/>.
- [56] Razorfish. Les Français face aux cookies, une étude signée Razorfish France, 2022. <https://france.publicisgroupe.com/download?n=CP%20Les%20Francais%20face%20aux%20cookies%20copie-pdf&picid=16019>.
- [57] Cristiana Santos, Midas Nouwens, Michael Toth, Nataliia Bielova, and Vincent Roca. Consent Management Platforms under the GDPR: processors and/or controllers? In *Annual Privacy Forum (APF'21)*, 2021. <https://hal.inria.fr/hal-03169436>.
- [58] Cristiana Santos, Arianna Rossi, Lorena Sanchez Chamorro, Kerstin Bongard-Blanchy, and Ruba Abu-Salma. Cookie banners, what’s the purpose? analyzing cookie banner text through a legal lens. In *Proceedings of the 20th Workshop on Privacy in the Electronic Society*, WPES '21, page 187–194, New York, NY, USA, 2021. Association for Computing Machinery.
- [59] Daniel J Simons and Christopher F. Chabris. *The Invisible Gorilla: And Other Ways Our Intuitions Deceive Us*. Crown Publishing Group, May 2010.
- [60] Ashutosh Kumar Singh, Nisarg Upadhyaya, Arka Seth, Xuehui Hu, Nishanth Sastry, and Mainack Mondal. What cookie consent notices do users prefer: A study in the wild. In *Proceedings of the 2022 European Symposium on Usable Security*, EuroUSEC '22, page 28–39, New York, NY, USA, 2022. Association for Computing Machinery.
- [61] Edith G. Smit, Guda Van Noort, and Hilde A.M. Voorveld. Understanding online behavioural advertising. *Comput. Hum. Behav.*, 32(C):15–22, mar 2014.
- [62] Than Htut Soe, Oda Elise Nordberg, Frode Guribye, and Marija Slavkovic. Circumvention by design - dark patterns in cookie consent for online news outlets. In *Proceedings of the 11th Nordic Conference on Human-Computer Interaction: Shaping Experiences, Shaping Society*, NordiCHI '20, New York, NY, USA, 2020. Association for Computing Machinery.
- [63] Surveys for money: What you can expect to earn. <https://www.nerdwallet.com/article/finance/surveys-for-money-site-reviews>, year = 2023,.
- [64] Behavioural Insights Team. Sciences comportementales appliquées : mieux protéger le consommateur en ligne, 2021.

- [65] Richard H. Thaler, Cass R. Sunstein, and John P. Balz. *Chapter 25. Choice Architecture*, pages 428–439. Princeton University Press, Princeton, 2013.
- [66] The European Parliament and the Council of the European Union. Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation), 2016.
- [67] Amos Tversky and Daniel Kahneman. Judgment under uncertainty: Heuristics and biases. *Science*, 185(4157):1124–1131, 1974.
- [68] Christine Utz, Martin Degeling, Sascha Fahl, Florian Schaub, and Thorsten Holz. (Un)informed Consent: Studying GDPR Consent Notices in the Field. In *Conference on Computer and Communications Security*, 2019.
- [69] René van Bavel and Nuria Rodríguez-Priego. Testing the effect of the cookie banners on behaviour. JRC Research Reports JRC103997, Joint Research Centre (Seville site), December 2016.
- [70] W3C. Web Content Accessibility Guidelines (WCAG) 2.1, 2018. <https://www.w3.org/TR/WCAG21/>.
- [71] Sebastian Zimmeck and Kuba Alicki. Standardizing and implementing do not sell. In *Proceedings of the 19th Workshop on Privacy in the Electronic Society*, pages 15–20, 2020.

A Survey questions

We designed our survey in French to be easily accessible for French speakers. Here we present English translation of the survey, made by the authors. For each question, the possible responses are shown in *italics*. The results of all survey questions are presented in the Supplemental Materials [5, Appendix E].

Introduction to the subject of cookies These retail sites also value the privacy of their users and seek to create cookie banners that comply with GDPR requirements and allow users to make informed choices about cookie banners. We will therefore ask you questions about this in the rest of the survey.

Q1 Please read the three descriptions below. Select the one that best describes what you think.

* *I am comfortable sharing my data on the internet, even without knowing exactly how it is being used*

* *I am only comfortable sharing my data if I know how it is being used*

* *I am not comfortable sharing my data at all*

Q2 How comfortable are you with sharing the following data? *Browsing history / Products purchased in the past / Bank data / Postal address / Telephone number / Geolocation / Emails sent and received / Contacts / Information shared on your social networks*

[Possible answers for **Q2**: Not at all comfortable / Not really comfortable / Somewhat comfortable / Very comfortable]

[Order of responses was randomized per each participant]

Recall and reasons for users' choices

Q3 On the cookies banner of the last site visited, you have: *Accepted / Declined / Customised your choices / I don't remember*

Q4 You saw three cookie banners overall, in your opinion, were they all the same? *Yes / No / I don't remember*

Q5 You chose to click “Decline all” on the first site when you saw the following banner. [A cookie banner is shown only to users who clicked “Decline all”] For what reasons did you click “Refuse all”? Tick all that apply [Random order of answers; multiple choice]

To avoid being followed or traced / Out of habit / Because it was the fastest option / Because the option to accept everything was not available / To avoid having personalised ads about this type of product / I clicked randomly / Other, please specify :

Q6 You chose to click “Accept all” on the first site when you saw the following banner. [A cookie banner is shown only to users who clicked “Accept all”] For what reasons did you click “Accept all”? Tick all that apply [Random order of answers; multiple choice]

Out of habit / Because it was the fastest option / Because the option to refuse everything was not available / To receive personalised advertisements about this type of product / For fear of not being able to access the site / I clicked randomly / Other, please specify :

Q7 You chose to customise cookies on the first site when you saw the following banner. [The cookie banner they saw in the first exercise is shown only to users who clicked “Customize”] For what reasons did you decide to personalise your choice? Check all that apply [Random order of answers; multiple choice]. *Out of habit / Because the option to refuse everything was not available / Because the option to accept everything was not available / To choose how my data will be used / I clicked randomly / Other, please specify :*

Q8 Are you satisfied with the choice you made on this cookie banner? [Shown to all] *Yes / No, I would have liked to refuse / No, I would have liked to accept / No, I would have liked to personalise my choices / I don't care, as long as the banner disappears*

Q9 What do you think of the following statements about this cookie banner? [Strongly disagree / Somewhat disagree / Somewhat agree / Strongly agree] *Navigating the banner is simple / The information is easy to understand / The choices are presented to me with my interests in mind / The banner allows me to easily control the level of data protection / More*

sites should use this format for banner cookies

Questions on knowledge Finally, we will ask you some questions about your general knowledge of cookie banners and data protection.

Q10 What data do you think can be collected if you accept cookies? *Tick all that apply* [Random order of answers, except the last two; multiple choice] *Browsing history / Products purchased in the past on the website visited / Bank data / Postal address / Telephone number / Geolocation / Emails sent and received / Contacts / Information shared on your social networks / None of the above / All of the above*

Q11 In your opinion, what are the sites' reasons for collecting data about you through cookies? (Tick all that apply) [Random order; multiple choice]. *To decide which ad to show me / To resell my data to other companies (third parties) / To improve my user experience on their sites (e.g. remembering my login details) / To allow me to view media content (e.g. YouTube video) on the sites I visit / To adjust the prices or promotions of the products I am looking for / To adapt the language and the display of the site to my geolocation / To report user data to government / To subscribe to their newsletter and send me marketing emails / To ensure security on the site I am visiting*

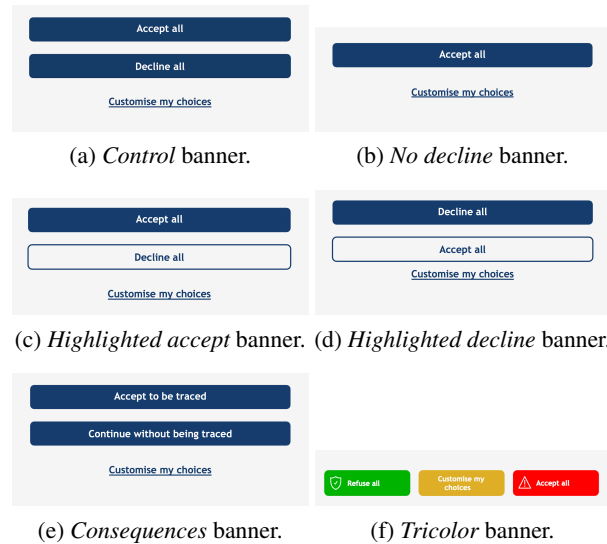
Q12 According to you, cookies collect data by : (Tick all that apply) [Random order of answers; multiple choice]. *Tracking what I do online (browsing history, services and products purchased) / Tracking information that I actively share (e.g. on social networks) / Listening to what I say through the microphones on my computer and phone / Collecting information that others have shared about me (e.g. on social networks)*

Opinions on alternative solutions

Q13 An extension to put on your browser could automatically refuse or accept cookies for you depending on the settings you select when you install it. Do you think you would use this type of extension to browse the internet if it existed? *Yes, most certainly / Yes, probably / I'm not sure / No, probably not / No, certainly not*

Q14 Last question!

Figure 6: Designs of the lowest part of the six consent banners, translated into English. A high-quality version of complete images of consent banners is included in Supplemental Materials [5, Appendix B].



Cookies are small files stored on your device that have several uses, including:

- * Memorise your login details with a merchant site,
- * Save the content in your shopping cart
- * Save the preferred display language and your country of connection
- * Track what you browsing for statistical or advertising purposes in order to offer you personalised content
- * Some of these uses are strictly necessary for the proper functioning of the site, others are used for commercial purposes.

Are you comfortable accepting all cookies (essential and non-essential) on the following sites? [Not at all comfortable / Not really comfortable / Somewhat comfortable / Very comfortable] [Random order of answers] *Banking sites / Health sites / Government sites / Online sales sites (shopping, clothing, concerts, transport) / Informational sites (Wikipedia, press, forums, etc.) / Social networks*

B Consent banners design variants

Figure 6 shows designs of the choice architecture (only the lowest part of the banners that differ across all six designs) for all banners in our experiment. The upper part (identical to all six designs) can be found in Figure 1.