

(Un)informed Consent: Studying GDPR Consent Notices in the Field

Christine Utz
Ruhr-Universität Bochum
Bochum, Germany
christine.utz@rub.de

Martin Degeling
Ruhr-Universität Bochum
Bochum, Germany
martin.degeling@rub.de

Sascha Fahl
Ruhr-Universität Bochum
Bochum, Germany
sascha.fahl@rub.de

Florian Schaub
University of Michigan
Ann Arbor, Michigan
fschaub@umich.edu

Thorsten Holz
Ruhr-Universität Bochum
Bochum, Germany
thorsten.holz@rub.de

ABSTRACT

Since the adoption of the General Data Protection Regulation (GDPR) in May 2018 more than 60 % of popular websites in Europe display *cookie consent notices* to their visitors. This has quickly led to users becoming fatigued with privacy notifications and contributed to the rise of both browser extensions that block these banners and demands for a solution that bundles consent across multiple websites or in the browser. In this work, we identify common properties of the graphical user interface of consent notices and conduct three experiments with more than 80,000 unique users on a German website to investigate the influence of notice position, type of choice, and content framing on consent. We find that users are more likely to interact with a notice shown in the lower (left) part of the screen. Given a binary choice, more users are willing to accept tracking compared to mechanisms that require them to allow cookie use for each category or company individually. We also show that the widespread practice of nudging has a large effect on the choices users make. Our experiments show that seemingly small implementation decisions can substantially impact whether and how people interact with consent notices. Our findings demonstrate the importance for regulation to not just require consent, but also provide clear requirements or guidance for how this consent has to be obtained in order to ensure that users can make free and informed choices.

CCS CONCEPTS

• **Security and privacy** → **Usability in security and privacy**;
• **Human-centered computing** → *Empirical studies in interaction design*; • **Social and professional topics** → *Governmental regulations*.

KEYWORDS

consent; notifications; usable privacy; GDPR

ACM Reference Format:

Christine Utz, Martin Degeling, Sascha Fahl, Florian Schaub, and Thorsten Holz. 2019. (Un)informed Consent: Studying GDPR Consent Notices in the Field. In *2019 ACM SIGSAC Conference on Computer and Communications Security (CCS '19)*, November 11–15, 2019, London, United Kingdom. ACM, New York, NY, USA, 18 pages. <https://doi.org/10.1145/3319535.3354212>

1 INTRODUCTION

In recent years, we have seen worldwide efforts to create or update privacy laws that address the challenges posed by pervasive computing and the “data economy”. Examples include the European Union’s General Data Protection Regulation (GDPR) [46], which went into effect on May 25, 2018, and the California Consumer Privacy Act (CCPA) [42], which becomes effective on January 1, 2020. These laws uphold informational self-determination by increasing transparency requirements for companies’ data collection practices and strengthening individuals’ rights regarding their personal data.

The GDPR’s impact was twofold. While the number of third-party services on European websites barely changed [41], websites now ask users for consent prior to setting cookies. In mid-2018, about 62 % of popular websites in the EU were found to display a *(cookie) consent notice*, often referred to as “cookie banner,” and in some countries an increase of up to 45 percentage points since January 2018 was observed [12]. The design and complexity of such consent notices greatly vary: Some merely state that the website uses cookies without providing any details or options, while others allow visitors to individually (de)select each third-party service used by the website. Paired with the fact that consent notices often cover parts of the website’s main content, this high prevalence has led website visitors to become fatigued with consent mechanisms [8]. Consequently, tools have emerged that provide pragmatic workarounds — one example is the “I don’t care about cookies” browser extension [21]. But oftentimes this only leads to data collection taking place without consent since the default on many websites is to employ user tracking *unless* the visitor has opted out [17], and 80 % of popular EU websites do not offer any type of opt-out at all [12].

Instead of adopting opt-in solutions or enforcing the existing Do-Not-Track standard, the online advertising industry has developed a consent framework [14] to reduce the number of consent requests. Notices using this framework ask website visitors if they consent to data collection for different purposes by up to 400 listed third-party

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

CCS '19, November 11–15, 2019, London, United Kingdom

© 2019 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-6747-9/19/11.

<https://doi.org/10.1145/3319535.3354212>

advertisers. Information about their consent decision is then passed down the online advertising supply chain.

Overall, consent notices have become ubiquitous but most provide too few or too many options, leaving people with the impression that their choices are not meaningful and fueling the habit to click any interaction element that causes the notice to go away instead of actively engaging with it and making an informed choice.

Most notice designs only partially use the available design space for consent notices. But we have also seen notices that, e. g., do not force users to accept cookies, ask for consent without hidden pre-selections, or provide visitors with granular yet easy-to-grasp mechanisms to control the website’s data processing practices. Hence, we expect that how a consent notice asks for consent has a large impact on how website visitors interact with it, and we are positive that there are design decisions that better motivate people to interact with consent notices in a meaningful way instead of annoying them.

In this paper, we systematically study design properties of existing consent notices and their effects on consent behavior. We systematize consent notices using a sample of 1,000 notices collected from live websites and identify common variables of their user interfaces. Our research goal is to explore the design space for consent notices to learn how to encourage website visitors to interact with a notice and make an active, meaningful choice. Over the course of four months, we conduct a between-subjects study with 82,890 real website visitors of a German e-commerce website and investigate their (non-)interaction with variants of consent notices. We collect passive clickstream data to determine how users interact with consent notices and invite them to participate in a voluntary follow-up online survey to obtain qualitative feedback. The study comprises three distinct field experiments to answer the following research questions:

- (1) Does the position of a cookie consent notice on a website influence visitors’ consent decisions? (Experiment 1, $n = 14,135$)
- (2) Do the number of choices and nudging via emphasis / pre-selection influence users’ decisions when facing cookie consent notices? (Experiment 2, $n = 36,530$)
- (3) Does the presence of a privacy policy link or the use of technical / non-technical language (“this website uses cookies” vs. “this website collects your data”) influence users’ consent decisions? (Experiment 3, $n = 32,225$)

In a short follow-up survey answered by more than 100 participants, we ask website visitors to voluntarily report the motivation for their selection, how they perceive the notice they have seen, and how they expect consent notices to function in general.

We find that visitors are most likely to interact with consent notices placed at the bottom (left) position in the browser window while bars at the top of the screen yielded the lowest interaction rates. This is mainly due to the (un)importance of the website content obstructed by the notices and suggests taking into account characteristics of the individual website to identify the notice position most likely to encourage user interaction. Interaction rates were higher with notices that provided at most two options compared to those that let users (de)activate data collection for different purposes or third parties individually, even though those notices do not allow visitors to express consent freely. We also show that the

more choices are offered in a notice, the more likely visitors were to decline the use of cookies. This underlines the importance of finding the right balance between providing enough detail to make people aware of a website’s data collection practices and not overwhelming them with too many options. At the same time, nudging visitors to accept privacy-invasive defaults leads more visitors to accept cookies, whereas in a privacy-by-default (opt-in) setting, less than 0.1 % of visitors allow cookies to be set for all purposes. This suggests that the current data-driven business models of many webservices, who often employ dark patterns to make people consent to data collection, may no longer be sustainable if the GDPR’s data protection by default principle is enforced. Technical language (“This site uses cookies” instead of “This site collects your data”) appears to yield higher interaction rates with the consent notice but decreases the chance that users allow cookie use. We find that the presence of a link to the site’s privacy policy does not increase user interaction, underlining the importance of making information immediately actionable rather than pointing to further resources.

Survey feedback indicates that users favor category-based choices over a vendor-based approach, and they expressed a desire for a transparent mechanism. A common motivation to give consent is the assumption that the website cannot be accessed otherwise.

Based on the results of our field study, we conclude that opt-out consent banners are unlikely to produce intentional/meaningful consent expression. We therefore recommend that websites offer opt-in notices based on categories of purposes. Above all, we observed that the majority of website visitors does not accept cookies for all purposes, and feedback from our survey suggests that a unified solution that does not interfere with every single website yet provides more control than a simple yes–no decision would best fit users’ needs.

2 CONSENT NOTICES

We first describe the legal background of consent notices and current challenges for their practical implementation. Then we identify and analyze variables of the graphical user interface of commonly used types of consent notices.

2.1 Background

Cookie consent notices emerged in the wake of the European Union’s Directive 2009/136/EC [45]. The directive changed Article 5(3) of the ePrivacy Directive (2002/58/EC) [44] to require that data is stored on users’ devices only after having obtained user consent based on “*clear and comprehensive information [...] about the purposes of the processing.*” An exemption to this consent requirement is storing of information that is “*strictly necessary,*” such as session or authentication cookies.

On May 25, 2018, the European Union’s General Data Protection Regulation (GDPR; Regulation (EU) 2016/679) went into effect. Its Article 6 contains six legal bases for the processing of personal data of European residents, including that “*the data subject has given consent to the processing of his or her personal data for one or more specific purposes*”. Recital 32 of the GDPR and guidelines published by EU data protection authorities [6] require for valid consent “*a clear affirmative act*” that is a “*freely given, [purpose-]specific, informed and unambiguous indication of [...] agreement*

to the processing of personal data.” Another document clarifies the relationship between the ePrivacy Directive (2002/58/EC) and the GDPR for the use of cookies: Article 5(3) of the directive governs access to non-necessary cookies in the user’s browser, whether it contains personal data or not, while the GDPR applies to subsequent processing of personal data retrieved via cookies [15].

Degeling et al. found that after the GDPR went into effect 62.1% of 6,579 popular websites in Europe displayed cookie consent notices, compared to 46.1% in January 2018 [12].

This high prevalence has sparked efforts to reduce the number of consents required. The most widely used solution, supported by the online advertising industry, is the Transparency & Consent Framework by IAB Europe [14]. This framework has been criticized for its bundling of purposes [36] and a lack of transparency regarding the parties the website visitor’s personal data could be shared with [12, 36]. An October 2018 decision by the French data protection authority CNIL [11] pointed out a lack of consent verification in the framework, and in April 2019 a formal complaint was filed against the IAB for showing a consent notice on its own website that forces visitors to consent if they want to access the website [37], which is not allowed under GDPR.

Another suggestion to decrease the number of consent prompts is to move consent decisions to the browser and let users locally specify their data collection preferences [31]. The browser then sends adequate signals to the websites requesting data collection. This would require websites to respect the opt-out signals requested by the browser — something that has not worked out in the past with the Do-Not-Track standard [29].

2.2 Properties of Consent Notices

Consent notices currently found on websites vary both in terms of their user interface and their underlying functionality. Regarding the latter, Degeling et al. identified distinct groups within existing implementations of consent notices [12]. Some are only capable of displaying a notification that the website uses cookies or collects user data without providing any functionality to make the website comply with the visitor’s choice. In contrast, other cookie notices are provided by third-party services that offer complex opt-in choices and block cookies until the user consents explicitly.

Our study focuses on the *user interface* of consent notices, a topic which has not been systematically studied before. In order to identify common properties of consent notices currently used on websites, we analyze a random sample of 1,000 notices drawn from a set of 5,087 we collected in a previous study [12]. To obtain that set, the following steps were taken: First we created a list of websites containing the 500 most popular websites for each member state of the European Union as identified by the ranking service Alexa [4]. This yielded a list of more than 6,000 unique domains. Using a Selenium-based automated browser setup, we visited all of them in an automated way in August 2018 from an IP address within the EU and took screenshots of each website’s home page. We then manually inspected these screenshots if they contained a consent notice. In our previous study, we identified six distinct types of choices consent notices offer to website visitors, as described below. In this work, we extend our prior analysis to other variables of the graphical user interface of consent notices. For this, we took the

5,087 consent notices collected previously, drew a random sample of 1,000 notices, and manually inspected how they differed in their graphical user interface. We identified the following eight variables, whose possible values, along with their frequency in our random sample, are listed in Table 1:

Size. The size of the consent notice as displayed in the browser. We found the value of this variable to vary widely depending on the implementation of the notice, from small boxes that only cover a fraction of the viewport to notices taking up the whole screen. Responsive web design may result in the same notice using up different shares of the viewport, depending on the screen size and orientation of the device used to view the website. Typically notices take up a larger percentage of the viewport on smartphones than on desktop computers and tablets. The size of a consent notice may also be fixed by design, i. e., to cover the whole viewport of any device.

Position. We observed the consent notices in our dataset to be displayed in seven distinct positions: in one of the four corners of the viewport (*dialog* style; 6.9%), at the top (27.0%) or bottom (57.9%) like a website header or footer (*bar* style), and vertically and horizontally centered in the middle of the viewport (7.8%). On smartphones in portrait mode, the limited space reduces the number of options to the top, bottom, and middle of the screen.

Blocking. Some consent notices (7.0%) prevent visitors from interacting with the underlying website before a decision is made [39]. The site’s content may also be blurred out or dimmed [17]. All consent notices shown in the center position were blocking. We also observed some blocking consent notices at the top or bottom position.

Choices. Consent notices offer website visitors different choice options. We identified the following mechanisms for user interaction [12]:

- **No option** notices simply inform the user that the website uses cookies without any option for interaction. The user continuing to use the website is interpreted as agreement to the notice.
- **Confirmation-only** banners feature a button with an affirmative text such as “OK” or “I agree”, clicking on which is interpreted as an expression of user consent.
- **Binary** notices provide two buttons to either accept or decline the use of all cookies on the website.
- **Category-based** notices group the website’s cookies into a varying number of categories. Visitors can allow or disallow cookies for each category individually, typically by (un)checking a checkbox or toggling a switch. For transparency reasons, the category of “strictly necessary” cookies (whose use does not require consent according to Article 5(3) of Directive 2002/28/EC) is often also listed but the switch to deactivate it is greyed out. Some notices use a **slider**: Instead of (de)selecting categories individually the user can move a slider to select one of the predefined levels, which implies consent to all of the previously listed categories.
- **Vendor-based** notices offer even more fine-grained control by allowing visitors to accept or decline cookies for each third-party service used by the website. Such notices are part of IAB Europe’s Transparency and Consent Framework [14], which refers to its advertising partners as “vendors.”

Table 1: Variables of the graphical user interface of consent notices and their values across a sample of 1,000 drawn from 5,087 consent notices collected from the most popular websites in the European Union in August 2018

Position	Choices (visible)	Choices (hidden)	Blocking	Nudging
top	27.0 % no option	27.8 % no option	26.3 % yes	7.0 % yes
bottom	57.9 % confirmation	68.0 % confirmation	59.9 % no	93.0 % no
top right	0.2 % binary	3.2 % binary	4.0 %	n/a ^a
bottom right	3.0 % categories	1.0 % slider	0.2 %	27.8 %
top left	0 % vendors	0 % categories	8.1 %	
bottom left	3.7 %	vendors	1.1 %	
center	7.8 %	other	0.4 %	
other	0.4 %			
Link to privacy policy	Text: Collection	Text: Processor	Text: Purposes	
yes	92.3 % "cookies"	94.8 % unspecified	75.5 % generic	45.5 %
no	6.6 % "data"	1.4 % first party	0.7 % specific	38.6 %
other	1.1 % both	1.6 % third party	2.6 % & none	16.9 %
	none	0.9 % both	21.1 %	
	other	1.3 % other	0.1 %	

^a Nudging is not available for "no option" notices.

Text. The text displayed by consent notices also varies widely. It should inform the website visitor of the fact that the website uses cookies or similar tracking technology and may list additional information such as the purpose of the data collection. Depending on the choices offered, the notice may provide instructions for consenting to (or denying) the use of cookies. Table 1 provides an overview of common text contents of consent notices for the following typical pieces of information:

- **Collection.** What the visitor consents to, which can be the use of cookies (94.8%), the collection of their personal data (1.4%), both (1.6%), neither (0.9%), or something else (such as the website’s privacy policy; 1.3%).
- **Processor.** Who collects this information, which can be specifically limited to the first party (0.7%), third-party services (2.6%), both (21.1%), or refer to an unspecified party (usually denoted by the pronoun “we” or the domain/website name; 75.5%).
- **Purposes.** These may be specific (e. g., “audience measurement” or “ad delivery”; 38.6%), generic (e. g., “to improve user experience”; 45.5%), or not specified at all (16.9%).

Nudging & Dark Patterns. Consent notices often (57.4%) use interface design to steer website visitors towards accepting privacy-unfriendly options. Typical techniques include color highlighting of the button to accept privacy-unfriendly defaults, hiding advanced settings behind hard to see links, and pre-selecting checkboxes that activate data collection [10]. We observed all of these techniques in our sample.

Formatting. We found that, unless predetermined by the consent library used, the choice of fonts and colors typically matched that of the underlying website. The formatting of consent notices may also be influenced by the website’s business requirements [17], e. g., sites relying on monetization via online behavioral advertising (OBA) are unlikely to steer their visitors towards an opt-out mechanism by making this option highly visible.

Link to additional information. Consent notices may include a link to the website’s privacy policy, a designated cookie policy, or

a website providing additional information about cookies – 92.3% of the notices in our sample contain such a link to additional information. In Table 1, we marked as “other” consent notices where the full privacy policy was already included in the notice itself (1.1%).

Table 1 shows that the majority of consent notices are placed at the bottom of the screen (58%), not blocking the interaction with the website (93%). They offer no options besides a confirmation button that does not do anything (86%), and most try to nudge users towards consenting (57%). While nearly all notices (92%) contain a link to a privacy policy, only a third (39%) mention the specific purpose of the data collection or who can access the data (21%).

3 METHOD

Given the legal requirements for explicit, informed consent, the vast majority of cookie consent notices we analyzed are likely not compliant with European privacy law. To further investigate the effects of different combinations of these properties on consent behavior, we conducted a field study with consent notices on a German e-commerce website.

We investigated the effect of the following parameters on users’ interactions with consent notices:

- (1) The *position* of the notice, as notices displayed in some parts of the screen are more likely to be ignored.
- (2) The number of *choices* offered by the notice, which is influenced by legal requirements and the need to give users actual control over the website without overwhelming them with too many options.
- (3) *Nudging* visitors towards giving consent through highlighting and preselection, since this may cause people to consent who would not have made the same decision otherwise.
- (4) The presence of a *privacy policy link* and whether the notice refers to “cookie use” (technical language) or “data collection” (non-technical language). These differences in wording may influence people’s expectations of the website’s data processing practices and thus their consent decision.

We did not evaluate the effects of the following parameters: blocking (because the owner of our partner website asked us not to block access to the site), formatting (because of the multitude of options – we chose the same color scheme as in the notice previously used on the website), and size (which is difficult to vary consistently across devices).

From the end of November 2018 to mid-March 2019, we conducted three between-subjects experiments to determine if, and how, different parameters of consent notices influence interaction rates. In each experiment, we tested variants for one or two of the parameters described in Table 1: position in Experiment 1, choices and nudging in Experiment 2, and wording and the presence of a privacy policy link in Experiment 3. The respective other parameters were kept constant in an experiment.

3.1 Study Setup

We partnered with a German-language e-commerce website based on WordPress. The website has 15,000–20,000 unique visitors per month, most of which are single-page visitors that reach the site from a search engine looking for product information and reviews. The third-party services used by the website are Google Fonts and the CSS framework Ionic for design, Google Analytics embedded via Google Tag Manager for audience measurement, Facebook social media buttons, embedded YouTube videos, and targeted advertisements delivered by Google Ads. All of these services store cookies in the visitor’s browser.

We modified a WordPress plugin, Ginger – EU Cookie Law [26], to test different notice variants. Ginger was selected because it can block cookies before opt-in, log users’ consent, and because it was released under a GPLv2 license. By the time of publication of this paper, the original version of the plugin had been discontinued. We added support for checkbox-based and “no option” notices. We did not implement “slider” notices because we considered them a less compliant variant of the “categories” type.

The plugin was further modified to function as follows in our study: When a user first visited our partner website, they were shown one consent notice. Which notice of the n test conditions in the current experiment was displayed was determined in round-robin fashion. The ID of the displayed notice was stored in a cookie in the participant’s browser to ensure visitors who did not click the notice would continue to see the same notice across subpages and recurring visits. Each participant was assigned a unique identifier: $pid=SHA-256(ip_address||user_agent)$. The participant’s IP address was discarded after computation of pid . The participant ID was stored in another cookie, together with the participant’s consent as required by Article 7 GDPR¹.

If the visitor clicked any interaction element that would usually cause a consent notice to disappear, i. e., the ‘X’ discard button, “Accept,” “Decline,” or “Submit,”² the notice did not disappear instantly. Instead, the notice content was replaced with an invitation

to take an online survey about their experiences with this and other consent notices (see Appendix B). The invitation disclosed that this was a university study and that participants could win one of 15 25-euro shopping vouchers. Users could either click “Discard” to close the notice, or select “Participate” to open the survey in a new browser tab. The survey was created in a LimeSurvey instance running on a web server hosted by the authors.

If the website visitor did not interact with the consent notice, the content of the notice was automatically replaced with the survey invitation 30 seconds after the page had fully loaded. This is because we also wanted to explore users’ reasons for *not* interacting with consent notices. Web analytics data for our partner website showed that 95 % of all users who had interacted with the website’s previous consent notice had done so within 30 seconds of accessing the site. Thus we assumed that website visitors who did not interact with the consent notice within 30 seconds would not have clicked it at a later point in time.

We modified the Ginger plugin’s logger add-on to create log entries whenever a participant clicked an interaction element on the notice. Log events were also triggered upon page load, when links to the privacy policy or survey were clicked, when the consent notice content was auto-replaced with the survey invitation, and when the participant dismissed this invitation. Each log entry consisted of a timestamp, the participant’s ID (pid), the ID of the consent notice they had seen, the event they had triggered, their screen resolution, operating system, browser, and whether an ad blocker had been detected.³

3.2 Experiment 1: Position

Experiment 1 ran from November 30 to December 18, 2018, i. e., for 19 days. We had observed consent notices being shown at various screen positions and wanted to determine the effect of placement on interaction with the cookie consent notice to inform our subsequent experiments. The research question for Experiment 1 was: *Does the cookie consent notice’s position on a website influence a visitor’s consent decision?* In order to encourage user interaction, we displayed a “binary” notice without nudging (see Figure 1(bb)), the simplest type offering an actual choice. We tested the notice in six different positions (see Figure 2). We could not test the center position as our partner asked us to not block access to their website.

3.3 Experiment 2: Number of Choices, Neutral Presentation vs. Nudging

From December 19, 2018 to January 28, 2019, we conducted Experiment 2, which focused on the effects of given choices and pre-selections on consent. In our analysis of consent notices, we had identified various complexity levels of choices offered and methods to emphasize certain options. Prior work has shown that the design and architecture of choices heavily influences people’s decisions [43, 50]. While this effect has also been shown successful in improving user privacy [1, 2], in practice it is most often used to make users share more information [10]. Website owners often have an interest in getting visitors to agree to the use of cookies and hence highlight certain choices in the consent notice to nudge

¹The legal bases for storing the cookie that remembers the banner ID are Article 6(1)(e) GDPR (public interest in conducting this study) and Article 6(1)(c) GDPR (compliance with a legal obligation) for storing the consent cookie.

²In all experiments, all texts in the consent notice and survey were in German to match the website’s language. Survey responses were also in German. The authors translated all texts and responses into English for this paper. Both the original and the translated consent notices and the survey are available in our GitHub repository at <https://github.com/RUB-SysSec/uninformed-consent>.

³We used BlockAdBlock 3.2.1 (<https://github.com/sitexw/BlockAdBlock>) to detect ad blocking functionality in the visitor’s browser.

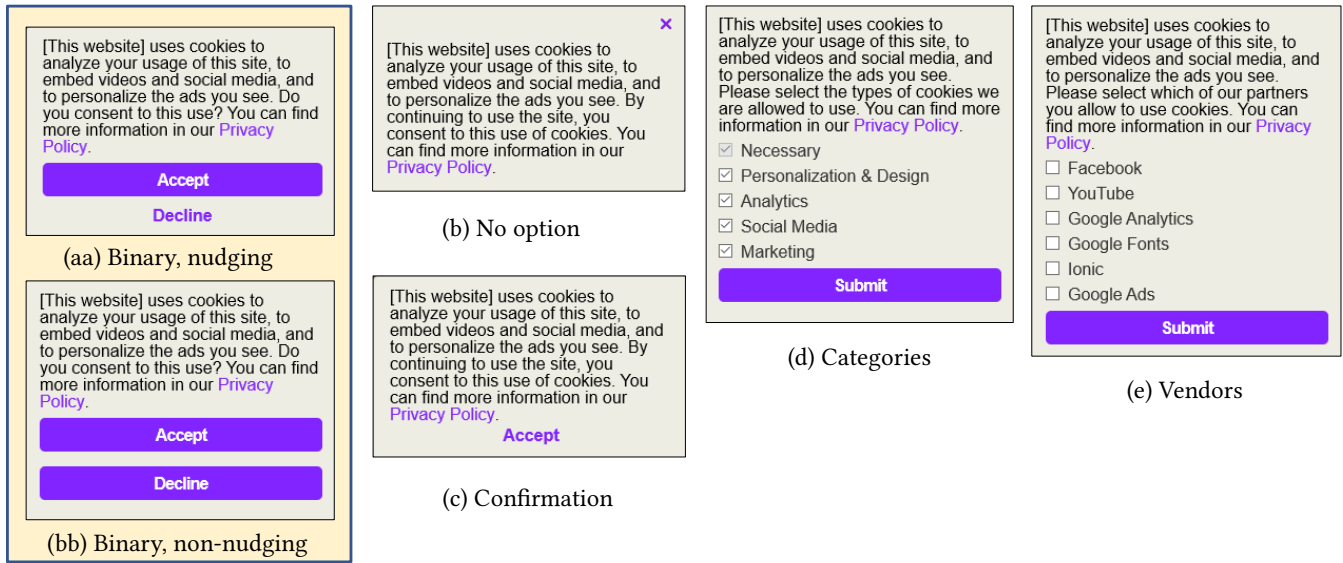


Figure 1: Cookie consent notices with different choice mechanisms and nudging used in our experiments: (a) a binary notice in two variants, one nudging visitors to click “Accept” (aa) and one presenting both choices equally (bb); (b) a no-option notice (nudging not applicable); (c) a confirmation-only notice (shown without nudging); (d) a category-based notice with pre-selected checkboxes (nudging); and (e) a vendor-based notice with unchecked checkboxes (non-nudging).




Figure 2: Positions tested in Experiment 1.

visitors towards accepting. We observed this for 57.4% of the notices in our sample. Our research question therefore was: *Does the number of choices and nudging through emphasis or pre-selection in consent notices influence user’s consent decisions?*

For nudging, we used pre-checked checkboxes and buttons highlighted in contrasting colors, techniques often used to nudge users towards accepting default settings [10]. While we observed that most category- and vendor-type notices in practice display such fine-grained controls only after the visitor clicked “Settings,” we chose to immediately display all available options to ensure that our conditions only varied in the number and framing of choices.

In Experiment 2, we displayed the following consent notices at the position determined in Experiment 1 to yield the highest interaction rates:

- **No option** (Figure 1 (b)): In line with many notices we observed, we added an ‘X’ in the top-right corner to dismiss

the banner. There is no nudging variant because the notice does not offer any choice.

- **Confirmation–Non-nudging** (Figure 1 (c)): This notice has an “Accept” button which is not highlighted.
- **Confirmation–Nudging**: Same as the Confirmation–Non-nudging notice, but the “Accept” button is highlighted (like the “Accept” button in Figure 1 (a) (aa)).
- **Binary–Non-nudging** (Figure 1 (a) (bb)): The “Accept” and “Decline” buttons are formatted the same way, neither is emphasized.
- **Binary–Nudging** (Figure 1 (a) (aa)): Same as Binary–Non-nudging but only the “Accept” button is highlighted in a contrasting color.
- **Categories–Non-nudging**: Same as notice (d) in Figure 1, but with unchecked checkboxes. The “Necessary” category cannot be unchecked, as is common practice.
- **Categories–Nudging** (Figure 1(d)): Same as Categories–Non-nudging but with pre-checked checkboxes for all categories.
- **Vendors–Non-nudging** (Figure 1(e)): Similar to the categories variant, but the checkboxes correspond to the third-party services used by our partner website.
- **Vendors–Nudging**: Same as Vendors–Non-nudging but with pre-selected checkboxes.

For the category-based notices, we had to map the third-party services used by the website to different categories. We manually inspected the 434 category-based notices in our initial set of 5,087 consent notices for common category wording. For example, we found advertising cookies to be categorized as “marketing” or “advertising”; web analytics was also referred to as “performance

cookies,” “statistics,” or “audience measurement.” This yielded the following category–third party mappings:

- **Necessary:** Cookies to remember the displayed notice and the website visitor’s consent decision.
- **Personalization & Design:** Ionic, Google Fonts
- **Analytics:** Google Analytics
- **Social Media:** Facebook, YouTube
- **Marketing:** Google Ads

For all category- and vendor-based notices in Experiments 2 and 3, the available options were displayed in random order, except for the “Necessary” category, which was always displayed first as in the majority of category-based notices we had observed.

In Experiments 2 and 3, we increased the font size of the banner message, resulting in larger notices. We did this to fix an implementation bug of the Ginger plugin that had caused the text to be displayed in a very small font on some smartphones in portrait mode.

3.4 Experiment 3: (Non-)Technical Language and Privacy Policy Link

Experiment 3 was conducted from January 29 to March 15, 2019. In this experiment, we tested the influence of the presence of a link to the website’s privacy policy. Previous research suggests that (American) Internet users have consistent misconceptions about privacy policies, indicated by the fact that a majority believes the existence of a privacy policy means that a website cannot share personal data with third parties [47]. At the same time, Martin [27] showed that the existence of a reference to a privacy policy in the context of data sharing explanations increases mistrust in a website. There are further known misconceptions about what cookies actually are and what they are used for [19, 30]. To learn more about the influence of these factors in the context of consent notices, our research question was: *Does the presence of a privacy policy link or mention of cookies influence users’ consent decisions?*

The base notice for this experiment was the Category–Non-nudging notice from Experiment 2 because of GDPR’s data protection by default requirement and the ability to provide consent for specific purposes with checkboxes. We chose a category-based notice over a vendor-based one due to the results of Experiment 2 (see Section 4.3). The notice text for this experiment was: “This website [uses cookies | collects your data] to analyze your usage of this site, to embed videos and social media, and to personalize the ads you see. Please select for which purposes we are allowed to use your data. [You can find more information in our privacy policy].” We tested the following conditions:

- **Technical–PP Link:** The original Categories–Non-nudging notice from Experiment 2. It uses both technical language (“collects cookies”) and a sentence with a link to the website’s privacy policy.
- **Technical–No PP Link:** Same as Technical–PP Link, but the privacy policy sentence was replaced with whitespace to keep the size of the notices consistent.
- **Non-Technical–PP Link:** Same as Technical–PP Link, but using non-technical language (“your data” instead of “cookies”).

- **Non-Technical–No PP Link:** Same as Non-Technical–PP Link, but with the privacy policy sentence replaced with whitespace.

For participants who saw a notice with non-technical language, we replaced other occurrences of the term “cookie” in our setup: In the study invitation, “cookie notice” was replaced with “privacy notice,” and we adjusted the wording of some survey questions and response options as described in Appendix B.

3.5 Research Ethics

Our study was conducted on a website with real users, which raises ethical concerns as we did not ask for consent prior to measuring their interactions with consent notices. We did so to ensure ecological validity and be able to capture non-biased results as we expected the majority of visitors to not pay attention to a study consent notice asking them to opt in, which was supported by our findings.

While our institution does not require IRB review for minimal risk studies, we ensured that we did not deceive or harm website visitors and their privacy. All displayed consent notices functioned as described and respected the visitor’s choice.

To test the effect of no-option consent notices, we had to offer fewer choices than we believe is required by the GDPR. We added a paragraph describing our study to the website’s privacy policy. The data we collected was pseudonymized. Logs were stored on the website’s server and access was limited to two researchers conducting the analysis and the website’s owner. After the study, the data was removed from the server and copied to the researchers’ data center.

All visitors were informed about the study after 30 seconds when we showed a notice asking them for participation in the survey. Survey participants were asked for explicit consent and to confirm they were over 18 and wanted to participate. Email addresses of participants who opted to participate in the prize draw were stored separately from the dataset, without the participant ID.

3.6 Data Analysis

3.6.1 Event logs. When we started the data analysis, we noticed inconsistencies in some entries. The event logs created by our plugin indicated that some website visitors had seen multiple notice versions. This could have happened because users had deactivated cookies completely, visited the website in multiple sessions using private browsing mode, or opened the website in multiple tabs simultaneously. For another set of users, we detected multiple screen resolutions, mostly because the screen orientation had changed. Rotating the screen could lead to the notice covering different parts of the website, so we removed these participants to preserve consistency. In total, we removed 2,1 % of participants across all experiments.

3.6.2 Survey. We considered a survey response complete if the participant had at least answered Q1–Q6 but did not provide a free-text answer to Q7 and Q8. Due to a low survey response rate we received few responses for some conditions. We therefore refrained from a quantitative analysis of survey responses. In Section 4, we evaluate responses to the open-ended questions (parts of Q1; Q6–Q8). We coded these responses using emergent thematic coding.

Two of the authors independently devised a set of codes for each question and coded the responses. The results were discussed and yielded a final codebook, which was used to re-code all responses. Any remaining disagreements were reconciled by the two coders. We report the codes and their distribution in Appendix B, along with the answers to all closed-ended questions.

4 RESULTS

4.1 Dataset and Website Visitors

Our cleaned dataset contained event logs of 82,890 unique website visitors: 14,135 in Experiment 1, 36,530 in Experiment 2, and 32,225 in Experiment 3. 21.72% of all visitors accessed the website on a desktop or laptop computer and 78.28% with a mobile device (of which 5.1% were tablets)⁴. Overall, 6.95% of participants used an ad blocker. The rate was much higher on desktop (29.1%) than on mobile devices (0.8%). These numbers are consistent with a 2017 report for Germany [35], the highest rate of ad block users in Western Europe (20% on average), and North America (18% on average). For 16.45% of visitors, we could not detect whether they used an ad blocker. These visitors did not stay long enough on the website to complete ad blocker detection. On average, users spent a short time on the website. Pre-study Google Analytics data provided by the partner website showed that 84.81% of visitors spend less than 10 seconds on the site, 5.21% 11 to 60 seconds, and 5.83% up to 3 minutes. Our dataset includes all users for whom the event logs indicated a fully loaded site, regardless of how long they stayed on the page, resulting in a high number of “no action” visitors. As described in Section 4.3, the median time until an interaction with any version of the notice was 4 to 8 seconds. About 11,800 users stayed on the page for 10 seconds or more.

The link to our survey was clicked 804 times (168 in Experiment 1, 445 in Experiment 2, and 191 in Experiment 3). We received a total of 110 responses (16 in Experiment 1, 60 in Experiment 2, and 34 in Experiment 3), which means that 0.37% of the 29,712 visitors who interacted with the notice or stayed on the site for longer than 30 seconds participated in the survey. To get an impression of visitors’ expectations about the website’s data collection practices, we asked Q2: *What do you think – what data does [the website] collect about you when you access the website?* This question was answered by all participants. Across all three studies, the data most commonly expected to be collected were links clicked on the site (78%), IP address (65%), posts read on the site (61%), and the device used (59%). Less often mentioned were other sites visited (29%) and the visitor’s place of residence (25%). 13% thought the website collected their name, even though the site never asks for it. Only 5% thought the site did not collect any data about them. These answers indicate that the survey participants had a good understanding of what data websites can collect even without user accounts.

4.2 Experiment 1: Banner Position

4.2.1 Interaction rates. Figure 3 shows how visitors interacted with the consent notices displayed at different positions. Overall the notices shown at the bottom-left position received the most interactions, 33.1% of visitors interacted with them regardless of

⁴We count as “desktop computer” actual desktop machines as well as laptops. “Mobile” devices include smartphones and tablets; the latter were used by 5.1% of visitors.




Figure 3: Interaction rates in Experiment 1 (notice position), arranged pairwise for mobile and desktop users.

device type or choice made. The notice positions most commonly observed in practice, small bars at the top or bottom, resulted in low interaction (2.9% and 9.6%, respectively).

While we were mainly interested in position in Experiment 1, we also analyzed the influence of other variables, such as ad blocker use, screen resolution, browser, operating system, and device type (desktop/mobile). We estimated the effect size of different properties by calculating Cramér’s V (CV) and over all visitors the banner position showed the largest effect size (CV=.31). Unless noted otherwise, χ^2 -tests for effects in this experiment are statistically significant ($p < .001$).

Ad blocker use also had a small impact on whether someone interacted with the notice. While on average 15.8% of visitors without an ad blocker interacted with any notice, only 12.6% of ad blocker users did so, but the effect size was rather small (CV=.11). The impact of screen resolution was much higher on desktop (CV=0.33) than on mobile (CV=0.16): Only 5.5% of visitors with screen resolutions of 1,920 by 1,080 pixels or higher interacted with the notice, while the average was 25.6% for smaller screens. Although the decline/accept ratio varied between conditions, we could not identify a single factor to explain the differences. Across all conditions the number of users who accepted cookies was higher than the number of those that declined.

4.2.2 Discussion. A possible explanation for higher interaction rates with notices displayed at the bottom is that these notices are more likely to cover the main content of the website, while

notices shown at the top mostly hide design elements like the website header or logo. If one uses their thumb to navigate websites on a smartphone, it is also easier to tap elements on the bottom part of the screen than those at the top. An explanation for higher interaction rates with notices displayed on the left of the viewport might be the left-to-right directionality of Latin script: Line breaks cause the information density of a text to be skewed to the left, so consent notices positioned on the left are more likely to obstruct visitors’ reading and trigger an interaction with the notice.

We looked for qualitative feedback in the survey responses. In Experiment 1, we received 16 responses, with eight participants having interacted with the notice and another eight that did not. All six participants who answered they had clicked the notice “because it prevented them from reading the website content” had seen a notice shown at the bottom or left side.

Both on desktop and mobile, the notice positioned in the bottom-left corner received the most attention. Thus, we decided to display the notices in Experiments 2 and 3 in the bottom-left corner.

4.3 Experiment 2: Choices & Nudging

In Experiment 2 there were 36,395 participants in total. Each of the nine conditions was shown to 4,044 website visitors on average.

4.3.1 Interaction rates. Figure 4 provides an overview of the recorded visitor interactions. Compared to Experiment 1, the overall percentage of visitors who interacted with the notice increased (13,8%–55,3%), especially on mobile devices, likely because we had increased the font size, resulting in larger notices. The highest interaction rate (55%) was measured for binary notices on mobile devices.

The experiment revealed a strong impact of nudges and pre-selections. Overall the effect size between nudging (as a binary factor) and choice was $CV=.50$. For example, even for confirmation-only notices, more users clicked “Accept” in the nudge condition, in which it was highlighted (50.8% mobile, 26.9% desktop), than in the non-nudging condition, in which “Accept” was displayed as a text link (39.2% m, 21.1% d). The effect was most pronounced for category- and vendor-based notices, in which all checkboxes were pre-selected in the nudging conditions, but not in the privacy-by-default conditions. The pre-selected versions led around 30% of mobile users and 10% of desktop users to accept all third parties. In contrast, only a small fraction (< 0.1%) allowed all third parties when given the opt-in choice and 1 to 4% allowed one or more third parties (“other” in Figure 4), indicating that some users still engaged with the offered choices. No desktop visitors allowed all categories. Interestingly, the number of non-interacting users was highest on average for the vendor-based conditions, although they took up the largest amount of screen space due to six options being offered. We discuss qualitative survey feedback on the category- and vendor-based notices in Section 4.5.2.

4.3.2 Choices. Results were mixed in terms of the consent choices users made when given options (in all but the no-option and confirmation conditions). Surprisingly, more participants accepted cookies in both binary conditions, where they had the option to decline cookies, than in the non-nudging confirmation condition, where they could only accept cookies or not interact with the notice.




Figure 4: Visitors’ consent choices in Experiment 2. “Accept”/“Decline” indicate that (all) options were accepted or declined. “Other” includes those who accepted/declined only some options. Bold figures indicate default options.

Figure 5 lists the specific choices participants made on category- and vendor-based notices. Few visitors chose specific categories or vendors if they were not pre-selected (non-nudging conditions). Interestingly, more visitors selected specific vendors than categories. Vendors YouTube and Ionic were selected most, even though survey responses (Q6) indicated that Ionic was lesser known than other listed vendors. We observe a similar pattern for the de-selection of specific categories and vendors: More visitors unchecked one or more vendors (10.0%) than categories (6.9%).

6% of visitors who saw a category- or vendor-based notice clicked at least one of the checkboxes more than once. 48 visitors (0,08%) toggled an even number of times, reversing previous

Table 2: Comparison of interactions with category notices

Dataset	Decision	None pre-selected	all pre-selected
Cookiebot		(n = 1,135,090)	(n = 1,988,681)
	Accept	5.59 %	98.84 %
	Decline	94.41 %	1.16 %
Our Data		(n = 1,239)	(n = 1,380)
	Accept	0.16 %	83.55 %
	Decline	99.84 %	16.45 %

decisions. Interestingly, 47 of those users saw a “nudging” notice so that they actively reactivated one of the categories.

We also recorded how long it took visitors to submit their choice. The median time to submit for no-option, confirmation and binary-choice notices was 4–5 seconds; 7–8 seconds for category- or vendor-based notices.⁵ For details see Appendix A.

4.3.3 External validation. To verify the generalizability of our results, which are only based on visitors to our partner website, we compared our data to internal data from Cookiebot, a company offering cookie consent notices (similar to our category-based conditions) as a service to websites. Their dataset from February 2019 contains 3 million user logs for 2,000 different websites. The Cookiebot notices also show purpose categories, so we compare their data with our data for the category-type notices. In their case, some of the checkbox selections cannot be changed by users, as website owners can argue that the use of certain cookie categories is based on different legal grounds (e. g., “legitimate interest”, Art. 6 (1) (f) GDPR). Therefore (de)selecting all consent-based cookie categories in Cookiebot notices sometimes requires fewer clicks to be made, and we were not able to compare decisions we labeled as “other”. As shown in Table 2, Cookiebot has a slightly higher acceptance rate (5.6 % compared to 0.16 % in our dataset) and a lower decline rate when all boxes are pre-selected (1.2 % compared to 16.5 % in our dataset). This means that our findings are generally comparable, but specific results may differ based on website and category, which is what we would expect given that privacy preferences are highly contextual [3]. A related 2017 study (n = 300) found that about 3 % of users are willing to accept marketing cookies [34], which is between marketing acceptance in our non-nudging (0.6 %) and nudging (7.3 %) conditions.

4.3.4 Discussion. Experiment 2’s results show that nudges and pre-selection had a high impact on users’ consent decisions. It also underlines that the GDPR’s data protection by default requirement, if properly enforced, could ensure that consent notices collect explicit consent. We further find that most visitors make binary decisions even when more choices are offered by agreeing to all or no options. Only very few visitors selected specific categories or vendors, while even in the non-nudging binary condition a considerable number accepted the use of cookies. An explanation for this behavior might be that those who are somewhat OK with cookie

use are not willing to expend effort on enabling it. Another explanation, suggested by previous work [27], is that showing the actual practices decreases the trust in a website and therefore leads to more users making an informed decision to decline cookies.

4.4 Experiment 3: Language & Privacy Policy Link

In Experiment 3, we tested four conditions with combinations of (a) the notice including a link to the privacy policy (or not) and (b) the text either referring to “cookies” or “your data” more generally. All conditions were variants of the category-based, non-nudging notice from Experiment 2. Figure 6 summarizes the results. All conditions were shown to 6,032 visitors on average. Again, interaction rates were higher for mobile visitors. As in Experiment 2, very few visitors accepted all categories (0–0.1 %), but some visitors (0.3–1.4 %) explicitly allowed one or more. More people make a choice when technical language is used, i. e., “cookie” is mentioned in the notice. While this difference is significant (χ^2 -test, $p < .01$), the effect size is low ($CV = .08$), as are the differences between conditions. Presence of the privacy policy link had no significant effect ($p < .08$).

4.4.1 Discussion. Experiment 3 showed that mentioning of cookies has a minor influence on users’ consent behavior. However, differences between conditions are small. This is not surprising given that most users either submit the default choice or do not interact with the notice at all. We could not confirm previous studies [27] that showed a negative effect on trust in a website when a privacy policy was mentioned, but we found that more visitors decline the use of cookies if a privacy policy is linked. Our findings indicate that position and choice have a more pronounced effect on consent behavior than notice language or pointers to more privacy information.

4.5 Survey Results

4.5.1 Reasons for (Non-)Interaction with Notices. In the survey (see Appendix B), we asked participants why they did or did not click on the consent notice. Participants could select multiple reasons. 44 of 61 survey participants who had clicked the notice reported they had done so because they were annoyed by it. 16 thought the website would not work otherwise, and 13 stated they had clicked the notice out of habit. 11 participants interacted with the notice to protect their privacy, 6 for security reasons, and 5 to see fewer ads.

49 participants had not interacted with the consent notice, 20 of which reported they had not seen the notice. Nine thought clicking the notice would not have any effect, six did not care what cookies the website used or what data it collected, and three thought it did not offer enough choices. Two reported to not know what cookies were or what data the question was referring to. 13 participants selected “other” and provided a free-text response. Recurring themes in these responses include that the notices were “annoying [...]”, so I just ignore them out of frustration” (Participant 2-94)⁶ and that participants thought no cookies would be set if they did not interact with the notice. One participant mentioned that they “[found] all of the partners suspicious” (2-255). One had opened the website in

⁵We report the median as the data showed a high standard deviation since we had no way to check when the interaction with a notice started, and sometimes the choice was submitted minutes after the page had been loaded.

⁶The first digit in our participant identifiers denotes the experiment and the second the response ID assigned by LimeSurvey.




Figure 5: Decisions to allow or decline specific categories (1) or vendors (2) in the the specific conditions of Experiment 2. Subgraphs (a) show how many visitors checked specific boxes, subgraphs (b) how many unchecked pre-selected boxes.




Figure 6: Visitors' interactions with different consent mechanisms in Experiment 3. Notices contained technical language ("cookies") and a link to the privacy policy (or not).

a background browser tab, so they had only seen the invitation to take the survey, and two participants reported that the notice had been auto-replaced before they could click it.

4.5.2 Perception of Complex Consent Notices. We asked survey participants who saw a category- or vendor-based notice to elaborate on their choice selection (Q6), in order to learn how they perceived purpose-based consent mechanisms as required by the GDPR. We received 38 responses across Experiments 2 and 3. Appendix B lists the codes and their distribution for this and the following open-response questions.

A recurring theme in the responses was transparency, as mentioned by 5 participants who had seen a category-based notice: "[I liked] that I could directly select the options without going to the settings. It would be great if this was the default" (3-171), "What I like [here] is that only [the ...] necessary option is selected and all of the others are deactivated" (3-88). One participant with a

vendor-based notice stated: "Having options makes me feel secure" (2-619).

However, participants had diverging opinions regarding the notices' clarity. Some found the categories "self-explanatory" (3-118). Others pointed out that "Necessary [from a technical perspective] does not say much. Cookies aren't necessary to view a website" (3-215) and that "something could be hidden" (2-557) behind the Necessary category. 6 (of 7) participants who saw a vendor-based notice in Experiment 2 reported it had "too much text, too many options. I'm interested in the website's content, not in the consent notice" (2-116), and one suggested "it would be perfect to have a button to (de)activate all cookies" (2-199). Seven participants based their choices on privacy considerations: "I don't tick anything. I only need advice [from] the website" (3-108), "I don't want personalized web pages, ads, [... and] pointers to social media" (3-165).

These responses indicate that more complex notices are not necessarily problematic, as long as options are not pre-selected. While some express concerns, do not trust the categorizations, or find the choices too complex, others appreciate the privacy-by-default approach.

4.5.3 Understanding of Consent Notice Behavior. The survey further investigated participants' general understanding of how consent notices work and what it meant to accept or decline cookies. This section was identical in all three studies. The participant was shown the binary notice depicted in Figure 1 (a) (bb). Then we asked the following two free-text questions: Q7: *What do you think happens when you click "Decline"?* Q8: *What do you think happens when you click "Accept"?*

4.5.4 Declining Cookies. For Q7 (Decline), we received 94 responses across the three studies. We identified ten themes.

The most prominent expectation was that declining cookies would prevent access to the website (28 responses): "I don't get access to the desired information" (1-282), "The site closes itself and you are redirected to the search engine" (2-685). 17 other participants expected parts of the website not to work: "I won't be able to use some functionality because [...] cookies fund the website" (2-255). Only 4 participants explicitly mentioned that they would be able to access the site, stating, for example, "Normally I can continue to navigate the site. It has only happened twice that [a]

site has kicked me out. But online shopping [is] difficult if you don't agree" (2-94). 3 participants expected no collection or processing of personal data to take place when cookies are declined but still had doubts "I hope that no data is collected" (1-177, 1-121, 3-216). 12 expected the site to behave as if "Accept" were clicked: "I guess my data is still collected" (1-170), "Nothing, of course. Me not accepting cookies does not mean that the site uses less or no cookies or does not collect any data about me" (2-630). Other recurring themes in the responses include the expectation to see less ads, a focus on the technical aspects ("no cookies are evaluated" [3-217]), and if the notice would dis- or reappear. See Appendix B for details.

For Q8 (Accept), which was also answered by 94 participants (not all the same respondents as for Q7), we also identified 10 themes.

29 participants expected their personal data would be collected and/or processed: "my behavior on the website is stored and analyzed" (2-216), "my data is shared with who knows what third parties [...] Facebook, Google, marketing / market research / ad analytics [...]" (2-557). 19 responses focused on technical aspects: "a cookie is set which recognizes me when I revisit the website" (1-250). 21 participants stated the website would only work if they allowed cookies: "I can read the article" (2-53), "I can continue to use the website" (2-405). Other themes included effects on the consent notice only ("the banner disappears" [2-675]), personal data being collected for advertising, user profiling, and other purposes, e. g., "sale to third parties" (3-171), "influencing Internet algorithms" (1-269), and "any purpose" (1-207, 3-64). 7 participants believed it made no difference what was clicked but did not specify what that "default" behavior of the website would be.

These answers indicate that our participants had some understanding of how cookies are used, e. g., to recognize recurring visitors and for ad tracking and targeting. Concerningly, almost a quarter of participants thought they had to accept cookies before they could access a website – negative experiences on some sites may be influencing general expectations and behavior across websites. A transparent and GDPR-compliant consent notice should inform users which website functionality may not work as intended if cookies are declined.

5 RELATED WORK

Multiple measurement studies of varying scope have provided insights about the prevalence of consent notices [5, 12, 49]. Even though many consent notice libraries can be configured to only display a notice to EU visitors [12], van Eijk et al. [49] found that a website's top-level domain was the primary factor in whether a consent notice was displayed rather than a visitor's location.

Sanchez-Rola et al. [38] evaluated the functionality of consent notices and opt-out mechanisms under GDPR. They manually visited 2,000 popular websites, tried to opt out of data collection whenever possible, and studied the effects on the website's cookies. They found that 92 % of websites set at least one high-entropy cookie before showing any kind of notice. Only 4 % of notices provided an opt-out choice, and 2.5 % of websites removed some cookies upon opt-out. Degeling et al. [12] further found that many third-party consent libraries either lack the functionality to block or delete cookies, or require significant modification of a website to properly react to visitors' consent choices.

In Section 2, we presented a detailed analysis of variants in consent notices' graphical user interfaces. Previous work had only classified consent notices by the provided information [22], the choices offered [12, 38], and if the notice blocks access to the website [38]. Van Eijk et al. [49] report some statistics on the height and width of consent notices, their location offset, and notices' word and link/button counts.

Kulyk et al. [22] investigated users' perceptions of and reactions to differently worded cookie consent notices. They identified five categories of disclaimers based on the amount of information provided about the purposes of cookie use and the parties involved. In a qualitative user study, they found that the text of a cookie notice does not significantly influence users' decisions to continue using a website; their decision was rather based on the website's perceived trustworthiness and relevance. The participants perceived cookie consent notices as a nuisance or threat to their privacy, and reported lacking information about the implications of cookies and possible countermeasures.

Users' perceptions of consent notices' choice architectures have only been partially studied before. Boerman et al. [7], using Dutch panel data, explored how users protect their online privacy. Given the opportunity to decline cookies, many participants self-reported that they decline cookies "often" (16 %) or "very often" (17 %). Facing the decision to either accept cookies or leave the website, 12 % and 13 % reported to refrain from using the site "often" and "very often," respectively.

Previous work has shown that cookies are poorly understood by Web users. Ha et al. [19] studied the usability of two cookie management tools in focus groups, identifying misconceptions about cookies and risks associated with them. Kulyk et al. [23] developed and tested a privacy-friendly cookie settings interface for the Chrome browser and found that users appreciate tools that help them better understand the standard browser cookie settings, such as an assistant that transforms users' privacy preferences into cookie settings or additional explanations about the purpose and security/privacy implications of different types of cookies.

Consent notices are not the only way for Web users to opt out of targeted advertising. Previous work has evaluated the usability of different opt-out tools [18, 20, 24] and found that users find it difficult to locate, configure, and understand these mechanisms.

Schaub et al. describe the design space for privacy notices and controls, including consent notices and permission prompts on mobile devices [39].

Warning research and ad placement studies provide insights into the effects of user interface design choices on user attention and behavior; examples include color [40] and position [9]. Studies investigating different notice designs were conducted, for example, for SSL [16], browser security [33], and phishing warnings [13].

Mathur et al. [28] classified common *dark patterns* in web services. In their classification scheme the observed actions are described as "sneaking" (attempting to misrepresent user actions, or delay information that, if made available to users, they would likely object to), "misdirection" (using visuals, language, or emotion to steer users toward or away from making a particular choice), and "forced action" (forcing the user to do something additional in order to complete their task).

6 DISCUSSION

We conducted three experiments evaluating the effects of cookie consent notices’ position, choices, and content on people’s consent behavior. In the following we describe recommendations based on our findings and discuss limitations of our approach.

6.1 Recommendations

Our experiments investigated different notice positions, details of the choices offered, and the wording of cookie consent notices. Future guidelines for consent notices should consider the following recommendations:

Position. Experiment 1 showed that the position of a notice has a substantial impact on whether a website visitor engages with the notice. A dialog box in the lower left corner (on desktop) or the lower part of the screen (on mobile) significantly increases the chance that a user makes a consent decision. While we had expected higher interaction rates on mobile devices for this position since it is easy to reach with the thumb, we were surprised by the impact on desktop users, given the general wisdom that content in the top left receives the most attention in cultures with left-to-right writing. This result could be related to our partner website, like many websites, displaying a header which shifted content to lower parts of the screen. This experiment shows that the second most common notice position observed in practice, the top position (see Table 1), results in notices being ignored by users.

Choices. Our results from Experiment 2 showed that nudging (highlighting “Accept” buttons or pre-selecting checkboxes) substantially affects people’s acceptance of cookies, providing clear evidence for the interference of such dark patterns with people’s consent decisions. Given a binary choice, more visitors accepted cookies than declined them, which could be evidence for the adverse effects of consent bundling on consent decisions, which is not allowed under the GDPR. Surprisingly, rejection rates in the vendor- and cookie-based conditions were close to those in the binary condition, although visitors had to make five to six additional clicks to reach the same goal. This suggests that people who want to decline cookies are willing to expend extra effort.

Moreover, the survey answers show that participants think that no data is collected *unless* they make a decision, showing that privacy by default is the expected functionality, although this is not the current practice.

Text. While we did not see an effect in Experiment 3 from including a privacy policy link in the notice, we found that mentioning “cookies” made more users reject the data collection. The negative effect of mentioning cookies can very well be related to the fact that Internet users have in general a negative feeling about them [19, 22].

It is clear that the current ecosystem of mechanisms to prompt for user consent – with a plethora of combinations regarding the provided information, the granularity of user options, and how and if their choice is enforced – provides no real improvement for user privacy compared to pre-GDPR times. At the same time many things are still in flux, with regulators publishing differing guidelines on how to obtain consent, the online advertising industry developing and updating proposals for consent frameworks, and

legal and technical scholars evaluating them. While some claim [36] that many underlying principles of the online advertising industry are not compatible with the GDPR at all, the regulation so far has only partially affected how companies process personal data [48]. We hope that our results can inform future discussions, not only with recommendations for the design of consent notices. Given that at the moment very few users are willing to give consent to any form of processing of their personal data, we think that the business model of online behavioral advertising, which targets ads based on large amounts of personal data, should be challenged and alternative models like privacy-friendly contextual advertising or other ways of monetization for web services need to be developed.

6.2 Limitations

Our study has some potential limitations. First, our sample is biased as we conducted all experiments on a German-language e-commerce website whose visitors may not be representative of the general public. However, our partnership with this website gave us control over the notice implementation and access to a high number of unique visitors. We validated some of our results with data from Cookiebot which showed similar results (see Section 4.3.3). Overall it seems our sample is more inclined towards rejecting cookies. We have to assume that in general a higher percentage of users may allow cookies. Our field study did not allow us to collect more detailed information about visitors, such as their specific device, the size of the notice on the screen, or how long they stayed on the website, which could potentially have an effect on consent behavior.

Furthermore, many visitors did not interact with the notice at all and spent only a short period of time on the site. While this could be related to the notice, it is not unusual that most visitors leave a site after a few seconds. Liu et al. [25] showed that website dwell time has a negative aging effect. Users first skim a site to decide whether they will stay on it. Since we were not able to measure the exact time visitors stayed on the site, we included all users for whom the logged data indicated a fully loaded page, which results in a high number of “no action” visitors. From a legal perspective the time spent on the site does not affect the need to request consent. Our partner website also does not have user accounts. Past research has shown that visitors tend to underestimate the amount of personal data collected by websites on which they do not create an account and enter personal data [32]. This may cause them to underestimate the privacy implications of allowing cookie use, but we did not see evidence for this in the survey responses.

Responses to our voluntary survey are likely biased due to participants’ self-selection. Responses to the question about possible data collection suggest that participants had a good understanding of the technical background or an interest in privacy. Of the survey participants, 61 had previously interacted with our consent notices and 49 had not, showing that the results are only partially biased towards those who care about notices. We considered this bias when interpreting results.

7 CONCLUSION

We conducted the first large-scale field study on the effect of cookie consent notices on people’s consent behavior. Cookie notices have seen widespread adoption since the EU’s General Data Protection

Regulation went into effect in May 2018. Our findings show that a substantial amount of users are willing to engage with consent notices, especially those who want to opt out or do not want to opt in to cookie use. At the same time, position, offered choices, nudging, and wording substantially affect people’s consent behavior. Unfortunately, many current cookie notice implementations do not make use of the available design space, offering no meaningful choice to consumers. Our results further indicate that the GDPR’s principles of data protection by default and purposed-based consent would require websites to use consent notices that would actually lead to less than 0.1 % of users actively consenting to the use of third-party cookies.

ACKNOWLEDGMENTS

The authors would like to thank the owner of their partner website for allowing them to display different sets of consent notices on this site. Additional thanks to Yana Koval for her help with the implementation of the WordPress plugin and the classification of existing consent notices. This research was partially funded by the MKW-NRW Research Training Groups SecHuman and NERD.NRW, the German Research Foundation (DFG) within the framework of the Excellence Strategy of the Federal Government and the States (EXC 2092 CaSa – 39078197), and the National Science Foundation under grant agreement CNS-1330596.

REFERENCES

- [1] Alessandro Acquisti. 2009. Nudging Privacy: The Behavioral Economics of Personal Information. *IEEE Security & Privacy* 7, 6 (Dec. 2009), 82–85. <https://doi.org/10.1109/MSP.2009.163>
- [2] Alessandro Acquisti, Idris Adjerid, Rebecca Hunt Balebako, Laura Brandimarte, Lorrie Faith Cranor, Saranga Komanduri, Pedro Leon, Norman Sadeh, Florian Schaub, Manya Sleeper, Yang Wang, and Shomir Wilson. 2017. Nudges for Privacy and Security: Understanding and Assisting Users’ Choices Online. *Comput. Surveys* 50, 3 (Aug. 2017). <https://doi.org/10.2139/ssrn.2859227>
- [3] Alessandro Acquisti, Laura Brandimarte, and George Loewenstein. 2015. Privacy and human behavior in the age of information. *Science* 347, 6221 (Jan. 2015), 509–514. <https://doi.org/10.1126/science.aaa1465>
- [4] Alexa Internet, Inc. 2019. The top 500 sites on the Web. <https://www.alexa.com/topsites>
- [5] Article 29 Data Protection Working Party. 2016. *Cookie Sweep Combined Analysis – Report*. Technical Report 14/EN WP 229. European Commission, Brussels, Belgium.
- [6] Article 29 Data Protection Working Party. 2018. *Guidelines on consent under Regulation 2016/679*. Technical Report 17/EN WP259 rev.01. European Commission.
- [7] Sophie C. Boerman, Sanne Kruikemeier, and Frederik J. Zuiderveen Borgesius. 2018. Exploring Motivations for Online Privacy Protection Behavior: Insights From Panel Data. *Communication Research* 0, 0 (2018), 1–25. <https://doi.org/10.1177/0093650218800915>
- [8] Matt Burgess. 2018. The tyranny of GDPR popups and the websites failing to adapt. Retrieved April 22, 2019 from <https://www.wired.co.uk/article/gdpr-cookies-eprivacy-regulation-popups>
- [9] Virginio Cantoni, Marco Porta, Stefania Ricotti, and Francesca Zanin. 2013. Banner positioning in the masthead area of online newspapers: an eye tracking study. In *14th International Conference on Computer Systems and Technologies (CompSysTech ’13)*. ACM, New York, NY, USA, 145–152. <https://doi.org/10.1145/2516775.2516789>
- [10] Forbrukerrådet (Norwegian Consumer Council). 2018. *Deceived by Design – How tech companies use dark patterns to discourage us from exercising our rights to privacy*. Technical Report. Oslo, Norway.
- [11] Commission Nationale de l’Informatique et des Libertés (National Commission on Informatics and Liberty). 2018. Décision n° MED 2018-042 du 30 octobre 2018 mettant en demeure la société VECTAURY (Decision No. MED 2018-042 of 30 October 2018 giving notice to the company VECTAURY). Retrieved February 18, 2019 from <https://www.legifrance.gouv.fr/affichCnil.do?id=CNLTEXT000037594451>
- [12] Martin Degeling, Christine Utz, Christopher Lentzsch, Henry Hosseini, Florian Schaub, and Thorsten Holz. 2019. We Value Your Privacy ... Now Take Some Cookies: Measuring the GDPR’s Impact on Web Privacy. In *26th Annual Network and Distributed System Security Symposium (NDSS ’19)*. Internet Society.

- [13] Serge Egelman, Lorrie Faith Cranor, and Jason Hong. 2008. You’ve Been Warned: An Empirical Study of the Effectiveness of Web Browser Phishing Warnings. In *Conference on Human Factors in Computing Systems (CHI ’08)*. ACM, New York, NY, USA, 1065–1074. <https://doi.org/10.1145/1357054.1357219>
- [14] Interactive Advertising Bureau Europe. 2019. GDPR Transparency and Consent Framework. <https://iabtechlab.com/standards/gdpr-transparency-and-consent-framework/>. [Online; accessed 2 May 2019].
- [15] European Data Protection Board. 2019. *Opinion 5/2019 on the interplay between the ePrivacy Directive and the GDPR, in particular regarding the competence, tasks and powers of data protection authorities*. Technical Report 5/2019.
- [16] Adrienne Porter Felt, Alex Ainslie, Robert W. Reeder, Sunny Consolvo, Somas Thyagaraja, Helen Bettes, Alan ad Harris, and Jeff Grimes. 2015. Improving SSL Warnings: Comprehension and Adherence. In *33rd Annual ACM Conference on Human Factors in Computing Systems (CHI ’15)*. ACM, New York, NY, USA, 2893–2902. <https://doi.org/10.1145/2702123.2702442>
- [17] Vitaly Friedman. 2019. Privacy UX: Better Cookie Consent Experiences. Retrieved May 7, 2019 from <https://www.smashingmagazine.com/2019/04/privacy-ux-better-cookie-consent-experiences/>
- [18] Stacia Garlach and Daniel Suthers. 2018. ‘I’m supposed to see that?’ AdChoices Usability in the Mobile Environment. In *Hawaii International Conference on System Sciences*. University of Hawai’i at Mānoa, Honolulu, HI, USA, 3779–3788. <https://doi.org/10.24251/hicss.2018.476>
- [19] Vicki Ha, Kori Inkpen, Farah Al Shaar, and Lina Hdeib. 2006. An Examination of User Perception and Misconception of Internet Cookies. In *CHI ’06 Extended Abstracts on Human Factors in Computing Systems (CHI EA ’06)*. ACM, New York, NY, USA, 833–838. <https://doi.org/10.1145/1125451.1125615>
- [20] Hana Habib, Yixin Zou, Aditi Jannu, Neha Sridhar, Chelse Swoopes, Alessandro Acquisti, Lorrie Faith Cranor, Norman Sadeh, and Florian Schaub. 2019. An Empirical Analysis of Data Deletion and Opt-Out Choices on 150 Websites. In *Fifteenth Symposium On Usable Privacy and Security (SOUPS 2019)*. USENIX Association, 387–406. <https://www.usenix.org/conference/soups2019/presentation/habib>
- [21] Daniel Kladnik. 2019. I don’t care about cookies 3.0.0. <https://www.i-dont-care-about-cookies.eu/>. [Online; accessed 2 May 2019].
- [22] Oksana Kulyk, Annika Hilt, Nina Gerber, and Melanie Volkamer. 2018. “This Website Uses Cookies”: Users’ Perceptions and Reactions to the Cookie Disclaimer. In *3rd European Workshop on Usable Security (EuroSec 2018)*. London, England, 11.
- [23] Oksana Kulyk, Peter Mayer, Oliver Käfer, and Melanie Volkamer. 2018. A Concept and Evaluation of Usable and Fine-Grained Privacy-Friendly Cookie Settings Interface. In *17th IEEE International Conference On Trust, Security And Privacy In Computing And Communications (TrustCom 2018)*. IEEE, Piscataway, NJ, USA.
- [24] Pedro Leon, Blase Ur, Richard Shay, Yang Wang, Rebecca Balebako, and Lorrie Cranor. 2012. Why Johnny can’t opt out: a usability evaluation of tools to limit online behavioral advertising. In *Conference on Human Factors in Computing Systems (CHI ’12)*. ACM, New York, NY, USA, 589–598. <https://doi.org/10.1145/2207676.2207759>
- [25] Chao Liu, Ryan W. White, and Susan Dumais. 2010. Understanding Web Browsing Behaviors Through Weibull Analysis of Dwell Time. In *33rd International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR ’10)*. ACM, New York, NY, USA, 379–386. <https://doi.org/10.1145/1835449.1835513>
- [26] Manafactory. 2019. Ginger – EU Cookie Law. <https://wordpress.org/plugins/ginger/>. [Online; accessed 22 August 2019].
- [27] Kirsten Martin. 2016. Do Privacy Notices Matter? Comparing the Impact of Violating Formal Privacy Notices and Informal Privacy Norms on Consumer Trust Online. *The Journal of Legal Studies* 45, S2 (June 2016), S191–S215. <https://doi.org/10.1086/688488>
- [28] Arunesh Mathur, Gunes Acar, Michael Friedman, Elena Lucherini, Jonathan Mayer, and Marsh Chetty. 2019. Dark Patterns at Scale: Findings from a Crawl of 11K Shopping Websites. (2019). arXiv:1907.07032
- [29] Jonathan R. Mayer and John C. Mitchell. 2012. Third-Party Web Tracking: Policy and Technology. In *2012 IEEE Symposium on Security and Privacy (SP ’12)*. IEEE Computer Society, Washington, DC, USA, 413–427. <https://doi.org/10.1109/SP.2012.47>
- [30] Aleecia M. McDonald and Lorrie Faith Cranor. 2010. Americans’ Attitudes About Internet Behavioral Advertising Practices. In *9th Annual ACM Workshop on Privacy in the Electronic Society (WPES ’10)*. ACM, New York, NY, USA, 63–72. <https://doi.org/10.1145/1866919.1866929>
- [31] Mike O’Neill. 2018. Do Not Track and the GDPR. Retrieved May 15, 2019 from <https://www.w3.org/blog/2018/06/do-not-track-and-the-gdpr/>
- [32] Ashwini Rao, Florian Schaub, Norman Sadeh, Alessandro Acquisti, and Ruogo Kang. 2016. Expecting the Unexpected: Understanding Mismatched Privacy Expectations Online. In *Twelfth Symposium On Usable Privacy and Security (SOUPS ’16)*. USENIX Association, 77–96. <https://www.usenix.org/conference/soups2016/technical-sessions/presentation/rao>
- [33] Robert W. Reeder, Adrienne Porter Felt, Sunny Consolvo, Nathan Malkin, Christopher Thompson, and Serge Egelman. 2018. An Experience Sampling Study of User Reactions to Browser Warnings in the Field. In *Conference on*

- Human Factors in Computing Systems (CHI '18)*. ACM, New York, NY, USA. <https://doi.org/10.1145/3173574.3174086>
- [34] Johnny Ryan. 2017. Research result: what percentage will consent to tracking for... <https://pagefair.com/blog/2017/new-research-how-many-consent-to-tracking/>
- [35] Johnny Ryan. 2017. *The state of the blocked web – 2017 Global Adblock Report*. Technical Report. PageFair. Retrieved May 8, 2019 from <https://pagefair.com/downloads/2017/01/PageFair-2017-Adblock-Report.pdf>
- [36] Johnny Ryan. 2018. French regulator shows deep flaws in IAB’s consent framework and RTB. Retrieved May 8, 2019 from <https://brave.com/cnil-consent-rtb/>
- [37] Johnny Ryan. 2019. Formal GDPR complaint against IAB Europe’s “cookie wall” and GDPR consent guidance. Retrieved May 10, 2019 from <https://brave.com/iab-cookie-wall/>
- [38] Iskander Sanchez-Rola, Matteo Dell’Amico, Platon Kotzias, Davide Balzarotti, Leyla Bilge, Pierre-Antoine Vervier, and Igor Santos. 2019. Can I Opt Out Yet? GDPR and the Global Illusion of Cookie Control. In *ACM ASIA Conference on Computer and Communications Security (AsiaCCS '19)*. ACM, New York, NY, USA. <https://doi.org/10.1145/3321705.3329806>
- [39] Florian Schaub, Rebecca Balebako, Adam L. Durity, and Lorrie Faith Cranor. 2015. A Design Space for Effective Privacy Notices. In *Eleventh Symposium On Usable Privacy and Security (SOUPS '15)*. The USENIX Association, Ottawa, 1–17. <https://doi.org/10.1145/567752.567774>
- [40] Mario Silic. 2016. Understanding Colour Impact on Warning Messages: Evidence from US and India. In *2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems (CHI EA '16)*. ACM, New York, NY, USA, 2954–2960. <https://doi.org/10.1145/2851581.2892276>
- [41] Jannick Sørensen and Sokol Kosta. 2019. Before and After GDPR: The Changes in Third Party Presence at Public and Private European Websites. In *The 2019 World Wide Web Conference (WWW '19)*. ACM, New York, NY, USA, 1590–1600. <https://doi.org/10.1145/3308558.3313524>
- [42] State of California Legislative Counsel. 2018. Assembly Bill No. 375 – Chapter 55.
- [43] Richard H. Thaler and Cass R. Sunstein. 2009. *Nudge: Improving Decisions About Health, Wealth, and Happiness*. Penguin Books, New York, NY, USA.
- [44] The European Parliament and the Council of the European Union. 2002. Directive 2002/58/EC of the European Parliament and of the Council of 12 July 2002 concerning the processing of personal data and the protection of privacy in the electronic communications sector. Official Journal of the European Communities.
- [45] The European Parliament and the Council of the European Union. 2009. Directive 2009/136/EC of the European Parliament and of the Council of 25 November 2009 amending Directive 2002/22/EC, Directive 2002/58/EC and Regulation (EC) No 2006/2004. Official Journal of the European Union, L 337/11.
- [46] The European Parliament and the Council of the European Union. 2016. 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). Official Journal of the European Union, L 119/1.
- [47] Joseph Turow, Michael Hennessy, and Nora Draper. 2018. Persistent Misperceptions: Americans’ Misplaced Confidence in Privacy Policies, 2003–2015. *Journal of Broadcasting & Electronic Media* 62, 3 (July 2018), 461–478. <https://doi.org/10.1080/08838151.2018.1451867>
- [48] Tobias Urban, Martin Degeling, Thorsten Holz, and Norbert Pohlmann. 2019. Perspectives on Transparency Tools for Online Advertising. In *35th Annual Computer Security Applications Conference (ACSAC)*. ACM, San Juan, 14.
- [49] Rob van Eijk, Hadi Asghari, Philipp Winter, and Arvind Narayanan. 2019. The Impact of User Location on Cookie Notices (Inside and Outside of the European Union). In *Workshop on Technology and Consumer Protection (ConPro '19)*. IEEE.
- [50] Markus Weinmann, Christoph Schneider, and Jan vom Brocke. 2016. Digital Nudging. *Business & Information Systems Engineering* 58, 6 (Dec. 2016), 433–436. <https://doi.org/10.1007/s12599-016-0453-1>

A TIMING IN EXPERIMENT 2

Table 3: Average time in seconds until users submitted decision in Experiment 2, if decision was made within the first three minutes

Banner	Type	# Users	Mean	Median	SD
No Option		4174	6.51	4	15.55
Confirmation	non-nudging	2984	10.65	5	51.25
	nudging	3634	9.11	4	37.78
Binary	non-nudging	4134	15.36	4	72.47
	nudging	4097	13.51	4	75.59
Category	non-nudging	2523	17.93	8	87.16
	nudging	2798	13.98	7	64.01
Vendor	non-nudging	2346	13.76	8	42.38
	nudging	2741	21.18	7	115.26

B SURVEY AND RESPONSES

^R indicates answers displayed in random order. All questions and answers were translated from German as true to the original as possible.

Motivation for Interacting With the Cookie Consent Notice

Q1-clicked^a: You just clicked the cookie consent notice^b on the website [WEBSITE_NAME]. Which of the following statements describe your motivation to click the notice? I clicked the cookie consent notice ... [multiple choice]

	Exp. 1	Exp. 2	Exp. 3	Total	%
... to protect me from dangers from the Internet. ^R	0	3	3	6	9.8 %
... to protect my privacy on the Internet. ^R	0	5	6	11	18.0 %
... because the website does not work otherwise. ^R	2	11	3	16	26.2 %
... to see less ads. ^R	1	1	3	5	8.2 %
... out of habit. ^R	1	10	2	13	21.3 %
... because the notice distracts me from viewing the website. ^R	6	25	13	44	72.1 %
Other: [free text]	0	0	1	1	1.6 %
I do not know why I clicked the notice.	1	1	1	3	4.9 %
I prefer not to answer.	0	0	0	0	0 %
# Answers	11	56	32	99	162.3 %
# Participants	8	34	19	61	100.0 %

Q1-notclicked:^a You did not click the cookie consent notice^b on the website [WEBSITE_NAME]. Which of the following statements describe your motivation to not click the notice? I did not click the cookie consent notice ... [multiple choice]

	Exp. 1	Exp. 2	Exp. 3	Total	%
... because I have not noticed it. ^R	4	11	5	20	40.8 %
... because it did not offer enough choices. ^R	0	0	3	3	6.1 %
... because I do not know what happens if I click the notice. ^R	1	6	4	11	22.4 %
... because I think that my selection does not have any effect. ^R	1	4	4	9	18.4 %
... because I do not know what cookies are. ^R	0	2	0	2	4.1 %
... because I do not care which cookies the website uses. ^{Rc}	1	3	2	6	12.2 %
... Other: [free text]	1	10	2	13	26.5 %
... I do not know why I did not click the cookie consent notice.	1	0	0	1	2.0 %
... I prefer not to answer.	0	2	0	2	4.1 %
# Answers	9	38	20	67	136.7 %
# Participants	8	26	15	49	100.0 %

^a Q1-clicked and Q1-notclicked were only displayed to participants who clicked / did not click the notice, respectively.

^b In Experiment 3, “cookie consent notice” was changed to “privacy notice” in the conditions Non-Technical-PP Link and Non-Technical-No PP Link.

^c In Experiment 3, this answer was changed to “because I do not know what data this is about” in the conditions Non-Technical-PP Link and Non-Technical-No PP Link.

Expectation of the Website’s Data Collection

Q2: What do you think – what data does the website [WEBSITE_NAME] collect about you when you access the website?

	Exp. 1	Exp. 2	Exp. 3	Total	%
The posts I am reading on the website. ^R	10	40	17	67	60.9 %
My residence. ^R	6	14	7	27	24.5 %
The links I click on the website. ^R	14	45	27	86	78.2 %
My IP address. ^R	11	39	22	72	65.5 %
The device I am using to access the website. ^R	10	36	19	65	59.1 %
The website does not collect any data about its visitors. ^R	0	4	1	5	4.5 %
My name. ^R	2	9	3	14	12.7 %
Other websites I visit besides [WEBSITE_NAME]. ^R	5	17	10	32	29.1 %
Other: [free text]	3	2	1	6	5.5 %
I prefer not to answer.	0	0	0	0	0 %
# Answers	61	206	107	374	340.0 %
# Participants	16	60	34	110	100.0 %

Perception of the Cookie Consent Notice Displayed to the Participant

This is the cookie consent notice^a the website has shown you. [IMAGE]
Please rate the following statements about this notice.

Q3: I think the number of choices offered by the above cookie consent notice^b is ...

	Exp. 1				Exp. 2				Exp. 3					
	BIN-S1 ^c	NOP	CON-NN	CON-NU	BIN-NN	BIN-NU	CAT-NN	CAT-NU	VEN-NN	VEN-NU	TE-PP	TE-NP	NT-PP	NT-NP
... too low	9	3	3	5	3	1	1	2	1	2	1	1	0	1
... just right	7	1	0	3	7	3	2	3	1	2	4	8	6	6
... too high	0	0	1	1	0	0	3	2	0	3	2	0	3	0
... No answer	0	2	0	1	2	0	1	0	1	0	0	0	2	0
Total	16	6	4	10	12	4	7	7	3	7	7	9	11	7

Q4: The above cookie consent notice^a allows me to control the website's behavior.

	Exp. 1				Exp. 2				Exp. 3					
	BIN-S1	NOP	CON-NN	CON-NU	BIN-NN	BIN-NU	CAT-NN	CAT-NU	VEN-NN	VEN-NU	TE-PP	TE-NP	NT-PP	NT-NP
Strongly disagree	6	3	3	0	2	0	1	1	0	1	1	1	0	0
Somewhat disagree	3	2	0	3	3	2	2	1	1	3	2	0	3	0
Neutral	6	0	1	4	3	0	0	1	1	1	1	1	4	2
Somewhat agree	1	1	0	2	4	1	4	3	1	1	1	4	4	5
Strongly agree	0	0	0	0	0	1	0	1	0	1	2	2	0	0
No answer	0	0	0	1	0	0	0	0	0	0	0	1	0	0
Total	16	6	4	10	12	4	7	7	3	7	7	9	11	7

Q5^b: I think the decision which option to select in the cookie consent notice^a is ...

	Exp. 2				Exp. 3			
	CAT-NN	CAT-NU	VEN-NN	VEN-NU	TE-PP	TE-NP	NT-PP	NT-NP
... very easy	2	0	1	1	4	3	4	1
... easy	0	2	1	0	0	2	2	2
... neither easy nor hard	2	2	0	2	2	2	5	4
... hard	2	2	1	2	1	1	0	0
... very hard	1	1	0	2	0	0	0	0
No answer	0	0	0	0	0	1	0	0
Total	7	7	3	7	7	9	11	7

^a In Experiment 3, "cookie consent notice" was changed to "privacy notice" in the conditions Non-Technical-PP Link and Non-Technical-No PP Link.

^b Q5 was only shown to participants who had seen a category- or vendor-based notice on the website.

^c BIN-S1 = the binary notice shown at six different positions in Experiment 1; NOP = no option; CON = confirmation; BIN = binary; CAT = categories; VEN = vendors; NN = non-nudging; NU = nudging; TE = technical; NT = non-technical; PP = privacy policy link; NP = no privacy policy link.

Perception of the Cookie Consent Notice Displayed to the Participant (cont.)

Q6^a Please explain your answer to the previous question. [free text answers, coded by two authors]

Code	Explanation	Exp. 2	Exp. 3	Total	%
Transparent	The participant considers the consent notice to be transparent.	1	5	6	15.8%
Privacy	The participant's preferences are privacy-focused, i. e., the least invasive option is chosen.	2	5	7	18.4%
Options clear	The options offered by the consent notice are considered clear / easy to understand.	0	3	3	7.9%
Options unclear	The options offered by the consent notice are considered unclear / not easy to understand.	4	2	6	15.8%
Notice clear	The participant expressed that the mechanism was clear but did not specify which part.	1	3	4	10.5%
Notice unclear	The participant expressed that the mechanism was unclear but did not specify which part.	2	0	2	5.3%
Too complicated	The consent notice was considered too complex.	4	1	5	13.2%
Don't care	The participant stated they did not care which cookies the website used.	3	0	3	7.9%
Other		4	2	6	15.8%
# Participants		60	34	110	100.0%

^a Q6 was only shown to participants who had seen a category- oder vendor-based notice on the website.

General Understanding of Cookie Consent Notices

This is another cookie consent notice. [Image of the binary notice in Figure 1 (a) (bb)]

Q7: What do you think happens when you click "Decline"? [free text answers, coded by two authors]

Code	Explanation	Exp. 1	Exp. 2	Exp. 3	Total	%
Site blocked	The content of the website cannot be accessed at all.	6	13	9	28	29.8%
Functionality limited	The content of the website can be viewed, but some parts may not work.	2	10	5	17	18.1%
Site accessible	The content of the website can be accessed.	0	3	1	4	4.3%
No data collected	The website visitor's personal data is not collected or processed.	2	4	5	11	11.7%
No cookies set	The website does not store any cookies in the visitor's browser.	1	8	3	12	12.8%
Less ads	The website displays less or no ads.	0	3	2	5	5.3%
Notice	The participants only mentions effects regarding the consent notice.	0	2	3	5	5.3%
No change	Declining cookies does not have any effect.	4	7	1	12	12.8%
Don't know		2	0	1	3	3.2%
Other		0	2	4	6	6.4%
# Participants		15	51	28	94	100.0%

Q8: What do you think happens when you click "Accept"? [free text answers, coded by two authors]

Code	Explanation	Exp. 1	Exp. 2	Exp. 3	Total	%
Data collected	The participant's personal data is collected and/or processed.	9	10	10	29	30.9%
Cookies stored	Cookies are stored in the user's browser.	4	9	6	19	20.1%
Site accessible	The content of the website can be accessed.	0	16	5	21	22.3%
Notice	The participants only mentions effects regarding the consent notice.	0	3	2	5	5.3%
Ads	The participant is subject to advertising.	6	11	6	23	24.5%
Profiling	The participant's personal data is used to create a profile of their interests.	5	8	6	19	20.2%
Other purposes	The participant's personal data is used for other purposes.	2	0	2	4	4.3%
No change	Clicking "Accept" does not have any effect.	0	4	3	7	7.4%
Don't know		0	3	0	3	3.2%
Other		0	3	1	4	4.3%
# Participants		15	51	28	94	100.0%