

Dominique Machuletz\* and Rainer Böhme

# Multiple Purposes, Multiple Problems: A User Study of Consent Dialogs after GDPR

**Abstract:** The European Union’s General Data Protection Regulation (GDPR) requires websites to ask for consent to the use of cookies for *specific purposes*. This enlarges the relevant design space for consent dialogs. Websites could try to maximize click-through rates and positive consent decision, even at the risk of users agreeing to more purposes than intended. We evaluate a practice observed on popular websites by conducting an experiment with one control and two treatment groups ( $N = 150$  university students in two countries). We hypothesize that users’ consent decision is influenced by (1) the number of options, connecting to the theory of choice proliferation, and (2) the presence of a highlighted default button (“select all”), connecting to theories of social norms and deception in consumer research. The results show that participants who see a default button accept cookies for more purposes than the control group, while being less able to correctly recall their choice. After being reminded of their choice, they regret it more often and perceive the consent dialog as more deceptive than the control group. Whether users are presented one or three purposes has no significant effect on their decisions and perceptions. We discuss the results and outline policy implications.

**Keywords:** web privacy, user study, consent, cookies, controlled experiment, choice proliferation, deception, privacy paradox, privacy by design, dark patterns

DOI 10.2478/popets-2020-0037

Received 2019-08-31; revised 2019-12-15; accepted 2019-12-16.

## 1 Introduction

The European Union’s General Data Protection Regulation (GDPR) [1] came into force in May 2018. It stipulates that *data controllers* (e. g., website operators)

must have a legal basis for the collection and processing of *personal data*. One legal basis is *consent*: *data subjects* (users) agree to the data processing for *specific purposes*. While these requirements are not new,<sup>1</sup> the GDPR’s threat of sanctions and more effective enforcement led many website operators to rethink their cookie practices, or at least ensure compliance by obtaining consent before using cookies for purposes that are not covered by other legal bases [6].

Web cookies are key–value pairs stored on the client device for purposes ranging from session tracking, user recognition, counting unique users, third-party tracking to profiling and targeted advertising [7]. As every cookie can in principle serve many purposes at the same time, and *necessary* cookies not carrying any personal data do not require consent, a user generally cannot verify if a website complies with the agreed purposes.

Common methods for asking web users to decide on the cookie settings are pop-up banners or dialogs that appear at the beginning of each user’s first visit of a website. They typically include a notice on the data collection that asks users whether they consent to (parts of) the practices. Systematic longitudinal measurements are lacking, but one study reports that 62% of the websites in its sample used such notices in June 2018 [8]. It also shows that the implementation—specifically, the granularity of control offered to users—differs between websites. The authors of [8] conjecture that many cookie banners and dialogs are not very usable, and they provide early evidence from a series of field experiments with several variants of cookie banners placed on one website [9].

Independently, in November 2018, we noticed that some cookie consent dialogs seem to be designed to nudge users into accepting all displayed purposes. (This observation is meanwhile documented in the literature [e. g., 10]). It is understandable that the industry finds cookie banners disadvantageous as they add fric-

\*Corresponding Author: Dominique Machuletz: Independent, E-mail: mail@machuletz.com. Work carried out while at the University of Münster, Germany.

Rainer Böhme: University of Innsbruck, Austria, E-mail: rainer.boehme@uibk.ac.at

<sup>1</sup> The principles of consent and purpose binding appear in data protection laws since the 1970s. The specific case for web cookies was harmonized in the EU through the 2009 update of the ePrivacy Directive [2–4], but respected by only one in two websites, according to a recent measurement study [5].

tion to the user experience and might limit the ability to track users on and across sites. Hence, there is ample business interest in minimizing friction and maximizing positive consent decisions by optimizing interface design. Common design elements in the dialogs we observed (see Figure 1 for examples) are checkboxes for several purposes of data processing as well as buttons to either select all purposes at once or to confirm the manual selection before accessing the website.

We identify two features that might compromise usability. First, the highlighted button automatically accepts all purposes, regardless of whether any checkboxes have (or have not) been selected before the button is clicked. This button does not increase the users' choice options, but might rather “trick” them into accepting all purposes without actively selecting them. Second, the number of selectable purposes may influence users' choice as former studies in the field of psychology revealed that a high number of alternatives has adverse effects on individuals' decision making [11, 12]. This phenomenon has also been demonstrated in the context of privacy settings [13].

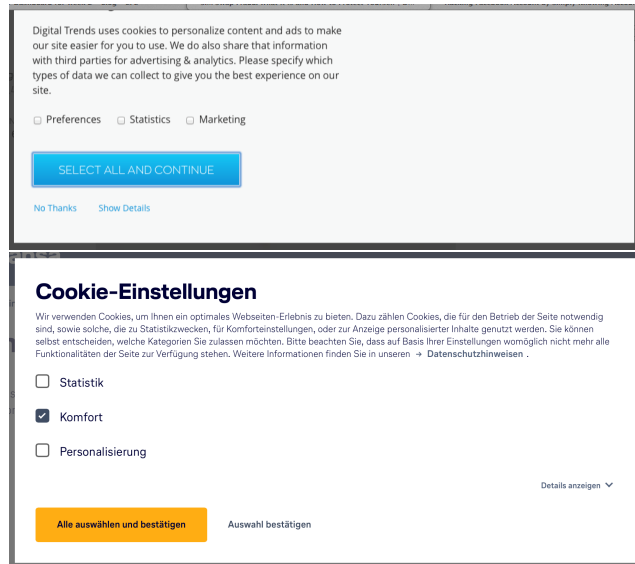
These considerations call for a user study, which we have carried out in the form of a controlled classroom experiment and report in this paper. Our general research question is:

*How do users react to design features of multi-purpose consent dialogs on the web in terms of actual behavior and stated perceptions?*

The rest of this paper is structured as follows. First, we review the literature on consent dialog designs in Section 2. Section 3 recalls the theoretical background on choice proliferation and deception, from which we derive our hypotheses. The instrument and the administration of the controlled experiment is described in Section 4. The results of our hypothesis tests (Section 5) precede the discussion of our findings (Section 6). We conclude with some recommendations for interface design and policy development in Section 7.

## 2 Background

We first summarize the legal requirements for GDPR-compliant consent dialogs in Section 2.1, before we review the literature on engineering solutions for specifying privacy preferences (with emphasis on the purpose) in Section 2.2.



**Fig. 1.** Examples of real-world cookie consent dialogs that motivated this study: a US technology news website (top) and a German airline website (bottom). Both dialogs are blocking and all items are unchecked initially (opt-in).

### 2.1 Legal Requirements for Consent Dialogs

Article 7 of the GDPR describes the requirements of legitimate consent: it needs to be (1) freely given, (2) unambiguous, (3) informed, and (4) withdrawable at any time [1]. In the event of a dispute, the data controller must prove that the subject has truly given consent to the processing practices [14]. Specifically, consent must be communicated “by a statement” or a “clear affirmative action” [1]. Regarding the clearness of this action, ticking a checkbox on a website is considered an acceptable form, while passiveness or predefined default settings that are not actively declined by the subject do not qualify as consent decisions. The European Court of Justice has just reconfirmed this interpretation [15].

If personal data is collected for more than one purpose, data subjects need to be informed and provided with distinct opt-in choices for every purpose [1]. Besides stating these principles, the GDPR intentionally does not specify any design template or rules, and thus leaves the exploration of the design space for consent dialogs to the market participants.

For the specific case of web cookies, the market has adopted a rough classification of purposes into strictly *necessary* (which presumably do not require consent), *preferences*, *statistics*, and *marketing* (which includes third-party tracking) [16][Fig. 4(d) of 8]. This mirrors

the approach taken in a user survey by Ackerman et al. as early as in 1999 [17]. The authors distinguish between cookies for “customized service”, “customized advertising”, and “customized advertising across many websites”. They report a decreasing willingness to agree, from 96% to 77% for users classified as “marginally concerned” about privacy, and from 43% to 14% for so-called “privacy fundamentalists” in a sample of 381 US internet users (Fig. 3 of [17]). While the former classification is implemented in popular content management systems, it is by no means the only way of defining purposes. As a result, website operators who can afford specialized lawyers enjoy more freedom in the design of consent dialogs. Others follow common practices in order to minimize legal uncertainty, or to comply with the terms of services of third parties who provide content or code to embed (e.g., Google Analytics). The bulk of the burden lands on privacy-aware users, who need to understand and navigate each site’s specific model.

## 2.2 Technical Solutions for Seeking Consent

Researchers have studied ways to effectively inform users about privacy policies and seek their consent to data processing long before the GDPR. For example, a CHI paper from 2001 provides design recommendations for cookie consent dialogs after evaluating design changes of the then popular browsers over time [18]. The authors criticize browsers in which users had to invest great effort when searching for an alternative to the “accept all cookies” default setting. Consent dialogs are specific forms of privacy notices, a topic so profoundly researched that Schaub et al. [19] saw the need to systematize the literature. According to their proposed taxonomy, the design space can be divided along the dimensions *timing*, *channel*, *modality*, and *control*. In the following, we use this terminology when applicable.

Bergmann [20] addresses the problem of complex and incomprehensible privacy choices. The author suggests a design for generic predefined privacy settings (*timing: at setup*) that are summarized in a limited number of categories. He defines four privacy profiles that differ in the acceptance level of transmission and processing of personal data. The suggested solution aims at decreasing the user’s cognitive effort when selecting suitable privacy settings, but we are not aware of any empirical study to evaluate this approach.

Pettersson et al. [21] discuss a similar design with predefined settings. They suggest the adoption of a pri-

vacuity management system that asks users for consent before transmitting their personal data (*timing: at setup*). Moreover, users’ acceptance of data processing practices can be configured in advance and apply to future website visits. However, the authors point out that designing consent forms that are applicable to a large number of different websites is a complex task. It might require compromises on usability as many different settings need to be offered by the system. More specifically, Pettersson et al. [22] propose design paradigms that include suggestions for consent dialogs. Incorporating recommendations by data protection commissioners and legal experts as well as standards established in the PISA project [23], the authors present a dialog window with several mandatory and optional fields, an expandable privacy notice, information about the data recipients, and an “I agree” button. They also propose methods to overcome habituation by, for instance, using drag-and-drop actions for consent. The authors qualitatively evaluate their usability tests and find that some users did not fully trust the privacy management system.

In a follow-up study, Bergman [24] empirically explores how to successfully communicate websites’ privacy policies to users. Specifically, he compared a conventional interface for online forms to an extended version with additional explanations of privacy information that pops up in tooltips (so-called “privacy bars”) while filling the form (*timing: just-in-time*). He finds that participants who saw the extended version were significantly more likely to be aware of the policy than the control group. But he did not measure the cost of this sophistication in terms of response time or frictions to usability. Moreover, the screenshots of the extended dialog (Fig. 2 of [24]) bears a risk of information overload. Finally, as the dialog was only tested on desktop computers, it remains unclear how this information can be perceived on small mobile displays.

Tiny displays raise the need for non-interactive forms of privacy preference negotiations. An established (but meanwhile discontinued) standard for expressing privacy preferences on the web is P3P. The standard lets websites communicate their privacy policies in machine-readable XML format (*modality: machine-readable*). Each XML element represents a component, such as the type of data, the purpose for data collection, and third party recipients [25]. A language called Appel has been developed for enabling users to express their privacy preference through predefined rules (*timing: at setup*), so that automated privacy decisions can be based on the user’s specific settings [26].

A recent approach towards facilitating informed and GDPR-compliant user consent is proposed by Ulbricht and Pallas [27]. The authors present a privacy preference language, called YaPPL, that is targeted on consent for data practices on the Internet of Things (IoT). For the development, they analyze legal requirements for consent and transform them to technical standards that suit IoT devices (*modality: machine-readable, channel: primary or secondary*). The language is prototypically tested in real-world IoT applications. The authors hope that the underlying approach of YaPPL will also be implemented in IoT applications that do not have to meet the standards of GDPR, but require a technical representation of users' privacy preferences.

Dissatisfied by the observation that many users tend to ignore notices with privacy impact [28, 29], perceive them as a threat to their privacy [30], and have been habituated to “click away” consent dialogs [31], several researchers investigated how to design more effective privacy notices. For instance, Felt et al. [32] propose design guidelines that aid mobile application developers in appropriately asking for permissions. They find that more than half of all permission requests can be automated while 16% require consent dialogs. By minimizing the number of runtime consent dialogs, the authors intend to decrease the required user attention. While technical permissions differ from legal purposes in several respects, it is conceivable that similar effects also apply to purposes. To our knowledge, this link is still unexplored.

Most closely related to the present work is the concurrent effort by Utz et al. [9], which draws on data from a field experiment exploring the design space for cookie banners. Both works share the experimental method, inquiry period (Q1/2019), and language (German). Some of their treatments and findings relate to our research questions. We shall comment on specific similarities and differences where it applies. The most salient differences between our colleagues' and this work are the mode of data collection (field vs lab), the type of cookie notice studied (non-blocking banner vs blocking dialog), the emphasis of the analysis (behavioral traces vs stated attitudes and beliefs), and the context of scientific discovery (inductive vs deductive). Both works leave many questions open, indicating that we are at the beginning of a relevant and potentially fruitful strand of research.

The works discussed in this section are selected pieces of the literature. They are representative in that the field focuses on technical and human aspects in many facets, but (with a few exceptions) it largely ignores economic interests [33]. In practice, we must expect that businesses use the flexibility in the design of

consent dialogs for their own interest by maximizing data disclosure instead of helping users to make privacy-conscious decisions.

## 3 Theory

User studies integrate better into the body of knowledge (and, arguably, generalize better), if the hypothesized causal links are derived from established theory. Therefore, we revisit relevant theories for explaining the effect of the two characteristic components in the consent dialogs inspiring this work (Fig. 1). Specifically, we review choice proliferation in Section 3.1 to reason about the number of purposes, and social norms in combination with deception in Section 3.2 to predict the effect of the default button. Then, we formulate our hypotheses in Section 3.3.

### 3.1 Choice Proliferation

Choice proliferation is a line of research in psychology that analyzes the influence of an increasing number of alternative choices on the human decision-making process. The phenomenon that more options result in negative effects, such as dissatisfaction, has mainly been studied in a marketing context [34, 35] and is sometimes referred to as “too much choice”, “tyranny of choice”, or “choice overload”.

As pointed out by Johnson et al. [35], two main aspects have to be considered when evaluating the number of choices offered. On the one hand, a high number of alternatives increases the cognitive load while causing individuals to feel stressed, overwhelmed, and more likely to regret one's decision [11, 12]. On the other hand, the likelihood that the choice suits the individual's preferences increases when more options are given. Thus, the practical challenge is to find the right balance.

A few works investigate the effect of increasing privacy choices on users' decision making. Korff and Böhme [13] experimentally study the influence of choice amount and choice structure in the context of privacy preferences on a fictitious business networking website. They find that participants who were confronted with a larger number of privacy settings to choose from were less satisfied with their choice and experienced more regret. The works by Knijnenburg et al. [36] and Tang et al. [37] investigate the number of privacy choices in the context of mobile location sharing. Both studies find

that the structure of presented choices significantly impacts users' tendency to disclose personal data. Utz et al. vary the number of choices of a cookie banner in their field study, but they neither relate this treatment to choice proliferation nor collect the relevant dependent variables. Since the instrument confounds the number of options with their type (5 categories, with one pre-selected, and 6 vendors; see Fig. 1 (d) and (e) in [9]), it is not easy to interpret the results. Krasnova et al. discuss the effect of an increasing amount of information items in mobile applications' permission requests [38]. The results of their experiment show that users tend to be more concerned if the permission request asks for more information items. This aspect of choice proliferation seems to be specific to privacy, because options in privacy dialogs often remind users of threats. This is rarely the case in the marketing literature on choice proliferation, where the typical study varies the number of forms of a retail product (e. g., flavors of jam).

Like almost any social science theory, choice proliferation is not undisputed. Critics argue that more options can lead to higher satisfaction since one's individual needs can be matched more precisely [39]. Moreover, more choice enables easier comparison of differences, which leads to more confident decision making [40].

Broadly related to the number of options is the number of occasions for privacy decisions. Böhme and Grossklags [41] discuss the averse effects of escalating too many decisions to users. They postulate that only the most important decisions should be made by users, so that they do not get habituated to ignore notices as a consequence of too high complexity. Several empirical studies support this interpretation. For example, the null result in an experiment on more or less verbose variants of the well-known consent dialog of Facebook Connect is attributed to habituated ignorance [42].

Our study connects to the literature on choice proliferation by experimentally varying the number of purposes. We adapt established constructs to measure perceived task difficulty and regret.

### 3.2 Deception and Social Norms

The concept of deception is often described as being misled due to unfair practices and can occur in many contexts when interests of different parties collide [43]. Deception has been studied in several areas such as marketing [44–47] and organizational research [48–50]. A deceptive practice is being conducted if the targeted individual receives false information that lead to false im-

pressions of a situation. Such false impressions may trigger decisions or opinions that would have been formed in a different way without the deceiving act.

However, deception is not always based on lying, as it may also comprise purposeful evocation of specific actions by the targeted party; for instance, by increasing the complexity of information, or by making use of behavioral clues or clue patterns. A study by Nochenson and Grossklags [46] investigates how users of web shops are tricked into falling for post-transaction marketing tactics due to specific design elements in notices. In an experiment, they test the purchasing behavior of more than 500 users and find that above 40% signed up unintentionally for an extra service with costs. The authors find that opt-in and opt-out default buttons significantly impact the users' tendency to fall for the trick.

Citing usability guidelines [51], Böhme and Köpsell [31] underline that the default option should include the most frequently selected settings so that inexperienced users can be assisted by the decision of the majority. In this sense, default buttons can be interpreted as a descriptive social norm. However, as highlighted in a study on default privacy settings on social media websites, the preset or default options are often very disclosing and might not reflect the majority of users' privacy preferences [52]. It seems that the default button has mutated from a usability tool that improves efficiency when selecting the typical choice to a strategic tool that supports the interests of the system designer.

For several decades, scholars in the behavioral sciences have identified and quantified cognitive and social effects, some of which cause successful persuasion or deception [53]. A shared objective in these disciplines was to isolate effects, which required substantial effort given that stimuli to human subjects often confound many factors. By contrast, the recent literature that criticizes the deliberate exploitation of these biases in favor of the designer typically looks at bundles of features as they appear in practice [54]. The term “dark pattern” [55], coined in 2010, classifies designs that trick users into making decisions they do not mean to make. Bösch et al. [56] were among the first to systematize dark patterns commonly adopted for privacy invasions. For instance, users typically do not read privacy notices completely [57] and often intuitively accept the presented conditions. This behavior can be exploited by hiding undesirable terms in privacy notices. Mathur et al. [58] structure common characteristics of dark patterns along five dimensions: (1) *asymmetric* (unequal emphasis or obstacles for specific choices), (2) *covert* (hidden interface design choices), *deceptive* (induce false beliefs), (3)

*hides information* (obscure or delay the communication of relevant information), and (5) *restrictive* (limitation of choices). The authors specifically name cookie consent dialogs which make use of a highlighted “accept” button as an example for the *asymmetric* dimension.

In the context of the GDPR, one could argue that tactics involving increased complexity, hidden information, or unwanted default settings—if effective—violate the requirements for clear and informed consent. Our study adds empirical evidence on the effectiveness of these tactics in the specific context. We vary the presence of a potentially misleading default button and measure perceived deception, unlike the wealth of studies that quantify this bias by merely observing the behavioral reaction to default buttons.<sup>2</sup> Since decisions in the privacy context often involve high cognitive load, we devise a combined (but not confounded) experiment with choice proliferation. This allows us to interpret perceived difficulty and response time—both proxies for cognitive load—in relation to perceived deception.

### 3.3 Hypotheses

Against the backdrop of the features in consent dialogs used by popular websites and the underlying theoretical considerations, we postulate four hypotheses:

- H1** If consent dialogs include a highlighted default button that selects all purposes, users effectively consent to more purposes than without this button.
- H2** If consent dialogs include a highlighted default button that selects all purposes, users
  - (a) regret their decision more and
  - (b) perceive the website as more deceptive
 than without this button, after being informed about the purposes they effectively consented to.
- H3** If consent dialogs present multiple purposes, users require more effort than for dialogs with a single purpose, as indicated by longer response times.
- H4** If consent dialogs include multiple purposes, users perceive the task as more difficult than reacting to dialogs with a single purpose.

In the hypotheses and the following, we shall use the term “effective consent” to refer to the consent statement recorded by the website, independent of whether this corresponds the user’s true intention.

<sup>2</sup> The default effect is in the order of 5 %-pts. for a consent dialog where about one of two participants agrees [31].

## 4 Method

To test the proposed hypotheses, we conducted a controlled experiment. We describe and justify the instrument in Section 4.1, then report from our pretests (Section 4.2) and the survey administration (Section 4.3). Ethical considerations are discussed in Section 4.4. Descriptive statistics are presented in Section 4.5.

### 4.1 Instrument

The survey instrument has two main components: a functional mock-up website offering flight search, and an exit questionnaire. As experimental factor, the mock-up randomly presents the user one of the three consent dialogs depicted in Figure 2. When categorizing these dialogs along the dimensions proposed by Schaub et al. [19], they constitute privacy notices which appear at setup (*timing*), in the primary *channel*, as visual pop-ups (*modality*) that include a blocking *control*. We copied the three purposes (statistics, comfort, personalization) from the airline website in verbatim in order to maximize external validity, noting that they differ from the convention discussed in Section 2.2. Users could learn more about the purposes by clicking on a small roll-down button labelled “show details” (see screenshot in Fig. 10 in the Appendix). Accordingly, *comfort* corresponds to *preferences*, and *personalization to marketing*, however without an indication whether this includes third-party tracking.

The treatment of the first group (T1) is a deceptive dialog, which closely resembled the one we saw on the German airline website (cf. Figure 9). It contains an explanation text about different cookie settings, three selectable purposes with initially unchecked checkboxes, an expandable part providing more details about the categories, and two buttons. The first button with the text “Select all and confirm” stands out due to its yellow color. The second button is colorless and says “Confirm selection” in gray font. If the yellow button is clicked, the user (effectively) consents to all three purposes, regardless of which boxes are checked. In contrast, a click on the second button only confirms the settings that have actively been selected by the user.

The second treatment (T2) differs in the reduction of selectable categories. Specifically, it only includes the *personalization* purpose. Pretests have shown that personalization is perceived as the most sensitive purpose, thus we deemed it plausible to make this purpose op-

tional. We could confirm this post-hoc: only 23% of the users in the control group consent to *personalization*, versus 35% for *comfort* and 46% for *statistics*. The results of Utz et al. [9] corroborate this further.<sup>3</sup> Arguable, the T2 dialog appears somewhat artificial, but it was the best way we could think of reducing the number of choices without changing the dialog to a yes/no question. We could not spot any indication that users perceived this dialog as odd in the responses to an open-ended question in the exit survey.

In contrast to the two treatments, the control group did not see a highlighted default button. The control dialog offers the same three purposes as observed in reality. We refrained from presenting a version with one purpose and no default button for the lack of hypotheses on potential interaction effects, and to increase the number of subjects in the interesting three groups. Therefore, our study technically combines two  $1 \times 2$  experiments with one overlapping group rather than realizing a complete  $2 \times 2$  design.

We decided against additional treatments with opt-out (i. e., where purposes are pre-selected) because they are almost certainly not compliant with the GDPR [15]. For the same reasons, we see little prospect for non-blocking cookie banners if the website to some extent depends on consent as the legal basis to process personal data. For comparison, Utz et al. [9] test two opt-out conditions in their field study of non-blocking banners.

The actual flight search website has a simplistic design and only contains text fields and date selectors for the search input. To increase realism, some “special offers” for specific destinations are depicted next to a photo of the respective city. These measures were intended to draw the focus away from the cookie dialog. The participants’ interaction on the website is captured and continuously transmitted to our server. This allows us to analyze response time, click trajectories, and possible dropouts post-hoc.

We measure the participants’ perceptions of the website in an exit questionnaire. At first, participants are asked to freely list positive and negative aspects of the website. Thereafter, they should recall their chosen cookie settings in the dialog; first in free-text form and followed by closed questions. Besides general questions on the cookie dialog, four established constructs are measured through multi-item scales. Such scales are common in psychometrics to attenuate the mea-

**Table 1.** Constructs and corresponding items.

Item	Item text (translated from German)
<i>Perceived Deception (PDE)</i>	
PDE1	When it comes to cookie settings, the website is dishonest towards its users.
PDE2	The website tries to mislead users towards selecting cookie settings which they do not intend to select.
PDE3	The website makes use of misleading tactics so that users select cookie settings which they do not intend to select.
<i>Perceived Difficulty (PDI)</i>	
PDI1	It was incomprehensible to select cookie settings.
PDI2	It was frustrating to select cookie settings.
PDI3↔	It was easy to select cookie settings.
<i>Regret (RE)</i>	
RE1	I regret my choice of cookie settings.
RE2	I would change my cookie settings if it was possible.
RE3↔	I am satisfied with my choice of cookie settings.
<i>Privacy Attitudes (PA)</i>	
PA1	It is important for me to protect my privacy online.
PA2	If websites use cookies, my online privacy is impaired.
PA3	I am concerned about my online privacy being impaired by website cookies.

Items marked with ‘↔’ use inverted scales.

surement error of individual items. All construct items are reported in Table 1. Answers were collected on 5-point rating scales with semantic anchors “strongly disagree” (1) and “strongly agree” (5). Perceived deception (PDE) is assessed using three (of originally four) items by Román [44], adapted to the context of our study.<sup>4</sup> Additionally, we measure perceived difficulty (PDI), privacy attitudes (PA), and regret (RE). RE is measured twice in the questionnaire: before and after reminding the participants of their effective cookies settings.

## 4.2 Pretests

Two pretests were conducted in order to assess the clarity of the instructions and survey questions. First, we carried out two one-on-one tests using *verbal probing* and *think aloud* techniques. Specifically, test subjects were asked to express their thought process and potential obstacles while going through the survey. Since we found that it is confusing to first open a link with a cookie dialog, and then receive the flight search task,

<sup>3</sup> See Fig. 5 (1a) of [9], although the precision is low and the baseline not comparable.

<sup>4</sup> The fourth item was dropped because it was too specific to the domain of online shopping.




Fig. 2. Variations of consent dialogs shown in the study. All dialogs are blocking. The participants saw German versions (see Fig. 8).

we decided to rearrange the instructions. This way, the participants are even more focused on flights than on cookies before visiting the website.

In order to simulate the actual survey environment in a lecture hall, the second pretest was conducted with 20 Austrian undergraduate students in a computer lab. This way, we were able to estimate the required time for each survey step. During the test, we observed that several test subjects glanced at their neighbors' screens and talked to one another while completing the survey. As this behavior might reduce the data quality, we added the appeal to work quietly and by oneself to the instructions. Additionally, we found and fixed a bug concerning the collection of timestamps.

### 4.3 Survey Administration

The data collection took place on two days in January 2019 at the University of Innsbruck in Austria and the University of Münster in Germany. All parts of the instrument and the written and spoken instructions were provided in the local language (German). We report the original wording of scale items and selected screenshots in the appendix (Section 8) to facilitate error analysis and possible replication studies. The survey was administered at the beginning of lectures attended by undergraduate computer science students, mainly in the first year. Figure 3 summarizes the steps of the data collection, along with the order of the measured constructs.

During the briefing, we informed the participants that their data will be held confidential and cannot be

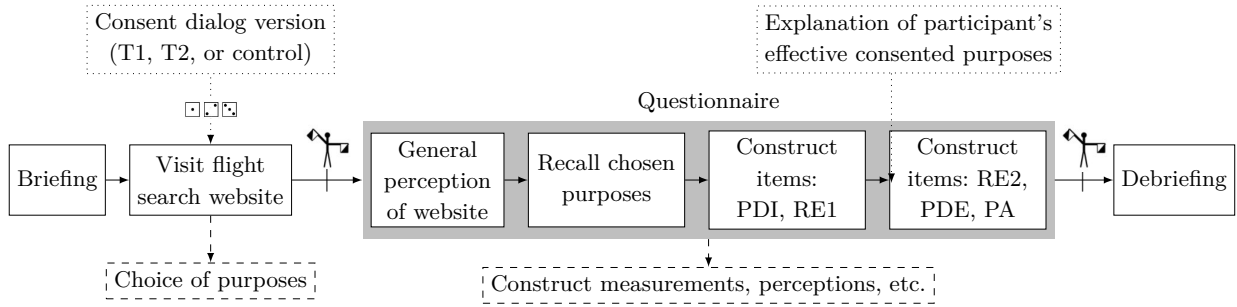
linked to their identity. We also pointed out the voluntary nature of participating in the study and asked them to conscientiously follow the instructions without interacting with one another. We communicated that the scope of the study is about the user experience of flight search websites, without mentioning the focus on cookie notices or privacy.

In the second step, participants were given the task to search for a flight with a specific departure, destination, and time. Then, we provided the link to the flight search website that is described above. When visiting the link, one of the three cookie dialogs depicted in Figure 2 was randomly assigned to each participant. After reacting to the dialog and entering the flight search, a modal window appeared, which asked to wait for further instructions.

When the vast majority had reached this step, a key combination for opening the questionnaire was displayed on the lecture hall's main projector. This way, participants started answering the questions almost simultaneously. It took the average participant 6' 34" to complete the questionnaire while 90% were done after 9' 06".

In the debriefing, we informed the participants about the topic of the study and showed them screenshots of cookie notices used by real websites. We ran the classroom experiment exactly once at each university, one in Austria and one in Germany, thereby minimizing the likelihood that earlier participants could tell later participants the true purpose of the study.





**Fig. 3.** Visualization of the study process with treatments (dotted boxes) and measurements (dashed boxes). Semaphores denote synchronization points: all participants in the classroom proceed to the next step simultaneously.

## 4.4 Research Ethics

In fulfillment of approved ethical standards, we clearly communicated that participation is voluntary and anonymous. Respondents could skip questions they did not want to answer. The search task itself and the surrounding stimulus material was chosen to not raise emotions or strong feelings. Independent of the participants' selected cookie settings, we did not store cookies and only transmitted data to our servers that are relevant for the research purpose.

In not revealing the true purpose of our study right away, we applied deception ourselves as part of the research method. This is common practice and was in accordance with the ethical oversight bodies at all universities involved. The practice is deemed acceptable in particular because of the low probability of causing harm and the fact that we revealed the purpose of our study in the debriefing, where we also provided contact information and offered the communication of results.

The experiment caused an opportunity cost of 10 minutes lost lecture time for everyone in the room, including about 8 students per session (15 altogether) who did not participate in the experiment. To minimize the harm, we chose the beginning of a Q&A session that had not used the entire allocated time in the previous years. Moreover, the main reason why students did not participate was that they arrived late in class.

## 4.5 Descriptive Statistics

Table 2 reports descriptive statistics of our sample. In total, 164 students took part in the study whereof 158 completed the survey. We deleted 8 records due to more than four missing responses on critical construct items. The remaining 13 records with missing answers contained a total of 20 missing values, which were replaced

by the mean of the observed item score. Consequently, the analysis uses 150 valid cases. As a consequence of the convenience sample, the ratio of female participants was below 20%, which is typical for German-speaking computer science undergraduates.

Even though we asked participants to use their laptops for completing the survey, we allowed those without one to chose another device they had at hand. While 42.7% followed the survey instructions on a screen

**Table 2.** Descriptive statistics.

Item	Number	Fraction
All	150	100.0%
<i>Group</i>		
T1	50	33.3%
T2	48	32.0%
Control	52	34.7%
<i>Location of the university</i>		
Austria	90	60.0%
Germany	60	40.0%
<i>Screen width</i>		
<500px	64	42.7%
500px–1000px	7	4.7%
>1000px	79	52.7%
<i>Browser</i>		
Chrome	81	54.0%
Firefox	30	20.0%
Safari	29	19.3%
Other	10	6.7%
<i>Knowledge about cookies</i>		
Self-reported knowledge	121	80.7%
Correctly described cookies	102	68.0%
<i>Privacy measures (self-report)</i>		
Regularly deletes cookies	64	42.7%
Has cookies disabled	36	24.0%
Uses ad-blocker	113	75.3%
Uses anti-virus software	80	53.3%

width below 500 pixels (i. e., likely smartphones), 52.7% had a screen width above 1000 pixels (i. e., likely laptops). Most participants opened the website on Chrome (54.0%); others used Safari (19.3%) or Firefox (20.0%). We did not observe noteworthy differences in results between device types, browsers, or locations and thus refrain from reporting breakdowns in the following.

By asking whether participants know what browser cookies are, we find that 68.0% are able to provide a correct explanation. Only 12.7% claim to know what cookies are, but provided either no explanation or an incorrect one. The remaining 19.3% stated that they have no knowledge about cookies.

On average, it took participants 11.8 seconds to respond to the cookie dialog (median: 7.3"). Only 8.7% expanded the dialog by clicking on "Show details", and 3.3% revised their initial choice by unselecting at least one purpose. In total, 41.3% did not consent to any cookie purpose, while 30.7% accepted all purposes offered. Of all participants in the two treatment groups ( $n = 98$ ), 56.1% clicked on the default button, which results in accepting all purposes regardless of which purposes were actively selected. Of these participants, 34.5% ( $n = 19$ ) still selected at least one purpose.

To evaluate the construct reliability of PDE, PDI, RE und PA, we examine internal consistency by calculating Cronbach's  $\alpha$ . As shown in Table 3, each constructs' Cronbach's  $\alpha$  value lies above 0.7, indicating that they are sufficiently consistent [59] and thus suitable for further analysis. We also check if the construct scores are sufficiently close to a normal distribution to justify the use of parametric inference statistics. Table 4 shows Q-Q plots for all constructs and reports the results of Kolmogorov–Smirnov (KS) tests for normality. Given that the deviations from normality are visibly caused by the range limits only, no KS-test rejects the null hypothesis at the 1% level, and the way we compute the scores cannot produce any outliers, we deem it safe to report hypothesis tests with parametric  $t$ -tests. To err on the side of caution, we report  $p$ -values for the two-sided test although all our hypotheses are directed.

## 5 Results

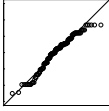
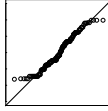
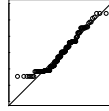
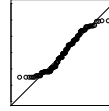
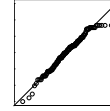
We begin with the deductive hypothesis tests, before we investigate additional aspects in a quantitative explorative way (Section 5.2).

**Table 3.** Construct reliability.

Construct	Cronbach's $\alpha$	Mean	Median	SD
PDE	0.79	3.52	3.67	1.13
PDI	0.73	2.82	2.67	1.17
RE-before	0.83	2.39	2.33	1.14
RE-after	0.74	2.62	2.33	1.20
PA	0.80	3.61	3.67	0.91

Each construct has 3 items.

**Table 4.** Normality of construct distributions.

PDE	PDI	RE-before	RE-after	PA
Normal Q-Q plots for the range $\pm 3$ SD				
				
One-sample Kolmogorov–Smirnov tests for normality				
$p = 0.10$	$p = 0.02$	$p = 0.04$	$p = 0.03$	$p = 0.47$

### 5.1 Hypothesis Tests

We test **H1** by analyzing whether the deception group and the control group differ in the number of purposes they effectively agreed to. To do so, we assign a score from 0 (no purposes, by clicking on "Confirm selection" without checking any box) to 3 (all purposes by either checking all boxes and then clicking any button, or by clicking the highlighted default button). Participants who saw the deceptive dialog effectively consented to more purposes. Table 5 presents the score values by treatment and control group. Since the score is a count variable, we use the non-parametric Kruskal–Wallis (KW) test, which indicates as strongly significant effect ( $\chi^2(1) = 7.2$ ,  $p < 0.01$ ) in the hypothesized direction. **This supports H1.** The effect of the deceptive default button on agreeing to no or all purposes is in the order of 20 %-pts., about four times larger than the plain default effect reported for an application consent dialog in [31]. We additionally check if participants in the deception group are more likely to consent to all three, instead of two or less purposes, than the control group. A Chi-squared test ( $\chi^2(1) = 9.05$ ,  $p < 0.005$ ) reveals a highly significant difference.

To test **H2a**, we compare the measurements of regret before (RE-before) and after (RE-after) the participants got informed about the purposes they effectively

consented to. Results of the paired *t*-test reveal a significant difference between the before/after states for the deception group ( $t(49) = 2.81, p < 0.01, d = 0.40$ ). **This supports H2a.** The difference in the control group is not significant ( $t(51) = 1.63, p = 0.11, d = 0.23$ ). Thus, we can attribute the regret to the misinformation caused by the deceptive design.

When analyzing perceived deception (PDE) for testing **H2b**, a notable difference between groups can be found (Figure 4). The *t*-test shows that PDE of the deception group is significantly higher than of the control group ( $t(96.8279) = 2.24, p < 0.05, d = 0.44$ ). Therefore, **H2b is also supported.**

Drilling down into the findings on H2a and H2b, we analyze if participants within the deception group who clicked on the deceptive default button perceive even more regret and deception after being informed about the consequence of their response. Indeed, our measurements of RE-after are significantly higher for those who clicked the default button compared to all other participants in T1 ( $t(43.962) = -3.82, p < 0.0005$ ). However, no significant differences for perceived deception can be found ( $t(47.73422) = 0.64, p = 0.64$ ). These two results can be explained by the presence of smart participants who debunk the default button as deceptive and do not fall for it. They have less to regret than those who only understand the button’s effect after the fact.

To test **H3**, we investigate the time needed to complete the consent dialogs. The measurement starts when the cookie dialog appears and ends when the participant clicks a button. This measure reflects the effort required for responding to the dialog. As shown in Figure 5, participants in the group with reduced choice spent on average five seconds less on their response than those who were presented with three purposes. The difference in medians shows the same trend, albeit less pronounced due to the skewed distribution. We choose non-parametric statistics to account for this fact. When only comparing the deception and reduced choice group, the KW-test reveals that the difference ( $\chi^2(1) = 8.89, p < 0.005$ ) is highly significant, which **supports H3.** We also find a significant difference between the reduced choice group and the control group ( $\chi^2(1) = 9.73, p < 0.005$ ). The difference between the deception and control group, which offer the same number of purposes, is not significant ( $\chi^2(1) = 0.17, p = 0.68$ ). The results indicate that the number of purposes is positively associated with cognitive load, even if the number of options is way below Miller’s “magic seven” [60].

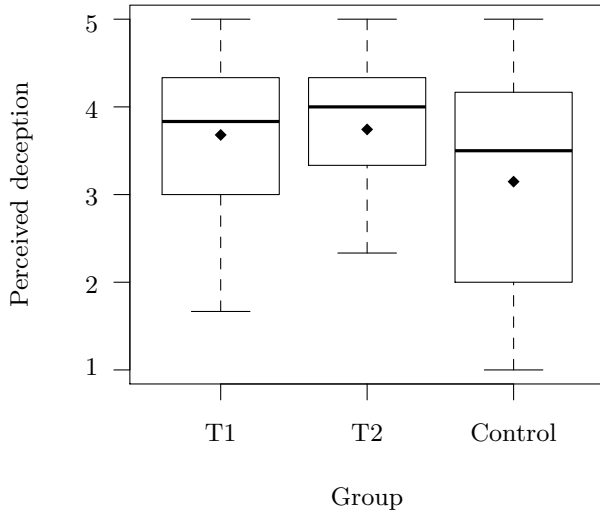
Regarding **H4**, we interpret perceived difficulty (PDI) as a measure of dissatisfaction. Specifically, we

**Table 5.** Overview of results by treatment group.


Dependent variable	Group			Test
	T1 (n = 50)	T2 (n = 48)	Control (n = 52)	
<i>Number of effectively consented purposes</i>				KW-test
0	<b>32.0%</b>	41.7%	<b>50.0%</b>	
1	<b>12.0%</b>	58.3%	<b>19.2%</b>	
2	<b>2.0%</b>	–	<b>7.7%</b>	
3	<b>54.0%</b>	–	<b>23.1%</b>	** H1
	100.0%	100.0%	100.0%	
<i>Construct means</i>				<i>t</i> -test
PDE	<b>3.68</b>	3.74	<b>3.15</b>	* H2b
PDI	<b>2.87</b>	<b>2.90</b>	2.69	n.s. H4
RE-before	<b>2.30</b>	2.64	<b>2.23</b>	n.s.
RE-after	<b>2.69</b>	2.79	<b>2.39</b>	n.s.
PA	3.50	3.70	3.62	
– “ – (only subjects who clicked the default button) (n = 27) (n = 28)				
PDI	2.78	2.72		
RE-after	3.20	3.48	–	
PA	3.32	3.56	–	
<i>RE-after minus RE-before (all subjects)</i>				paired <i>t</i> -test
	<b>0.39</b>	0.15	0.16	** H2a
<i>Response time for consent dialog (seconds)</i>				KW-test
Median	<b>5.36</b>	<b>3.16</b>	7.24	** H3
Mean	10.95	5.41	11.62	
<i>Correct recall of effective purposes</i>				$\chi^2$ -test
	<b>73.5%</b>	<b>60.9%</b>	<b>90.0%</b>	**
– “ – (only subjects who clicked the default button) (n = 27) (n = 28)				
	55.6%	53.6%	–	

Legend: \*  $p < 0.05$ , \*\*  $p < 0.01$ , n.s. not significant.  
The test results refer to the bold values in the same row.

test the difference between the two treatment groups in order to show whether the number of purposes in the consent dialog affect the participants’ perceptions. Since the *t*-test results in no significant difference ( $t(95.99) = 0.16, p = 0.88, d = 0.07$ ), **H4 must be rejected.** Unrelated to our hypotheses, we also tested for differences in PDI between the control group and T1, respectively, T2. No test result was even close to statistical significance.



**Fig. 4.** Perceived deception (PDE) by treatment. The  $t$ -test shows that PDE is significantly higher in T1 than in the control group ( $t(96.82) = 2.24, p < 0.05$ , Cohen’s  $d = 0.44$ ).



**Fig. 5.** Time spent on responding to the consent dialog by treatment. The KW-test shows a significant difference between T1 and T2 ( $\chi^2(1) = 8.89, p < 0.005$ ).

## 5.2 Post-hoc Analyses

At the beginning of the questionnaire, we asked participants to recall which purposes they have agreed to in the consent dialog. Thus, we are able to compare the accuracy of participants’ statements between groups. As reported in Table 5, the difference is highly significant between all three groups ( $\chi^2(2) = 11.01, p < 0.005$ ).

When looking at the proportion of participants who declined all purposes, it is notable that 50% of the con-

trol group, but only 32% of the deception group chose this option. After informing participants about their choice, we specifically asked those who agreed to at least one purpose whether they had been aware of the possibility to decline all purposes. Only 32% stated to be aware of this option. However, the proportion of aware participants does not differ significantly between the deception and control group ( $\chi^2(1) = 2.71, p = 0.10$ ). It seems that even the design of our control dialog, possibly in combination with learned expectations, imposes some pressure to select at least one option on a subset of the participants. This highlights that future research could seek to improve the communication of the “freely given” aspect of GDPR-compliant consent (cf. Section 2.1).

To test whether users’ privacy attitudes regarding cookies influence their reaction to the cookie dialog, we also test the relationship between the number of chosen purposes and PA. For this analysis we only consider the groups that were presented all three purposes. We find a weak but significant *negative* correlation between privacy attitudes and the number of consented purposes ( $r_s = -0.23, p < 0.05, n = 102$ ). However, the difference in PA between those who clicked the deceptive button and those who did not, is not significant ( $t(91.9) = -1.04, p = 0.30, n = 98$ ). Moreover, as expected, privacy attitudes do not differ significantly between groups as participants were randomly assigned to groups. This reassures us that the PA items measure trait rather than state.

## 6 Discussion

Next we reflect on the results, then discuss limitations (Section 6.2), and comment on recent developments in the space (Section 6.3).

### 6.1 Summary and Interpretation

Our experimental results confirm the common conjecture that design elements of consent dialogs can nudge users towards making specific choices. We show empirically that the selection of data processing purposes, as required by the GDPR, is not exempt: users accept more data collection purposes when consent dialogs integrate a highlighted default button that selects all purposes at once. Surprisingly, we observe a four times stronger effect for our multi-purpose consent dialog than previously reported for simple default buttons in binary con-

sent dialogs. Moreover, the fact that users who click this button are less likely to correctly recall the consented purposes casts doubt on the morality and legitimacy of this design element, as it might lead users to act against their intention. This interpretation is further supported by the finding that users tend to regret their decision after being informed about the effective purposes.

Besides the effect of deceptive default buttons, we present more encouraging results on the possibility of differentiating between consent decision for multiple purposes in one dialog: although the number of purposes significantly affects the response time, the difference in perceived difficulty is insignificant. This indicates that most users can handle three different purposes without experiencing the negative effects predicted by the theory of choice proliferation. Of course, more research is needed to investigate the critical number of purposes. Also choice structure, the other relevant determinant in choice proliferation, requires further attention [13].

Our analysis of control variables reveals that users with stronger stated privacy attitudes consent to fewer purposes. While this result seems to challenge the privacy paradox (a term for the often observed discrepancy between stated attitudes and privacy behavior [61–63]), it must be interpreted with caution. First, our instrument is not ideal to study the paradox. It confounds this relationship with the dominant effect of a deceptive default button and measures the privacy attitude only after recalling the effective purposes. Second, unlike in many studies that find the paradox, our items measure privacy attitudes quite narrowly for the specific domain: two out of three items mention cookies. According to the principle of compatibility, behavior is more predictable from attitudes if it is measured on the same level of specificity [64]. Third, the interpretation of attitude–behavior links is problematic if the behavior is partly unintentional, such as accepting undesired purposes. To some extent, this corroborates nuanced or critical perspectives on the privacy paradox [65].

## 6.2 Validity and Limitations

To gauge the relevance of our results, one may ask how prevalent the tested dialog is on the web. Unfortunately, reliable data in this dynamic space is scarce. The most recent data in [9] refer to a snapshot in August 2018 and thus predate the introduction of the dialog on the airline website, where we discovered it, and possibly elsewhere. According to this snapshot, only 7% of web consent notices are blocking, and 8% present multiple purposes

**Table 6.** Robustness of the main effects: *p*-values of hypothesis tests broken down by the location of the classroom experiment.

Hypothesis and contrast groups	Austria ( <i>n</i> = 90)	Germany ( <i>n</i> = 60)
H1: T1 vs Control	0.028	0.090
H2a: RE-before/-after in T1	0.036	0.098
H2b: T1 vs Control	0.019	0.564
H3: T1 vs T2	0.031	0.041
H4: T1 vs T2 (rejected)	0.572	0.560

(immediately or on request). These shares almost certainly increased with the adoption of consent managers in the course of 2019 (see Sect. 6.3 below).

However, our choice of stimulus was not driven by the most prevalent design, which we and other researchers [5, 9, 10] suspect to trivially violate the GDPR. Instead, we set out to study elementary design options of multi-purpose dialogs, the novel and most under-researched aspect of consent dialogs. Our dialog implements opt-in and does not proceed without an affirmative action (blocking) in order to anticipate future good practices. The fact that similar dialogs are used by respectable organizations with competent legal departments and millions of unique users per year<sup>5</sup> adds to the relevance. More importantly, since we study individual effects derived from theory, the prevalence of our stimulus material is of subordinate importance. We aim to identify generalizable effects, which could be studied on real or artificial dialogs. The choice of using a real dialog for inspiration along with a credible cover story is merely one of multiple measures to assure the external validity of our lab study.

To check for possible risks to external validity, we analyzed the participants’ free-text responses for prejudiced assumptions about our study. Only one participant exhibited demand characteristics [66]. The person wrote that he or she has agreed to all purposes because the website was part of a scientific experiment. All remaining participants answered as if they were dealing with an actual flight search website. Moreover, we do not find further indications that the participants might have perceived our stimuli as artificial.

It is important to mention that the study has limitations. First, the experimental setup may not fully reflect users’ actual behavior regarding consent dialogs. Even though we made an effort to hide the research

<sup>5</sup> Figures extracted from media data of the airline’s online services, available from the authors on request.

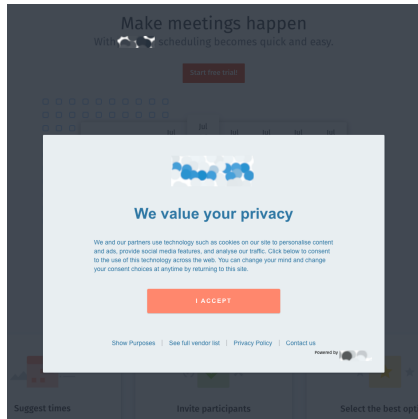


Fig. 6. Dialog of a commercial consent manager (August 2019).

purpose of our study, we cannot rule out that participants might have guessed our focus on cookie choices or privacy in general. Moreover, our sample is limited to German-speaking computer science students who are probably more educated about the functionality of cookies and the web in general. This *known* bias, however, does not compromise the upshot of this paper: if even computer literate populations fall for the deceptive design, we must assume that the outcome for the general public is even worse. We tried to mitigate all other disadvantages of convenience samples by replicating the experiment in two geographically distant universities. Table 6 confirms that H1–H3 are supported in both populations.<sup>6</sup> We chose a classroom experiment (and accepted its limitations) in order to reduce *unknown* biases due to participant self-selection, which is an acute problem of empirical privacy research [67, 68]. Precisely the attitudes and beliefs of interest correlate with non-response and dropouts. For perspective, our completion rate is above 88% in all sessions, whereas the concurrent field experiment received 110 completed surveys from more than 30,000 solicitations, translating to a response rate below 0.4% [9]. (The authors acknowledge this bias and chose not to analyze self-reported data quantitatively.)

<sup>6</sup> Some  $p$ -values for Germany are above 5% only because we conservatively use the two-sided test. The values for the one-sided test are half of the reported ones. Recall that group sizes in Germany alone may be below 20 subjects.

## 6.3 Recent Developments

A simple interface adjustment, meanwhile implemented in the airline website, is to change the button text from “Select all and confirm” to “Select and confirm” (and change the function accordingly) as soon as the first checkbox is selected. While this breaks with the design principle that button semantics should be stateless, it might avoid the severest mishaps where cognitive effort that went into selecting purposes is wasted. It requires another user study with more participants to gauge if this modification reduces the disappointment in the subset of users who select at least one but not all purposes.<sup>7</sup>

In the past months, we (and others [10]) have observed other “innovative” consent dialogs, such as page-long lists of affiliate partners for third-party tracking, that call for tailored user studies through the lenses of deception and choice proliferation. For example, a popular meeting scheduling service uses a modal dialog entitled “We value your privacy” with a prominent button labeled “I accept.” To access literally hundreds of options, one has to click on “Show purposes”, which is a text link next to three others (see Fig. 6). Interestingly, this dialog seems to be operated (and presumably evaluated) by an intermediary specialized on consent management. This fits into the picture where ENISA, a EU agency, mentions “consent management” as a new business opportunity for cybersecurity startups [69, p. 10].

## 7 Conclusion

This study presents new empirical evidence supporting that design elements used in consent dialogs of popular websites might deceive users into agreeing to more data processing purposes than intended. It complements the measurement studies [5, 6, 8, 10] that emphasize the wide adoption of such “dark patterns” [56] as well as a recent field study on cookie banners [9]. Based on these findings, we can derive recommendations for user interface designers and policy makers.

Our first recommendation reiterates calls respect the user’s interest: instead of nudging users towards agreements that mainly benefit the party who owns the website, defaults should reflect either a privacy-aware

<sup>7</sup> This description applied to the time of writing in mid-2019. The checkbox logic had been changed once again when we revisited the website in fall 2019 for the preparation of the camera-ready version. This highlights the dynamics in this space.

safe choice or elicit the majority's preferences as a descriptive social norm. This could be achieved by designing a set of best-practice consent dialogs, incorporating the body of knowledge from behavioral privacy research. These templates can be made available to organizations who value consumers' privacy or seek legal certainty without commissioning an intermediary.

However, past and ongoing efforts in the usable privacy research community towards understanding how to nudge users into making safer choices are void if the industry tries to achieve the opposite. Since the value of personal data increases with the number of possible secondary uses [33], businesses have incentives to maximize the number of consented purposes. It is tempting to call for a regulator or oversight body to step in and ensure that dialogs are designed in the users' interest. But we are hesitant about suggesting more (or more specific) regulations for two reasons. First, the GDPR stipulates freely given, unambiguous, and informed consent. It may take a court decision to provide clarity over the fact that the practices we observe *do not* meet these requirements and hence *cannot* provide a legal basis for personal data processing. However, such decisions must be based on further empirical research. Second, the time and cognitive effort millions of users regularly spend on consent dialogs may not justify the outcome at the societal level. Rather than mandating special forms of consent dialogs (which hardly work for devices without display or network services that are not customer-facing), a policy priority should be the establishment of a standard for non-interactive privacy preference negotiations.

It seems that P3P [25] was 20 years ahead of its time, and the do-not-track header too simple and polarized [70]. There could be a middle ground in which consent dialogs do not disappear. But their design moves from the hands of data controllers to developers of user agents, who compete for the best service in the data subject's interest. In order to foster competition, and not to repeat the mistakes of do-not-track, it is important that browser and app vendors must be required to interoperate with any privacy agent of the user's choice.

## Acknowledgements

We thank all participants of the pretests and the main study. We also thank Daniel Woods, Henry Hosseini and our shepherd Blase Ur for useful discussions as well as the anonymous reviewers for many constructive comments. This work received funding from the German

Bundesministerium für Bildung und Forschung (BMBF) under grant agreement 16KIS0382 (AppPETs) and the Archimedes Privatstiftung, Innsbruck.

## References

- [1] European Parliament and the Council of the European Union. Regulation (EU) 2016/679 of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation) (2016)
- [2] European Parliament and the Council of the European Union. Directive 2002/58/EC of 12 July 2002 concerning the processing of personal data and the protection of privacy in the electronic communications sector (Directive on privacy and electronic communications) (2002)
- [3] European Parliament and the Council of the European Union. Directive 2009/136/EC of 25 November 2009 amending Directive 2002/22/EC universal service and users' rights relating to electronic communications networks and services, Directive 2002/58/EC concerning the processing of personal data and the protection of privacy in the electronic communications sector and Regulation (EC) No 2006/2004 on cooperation between national authorities responsible for the enforcement of consumer protection laws. (2009)
- [4] R. Leenes, E. Kosta. Taming the cookie monster with Dutch law — A tale of regulatory failure. *Computer Law & Security Review* (2015) 31, 3, 317–335
- [5] M. Trevisan, B. E. Traverso, Stefano, M. Mellia. 4 years of EU cookie law: Results and lessons learned. In: *Proceedings on Privacy Enhancing Technologies (PoPETs)* (De Gruyter Open, 2019) 126–145
- [6] R. van Eijk, H. Asghari, P. Winter, A. Narayanan. The impact of user location on cookie notices (inside and outside of the European Union). In: *Workshop on Technology and Consumer Protection (ConPro)* (2019)
- [7] S. Englehardt, A. Narayanan. Online tracking: A 1-million-site measurement and analysis. In: *Conference on Computer and Communications Security (CCS)* (ACM, 2016) 1388–1401
- [8] M. Degeling, C. Utz, C. Lentzsch, H. Hosseini, F. Schaub, T. Holz. Measuring the GDPR's impact on web privacy. In: *Network and Distributed System Security Symposium (NDSS)* (Internet Society, 2019)
- [9] C. Utz, M. Degeling, S. Fahl, F. Schaub, T. Holz. (Un)informed consent: Studying GDPR consent notices in the field. In: *Conference on Computer and Communications Security (CCS)* (ACM, 2019) 973–990
- [10] I. Sánchez-Rola, M. Dell'Amico, P. Kotzias, D. Balzarotti, L. Bilge, P. Vervier, et al. Can I opt out yet?: GDPR and the global illusion of cookie control. In: *Conference on Computer and Communications Security (AsiaCCS)* (ACM, 2019) 340–351
- [11] J. R. Kling, S. Mullainathan, E. Shafir, L. Vermeulen, M. V. Wrobel. Misperception in choosing medicare drug plans. Harvard University working paper (2008)

- [12] H. Cronqvist, R. H. Thaler. Design choices in privatized social-security systems: Learning from the Swedish experience. *American Economic Review* (2004) 94, 2, 424–428
- [13] S. Korff, R. Böhme. Too much choice: End-user privacy decisions in the context of choice proliferation. In: *Symposium On Usable Privacy and Security (SOUPS)* (USENIX, 2014) 69–87
- [14] European Commission. The GDPR: New opportunities, new obligations. Tech. rep., Publications Office of the European Union, Brussels, Luxembourg (2018)
- [15] European Court of Justice. Judgement in Case C-673/17 in the proceedings *Bundesverband der Verbraucherzentralen und Verbraucherverbände vs Planet49 GmbH* (2019)
- [16] Cybot A/S. WordPress and GDPR, and how to deal with cookies and plugins. Copenhagen, Denmark (2019). <https://www.cookiebot.com/en/wordpress-cookie-plugin/>
- [17] M. Ackerman, L. F. Cranor, J. Reagle. Privacy in e-commerce: Examining user scenarios and privacy preferences. In: *Conference on Electronic Commerce (EC)* (ACM, 1999) 1–8
- [18] L. I. Millett, B. Friedman, E. Felten. Cookies and web browser design: Toward realizing informed consent online. In: *Conference on Human Factors in Computing System (CHI)* (ACM, 2001) 46–52
- [19] F. Schaub, R. Balebako, A. L. Durity, L. F. Cranor. A design space for effective privacy notices. In: *Symposium On Usable Privacy and Security (SOUPS)* (USENIX, 2015) 1–17
- [20] M. Bergmann. Generic predefined privacy preferences for online applications. In: *IFIP International Summer School on the Future of Identity in the Information Society* (Springer, 2007) 259–273
- [21] J. S. Pettersson, S. Fischer-Hubner, M. C. Mont, S. Pearson. How ordinary internet users can have a chance to influence privacy policies. In: *Nordic conference on Human-computer interaction: Changing roles (NordiCHI)* (ACM, 2006) 473–476
- [22] J. S. Pettersson, S. Fischer-Hübner, N. Danielsson, J. Nilsson, M. Bergmann, S. Clauss, et al. Making PRIME usable. In: *Symposium on Usable Privacy and Security (SOUPS)* (ACM, 2005) 53–64
- [23] J. J. Borking. Privacy incorporated software agent (PISA): Proposal for building a privacy guardian for the electronic age. In: *Designing Privacy Enhancing Technologies: International Workshop on Design Issues in Anonymity and Unobservability* (Springer, 2001) 130–140
- [24] M. Bergmann. Testing privacy awareness. In: *IFIP Summer School on the Future of Identity in the Information Society* (Springer, 2008) 237–253
- [25] L. F. Cranor. P3P: Making privacy policies more useful. *Security & Privacy* (2003) 99, 6, 50–55
- [26] M. Langheinrich, L. Cranor, M. Marchiori. Appel: A P3P preference exchange language. *W3C Working Draft* (2002)
- [27] M.-R. Ulbricht, F. Pallas. YaPPL – a lightweight privacy preference language for legally sufficient and automated consent provision in IoT scenarios. In: J. García-Alfaro, J. Herrera-Joancomartí, G. Livraga, R. Rios (Eds.) *Data Privacy Management, Cryptocurrencies and Blockchain Technology* (Springer, 2018), no. 11025 in *Lecture Notes in Computer Science*, 329–344
- [28] T. Vila, R. Greenstadt, D. Molnar. Why we can't be bothered to read privacy policies. In: *Economics of Information Security* (Springer, 2004), 143–153
- [29] J. Grossklags, N. Good. Empirical studies on software notices to inform policy makers and usability designers. In: *Financial Cryptography and Data Security (FC)* (Springer, 2007) 341–355
- [30] O. Kulyk, A. Hilt, N. Gerber, M. Volkamer. Users' perceptions and reactions to the cookie disclaimer. In: *European Workshop on Usable Security (EuroUSEC)* (2018)
- [31] R. Böhme, S. Köpsell. Trained to accept?: A field experiment on consent dialogs. In: *Conference on Human Factors in Computing System (CHI)* (ACM, 2010) 2403–2406
- [32] A. P. Felt, S. Egelman, M. Finifter, D. Akhawe, D. A. Wagner, et al. How to ask for permission. *HotSec* (2012)
- [33] S. Spiekermann, A. Acquisti, R. Böhme, K. L. Hui. The challenges of personal data markets and privacy. *Electronic Markets* (2015) 25, 2, 161–167
- [34] B. Scheibehenne, R. Greifeneder, P. M. Todd. Can there ever be too many options? A meta-analytic review of choice overload. *Journal of Consumer Research* (2010) 37, 3, 409–425
- [35] E. J. Johnson, S. B. Shu, B. G. Dellaert, C. Fox, D. G. Goldstein, G. Häubl, et al. Beyond nudges: Tools of a choice architecture. *Marketing Letters* (2012) 23, 2, 487–504
- [36] B. P. Knijnenburg, A. Kobsa, H. Jin. Preference-based location sharing: Are more privacy options really better? In: *Conference on Human Factors in Computing System (CHI)* (ACM, 2013) 2667–2676
- [37] K. Tang, J. Hong, D. Siewiorek. The implications of offering more disclosure choices for social location sharing. In: *Conference on Human Factors in Computing System (CHI)* (ACM, 2012) 391–394
- [38] H. Krasnova, N. Eling, O. Schneider, H. Wenninger, T. Widjaja, P. Buxmann, et al. Does this app ask for too much data? The role of privacy perceptions in user behavior towards Facebook applications and permission dialogs. In: *European Conference on Information Systems (ECIS)* (2013)
- [39] C. Anderson. *The long tail: Why the future of business is selling less of more* (Hachette Books, London, UK, 2006)
- [40] J. M. Hutchinson. Is more choice always desirable? Evidence and arguments from leks, food selection, and environmental enrichment. *Biological Reviews* (2005) 80, 1, 73–92
- [41] R. Böhme, J. Grossklags. The security cost of cheap user interaction. In: *New Security Paradigms Workshop (NSPW)* (ACM, 2011) 67–82
- [42] S. Egelman. My profile is my password, verify me!: The privacy/convenience tradeoff of Facebook Connect. In: *Conference on Human Factors in Computing System (CHI)* (ACM, 2013) 2369–2378
- [43] P. E. Johnson, S. Grazioli, K. Jamal, R. G. Berryman. Detecting deception: Adversarial problem solving in a low base-rate world. *Cognitive Science* (2001) 25, 3, 355–392
- [44] S. Román. Relational consequences of perceived deception in online shopping: The moderating roles of type of product, consumer's attitude toward the internet and consumer's demographics. *Journal of Business Ethics* (2010) 95, 3, 373–391
- [45] B. Xiao, I. Benbasat. Product-related deception in e-commerce: A theoretical perspective. *MIS Quarterly* (2011)



- 35, 1, 169–196
- [46] A. Nochenson, J. Grossklags. An online experiment on consumers' susceptibility to fall for post-transaction marketing scams. In: European Conference on Information Systems (ECIS) (Association for Information Systems, 2014)
- [47] D. M. Boush, M. Friestad, P. Wright. Deception in the marketplace: The psychology of deceptive persuasion and consumer self-protection (Routledge/Taylor & Francis Group, 2015)
- [48] P. Fleming, S. C. Zyglidopoulos. The escalation of deception in organizations. *Journal of Business Ethics* (2008) 81, 4, 837–850
- [49] K. A. Jehn, E. D. Scott. Perceptions of deception: Making sense of responses to employee deceit. *Journal of Business Ethics* (2008) 80, 2, 327–347
- [50] K. Yoon, K. Knight, D. Martin. Deceiving team members about competence: Its motives and consequences. *Western Journal of Communication* (2018) 1–22
- [51] B. Shneiderman, M. Leavitt, et al. Research-based web design & usability guidelines (Department of Health and Human Services, Washington, DC, 2006)
- [52] J. Watson, H. R. Lipford, A. Besmer. Mapping user preference to privacy default settings. *Transactions on Computer-Human Interaction (TOCHI)* (ACM, 2015) 22, 32
- [53] C. I. Hovland, I. L. Janis, H. H. Kelley. *Communication and Persuasion: Psychological Studies of Opinion Change* (Greenwood Press, 1953)
- [54] C. M. Gray, Y. Kou, B. Battles, J. Hoggatt, A. L. Toombs. The dark (patterns) side of UX design. In: Conference on Human Factors in Computing System (CHI) (ACM, 2018) 534:1–14
- [55] H. Brignull. Dark patterns. Tech. rep. (2019). <https://darkpatterns.org>
- [56] C. Bösch, B. Erb, F. Kargl, H. Kopp, S. Pfattheicher. Tales from the dark side: Privacy dark strategies and privacy dark patterns (De Gruyter Open, 2016), vol. 2016 237–254
- [57] A. M. McDonald, L. F. Cranor. The cost of reading privacy policies. *I/S: J. L.* (2008) 4, 3, 540–565
- [58] A. Mathur, G. Acar, M. J. Friedman, E. Lucherini, J. Mayer, M. Chetty, et al. Dark patterns at scale: Findings from a crawl of 11K shopping websites. *Proceedings of the ACM on Human-Computer Interaction* (2019) 3, 81
- [59] J. M. Bland, D. G. Altman. Cronbach's alpha. *British Medical Journal* (1997) 314, 7080, 570–572
- [60] G. A. Miller. The magic number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological Review* (1956) 63, 2, 81–97
- [61] P. A. Norberg, D. R. Horne, D. A. Horne. The privacy paradox: Personal information disclosure intentions versus behaviors. *Journal of Consumer Affairs* (2007) 41, 1, 100–126
- [62] S. S. Sundar, H. Kang, M. Wu, E. Go, B. Zhang. Unlocking the privacy paradox: Do cognitive heuristics hold the key? In: *Extended Abstracts on Human Factors in Computing Systems (CHI EA)* (ACM, 2013), 6 811–816
- [63] N. Gerber, P. Gerber, M. Volkamer. Explaining the privacy paradox: A systematic review of literature investigating privacy attitude and behavior. *Computers & Security* (2018) 77, 226–261
- [64] I. Ajzen. Models of human social behavior and their application to health psychology. *Psychology and Health* (1998) 13, 4, 735–739
- [65] T. Dienlin, S. Trepte. Is the privacy paradox a relic of the past? An in-depth analysis of privacy attitudes and privacy behaviors. *European Journal of Social Psychology* (2015) 45, 3, 285–297
- [66] M. T. Orne. On the social psychology of the psychological experiment: With particular reference to demand characteristics and their implications. *American psychologist* (1962) 17, 11, 776–783
- [67] J. P. Walsh, S. Kiesler, L. S. Sproull, B. W. Hesse. Self-selected and randomly selected respondents in a computer network survey. *Public Opinion Quarterly* (1992) 56, 2, 241–244
- [68] H. Cho, R. LaRose. Privacy issues in internet surveys. *Social Science Computer Review* (1999) 17, 4, 421–434
- [69] European Union Agency for Network and Information Security. Challenges and opportunities for EU cybersecurity start-ups (2019)
- [70] L. Olejnik. A second life for the 'do not track' setting – with teeth. *Wired* (2019) <https://www.wired.com/story/a-second-life-for-the-do-not-track-setting>

## 8 Appendix

The image shows a screenshot of a questionnaire pop-up dialog. The title is "Fragebogen". Below the title is a progress bar. The main question is "Wissen Sie, was Browser-Cookies sind?". There are two radio button options: "Ja" and "Nein". Below the radio buttons is a text input field with the prompt "Falls ja, bitte beschreiben Sie kurz, was Browser-Cookies sind.". At the bottom right of the dialog is a blue button labeled "Weiter". At the very bottom of the dialog, there is a timestamp: "02.06.2019 - 06.06.2019".

Fig. 7. Pop-up with questionnaire.




Fig. 8. Original German version of the stimulus material.

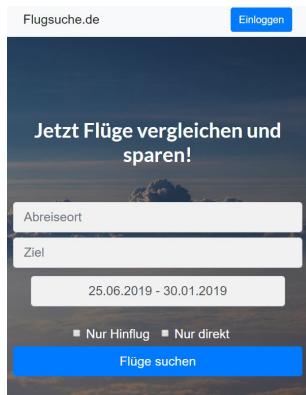


Fig. 9. Functional mock-up website offering flight search.

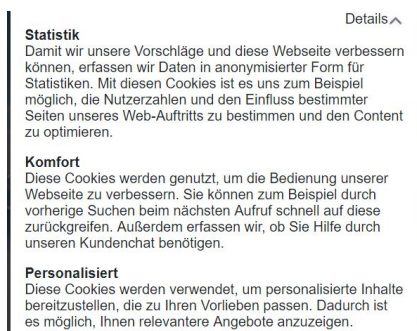


Fig. 10. German version of the expanded cookie dialog after clicking on “details”.

Table 7. Constructs and corresponding items in German.

Item	Item text (original)
<i>Perceived Deception (PDE)</i>	
PDE1	Die Seite ist bezüglich der Cookie-Einstellungen unehrlich gegenüber ihren Nutzern.
PDE2	Die Seite versucht Nutzer dazu zu führen, Cookie-Einstellungen zu wählen, die sie nicht wählen wollen.
PDE3	Die Seite benutzt irreführende Taktiken, damit Nutzer Cookie-Einstellungen wählen, die sie nicht wählen wollen.
<i>Perceived Difficulty (PDI)</i>	
PDI1	Es war unverständlich, eine Auswahl zu treffen.
PDI2	Es war frustrierend, eine Auswahl zu treffen.
PDI3↔	Es war einfach, eine Auswahl zu treffen.
<i>Regret (RE)</i>	
RE1	Ich bereue meine getroffene Auswahl.
RE2	Ich würde meine Auswahl ändern, wenn ich die Möglichkeit hätte.
RE3↔	Ich bin mit meiner Auswahl zufrieden.
<i>Privacy Attitudes (PA)</i>	
PA1	Der Schutz meiner Privatsphäre im Internet ist mir wichtig.
PA2	Wenn Webseiten Cookies verwenden, schränkt dies meine Privatsphäre ein.
PA3	Ich bin besorgt darüber, dass meine Privatsphäre durch Cookies von Webseiten eingeschränkt wird.

Items marked with ‘↔’ use inverted scales.