

Supporting Informed Choices about Browser Cookies: The Impact of Personalised Cookie Banners

Tom Biselli

Science and Technology for Peace and Security (PEASEC), Technical University of Darmstadt Darmstadt, Germany

Laura Utz

Science and Technology for Peace and Security (PEASEC), Technical University of Darmstadt Darmstadt, Germany

Christian Reuter

Science and Technology for Peace and Security (PEASEC), Technical University of Darmstadt Darmstadt, Germany

ABSTRACT

Browser cookies, especially those from third parties, pose a threat to individual privacy. While it is possible in principle to control the number of cookies accepted, this choice is often neither usable nor truly informed. To address this issue, this study used semi-structured interviews (N=19) to identify attitudes and user requirements to develop an alternative personalised cookie banner, which was evaluated in an online experiment (N=157). The cookie banner explanations were tailored to the privacy knowledge of three groups of users: low, medium and high. The online experiment measured cookie choices and perceived usability of the cookie banner across three groups: an experimental group that viewed the novel cookie banner with personalisation (personalised privacy assistant), a control group that viewed the novel cookie banner without personalisation (privacy assistant) and a control group that viewed the standard cookie banner provided by the website. The results indicate that the novel cookie banner (with or without personalisation) generally resulted in significantly fewer accepted cookies and increased usability compared to the standard cookie window. In addition, the personalised cookie banner resulted in significantly fewer accepted cookies and higher usability than the non-personalised cookie banner. These results suggest that tailoring cookie banners to users' privacy knowledge can be an effective approach to empowering users to make informed choices and better protect their privacy.

KEYWORDS

browser cookies, personalisation, user-centric, privacy assistant

1 INTRODUCTION

Nowadays, most people own and use a technical device to access the internet - and are therefore regularly confronted with cookie pop-ups. A cookie, also known as web cookie or HTTP cookie, is a small text package that is sent from the server to the browser and back when the server is accessed again. This is very useful and sometimes even essential to enable stateful functionalities in stateless protocols [30]. However, like most IT tools, cookies - in particular third-party cookies - have a dual-use potential: They can be used to track users' browsing activities, thus creating profiles and

identifying users [40]. As more and more applications move from desktop applications to web- and thus browser-based applications, the impact of cookies on user privacy is increasing.

While users are generally concerned about online tracking and want to remain as anonymous as possible, user behaviour is often different in reality - a discrepancy known as the privacy paradox [54]. A major factor here concerns the cognitive boundaries of users, as they often have only a limited ability to acquire, remember and process information. Most users try to resolve this by making trade-offs: disclosing some, but not all, data and thus trying to preserve their privacy. But with modern technologies such as big data analysis, it is easier than ever to collect and analyse huge amounts of private data, and users are often unaware of the possible consequences [27]. The discrepancy is not to be blamed entirely on the users. The main business model in today's online world is targeted advertising based on the collection of rich user profiles [14]. When the EU General Data Protection Regulation (GDPR) and the California Consumer Privacy Act (CCPA) were ratified in 2016 and 2018, respectively, websites were forced to give users a choice about cookies. As this threatened their business model, many companies developed cookie pop-ups to encourage users to accept all cookies, using specific interface designs such as *nudging* [18] or so-called *dark patterns* [23].

One way of dealing with this discrepancy is the use of privacy assistants (PAs) [13, 25, 33]. They facilitate the work for the user, providing simple choices and options in an easy-to-use way. Furthermore, PAs can provide a way to remind the user of their privacy attitude at the moment of decision, which has also been shown to reduce the aforementioned discrepancy [25]. Here, research also shows that using the same PA for all users is often not effective, as there are inter-individual differences between users [13]. Personalisation [27, 33] by tailoring privacy tools to individual users has proven to be an effective way to deal with the privacy paradox [25]. While personalised privacy assistants (PPAs) have been shown to be effective in the domains of app permissions [25, 33] and Internet of Things (IoT) devices [10], the use of both PAs and PPAs in the domain of browser cookies has yet to be explored.

With the aim of investigating whether PPAs are effective in enabling individuals to make informed privacy-relevant decisions, this paper investigates the research question: "Can privacy knowledge-based personalisation of cookie banners improve users' cookie choices?" To answer this, we (1) conducted 19 interviews with three privacy knowledge groups (low, medium, high) to *identify requirements* for a customisable cookie banner (see Section 3). We then (2) *designed and implemented* such a customisable cookie banner by tailoring the cookie explanations to the users' privacy knowledge

This work is licensed under the Creative Commons Attribution 4.0 International License. To view a copy of this license visit <https://creativecommons.org/licenses/by/4.0/> or send a letter to Creative Commons, PO Box 1866, Mountain View, CA 94042, USA.



Proceedings on Privacy Enhancing Technologies 2024(1), 171-191

© 2024 Copyright held by the owner/author(s).

<https://doi.org/10.56553/popets-2024-0011>

and tested it with a smaller sample of $N=39$ (see Section 4). Finally, we conducted an (3) *evaluation study* ($N=157$) to assess the effect of our novel cookie banner on cookie choice and perceived usability (see Section 5). Our results show that the novel cookie banner (PA) generally leads to fewer accepted cookies and improved usability. Furthermore, we show a specific effect due to personalisation (PPA) that leads to even fewer accepted cookies and higher usability compared to the novel cookie banner in its non-personalised version of the PA.

2 RELATED WORK

This section summarises relevant related work in the area of dark nudges in browser cookies, solutions to better manage cookies, and the potential of personalised support to enhance individual privacy.

2.1 The Landscape of Browser Cookies

In general, a lot of online user data is collected via browser cookies. In their 2016 study, Cahn et al. found that third-party cookies outnumber first-party cookies by a factor of two. They also found that less than one percent of the entities that place cookies can aggregate information across 75 percent of websites [9]. User interface design that acts as an intermediary between individuals and their cookie choices is, even in the most well-intentioned circumstances, restrictive of user choice [22]. Thus, users are often misled by the architecture of cookie banners [20]. Here, *dark nudges* come into play. Mathur et al. [37] define them as “interface design choices that benefit an online service by coercing, steering, or deceiving users into making decisions that, if fully informed and capable of selecting alternatives, they might not make”. One of the oldest and most influential categorisations of dark nudges is the “list of deceptive designs” by Harry Brignull [7]. Bösch et al. [6] use the term “privacy dark pattern”, essentially providing an updated taxonomy of malicious or deceptive design based on Hoepman’s privacy design strategies [24] and inspired by Brignull. While Brignull’s taxonomy is widely used as a reference in the user experience community, it has been criticised scientifically for conflating context, strategy and outcome, making it difficult to compare nudge types [17]. For this reason, Gray et al. [17] developed a new taxonomy based on an analysis of Brignull’s collection of examples and created and analysed a corpus of their own examples.

The entry into force of the GDPR in 2016 resulted in excessive use of dark nudges in cookie banners to ensure that users choose settings that allow the collection of vast amounts of personal data [37]. In a 2016 study, Machuletz et al. [35] demonstrated the effectiveness of dark nudges. When faced with a choice between an overwhelming selection of individual settings or a highlighted “select all”, users often opted for the latter option, even though this was associated with feelings of regret. The difference between one or three selection options had no significant effect. The fact that users who used the highlighted default button were less likely to correctly recall their consent led the authors to question the “morality and legitimacy of this design element”. In 2019, Utz et al. conducted a field study on more than 80,000 German participants analysing the influence of cookie banner design [61]. They found evidence that dark nudges such as pre-selected choices or a highlighted accept-button have a strong impact on users accepting

third-party cookies. In the same year, Trevisan et al. published a large-scale measurement campaign using the tool *CookieCheck* to automatically verify legislation violations [60]. They found that 49% of websites violate the ePrivacy Directive. In 2020, Nouwens et al. published an analysis of cookie banners [43]. They found that practices that are now illegal are nevertheless still common – for example practices such as pre-checking boxes or making it harder to opt out than to opt in. Vendors of cookie banner design even provide incentives for illegal cookie banners. Nouwens et al. also found evidence for the effectiveness of the dark nudge *Sneaking*: any control of information set deeper than the first layer of the pop-up window is effectively ignored by users.

In 2021, Kampanos and Shahandashti published a large-scale analysis of Greek and UK websites and cookie banners [26]. They found that while more than 60% of websites store third-party cookies, less than 50% show a cookie banner, thus more than 10% of websites do not comply with the law. They found little evidence for direct opt-out options and that an overwhelming majority of websites used dark nudges to influence users towards less private choices. The nudge most often found was *Interface Interference*: making accepting cookies much easier than rejecting them. Krisam et al. analysed the cookie banners of the 500 top visited German websites [29]. They found that more than 85% of the websites that allow the option to reject cookies visually nudge users towards acceptance and only 21.5% permit single-click rejection.

Taken together, this shows that the sovereign choice of browser cookies is often made difficult by dark nudges, among other things. Although not used by all websites, they are widespread and pose a significant threat to individual privacy by facilitating the uninformed, premature acceptance of too many browser cookies.

2.2 Cookie Management Support

Several solutions have been proposed and developed to overcome the inconveniences associated with browser cookies. In 2007 and 2010, Yue et al. published the scientific documentation of their *CookiePicker* software [64, 65]. *CookiePicker* is a Firefox browser plugin which categorises cookies into useful (actually used and contents changed by the website) and not-useful. It then proceeds to delete all cookies which are classified as non-useful. The entire process happens automatically and without user interaction. The advantage of this is the severe reduction of *Nagging*, as this leads to significantly fewer cookie banners. However, firstly, users cannot make an informed decision, and secondly, *CookiePicker*’s decision-making metric is the perceived change on the website, thus it can be circumvented, which renders the software useless.

In a 2012 survey on privacy enhancing web tools, Ruiz-Martínez gave a short summary on cookie managing options and their pros and cons [50]. The most extreme option, completely disabling all cookies, would lead to usability problems, as this would also include functional cookies. A tedious option would be to manually delete cookies at appropriate intervals. To be effective, the intervals would have to be short (e.g. after each session), which would be a nuisance for users. The use of plug-ins or other tools can be effective depending on intention and tool. The usage of anonymous web proxies is quite complex for the average user, but filters and processes the user’s data to make it anonymous. A simple option

used by many is to use the private browsing mode. Here, all cookies are automatically deleted as soon as the browser is closed. However, this has two main drawbacks: First, the cookies may be set initially and collect data during the session. Second, usability suffers as browsing history and other features are also disabled.

In 2015, Nosheen and Quamar proposed a cookie removing solution for android mobile devices [42], using the technology of self-destructing cookies. While this technology allowed a fine grained set-up, it required expertise that made it unsuitable for average users. Since then, several solutions to manage and reduce cookies have been developed. Browser plug-ins such as *I don't care about Cookies*, *Cookiebro* and *AdBlock* or more extensive systems such as *Ghostery* help users set cookies according to their preferences, but mainly offer a wealth of setting options and/or automatic technical support. These systems are mechanical tools of cookie management for which lay users lack understanding and knowledge and do not benefit from any gain in knowledge [51].

In 2018, Kulyk et al. proposed a concept for an interface including necessary information to allow even less tech-savvy users to make an informed decision [31]. Thus, they tackled the cookie problem not from a technical point of view, but with a usability approach. They used an assistant to ask users a set of questions about their preferences which then applied a general cookie setting to all visited websites. The concept was refined and evaluated with repeated feedback sessions and thoroughly analysed through a user study.

In summary, the benefits and dangers of cookies in general have been extensively discussed by the scientific community, but there is a lack of scientifically sound support. On the other hand, there are some browser plug-ins that make dealing with cookies much easier, but generally do not provide better information.

2.3 The Role of Personalisation

When trying to improve user privacy and security, a user-centred approach is usually recommended. Generally, most user interventions and privacy support tools are one-size-fits-all approaches and designed for use by average users only [13]. However, this does not take into account the fact that users and contexts of use are very different and can entail individual privacy requirements. Previous research has proposed that the “average user” might actually be a myth and that users differ substantially in their privacy preferences and needs [2, 5, 13, 16, 56]. There are also highly diverse usage contexts resulting in different requirements [32, 55]. Therefore, it is crucial to take into account inter-individual differences and different groups of end-users if truly effective support is to be provided. This is especially important regarding users who might be more vulnerable than others. While for some a personalised cookie banner might “only” be more convenient but not change the privacy implications, for others it might be crucial to inform them in a better and more tailored way so they can actually act according to their preferences.

It has previously been suggested that tailoring technology to the individual end-user, and thus recognising the vast inter-individual differences, might in fact lead to better outcomes [13]. This approach has now been tested in several privacy studies. In the area of mobile apps, it has been shown that users have a wide range of privacy preferences when it comes to granting permissions [34].

However, by condensing these diverse preferences into a smaller number of distinct groups, it may be possible to provide more specialised support to these user groups. In an applied study, a PPA was developed and used for the duration of one week to assess the extent to which personalised app permission recommendations were adapted. They demonstrated the effectiveness of these recommendations, which were personalised to pre-determined preferences, resulting in a high acceptance rate of around 80% of the proposed recommendations by the PPA [33]. Similarly, in another study, personalised privacy notifications drew attention to a discrepancy between general attitudes towards privacy and the actual granting of permissions for mobile apps. This successfully led to a higher alignment between attitude and behaviour [25].

In terms of nudging, tailored nudges in the area of passwords have shown promise [21, 46], while nudges in the area of privacy disclosure behaviour tailored to psychometric characteristics have not led to beneficial disclosure rates [62]. Finally, personalisation has also been studied in relation to IoT devices. For example, an interview study with 17 participants was conducted to assess the requirements for discovering and controlling the collection of potentially sensitive information from nearby IoT devices [10]. These requirements revealed three distinct user groups with specific preferences for the level of automation a PPA should provide in this area. The potential of PPAs was also highlighted in another study, which focused on the development of field PAs that aim to detect and inform about IoT devices in the users' vicinity based on an accurate understanding of different user needs [11].

2.4 Research Gap

Taken together, current cookie banners often make it difficult for users to make an informed and sovereign decision about which cookies to accept and which to reject. This can have a negative impact on user privacy and widen the gap between attitudes and actual behaviour. One obstacle here is dark nudges, which are widely used on the web [26, 29, 43, 60, 61]. This has led to a proliferation of software tools to assist users with cookie settings (see section 2.2). However, these tools are often technical and usually fail to provide personalised and relevant information to the user, and lack the ability to actively engage them in a meaningful deliberation process [42]. Due to inter-individual differences, browser cookie support should ideally be tailored to users privacy needs and knowledge. This approach has shown promise in mobile apps [25, 33, 34], nudging [46, 62] and IoT [10, 11]. These approaches included both the conceptual assessment of group-specific requirements [10] and the development of actual artefacts to evaluate the benefits of PPAs [11, 25, 33]. Overall, there are several approaches that harness the power of personalisation when individuals have to make privacy-sensitive decisions to protect their privacy and enable them to act more in line with their attitudes. However, the benefits of personalisation have yet to be explored in the context of customised cookie banners tailored to users' privacy knowledge. There are numerous studies on the design and perception of cookie banners and the wording of privacy notices [19, 38]. However, little research has been done on the wording of cookie banners tailored to users'

knowledge. Therefore, this study aims to fill this gap by investigating the impact of PAs and PPAs in the area of browser cookies to better inform and protect users' privacy.

3 PRE-STUDY: IDENTIFYING REQUIREMENTS

This section presents the pre-study with semi-structured interviews to determine what users expect from a browser cookie PA. Participants were classified into three privacy knowledge groups (low – medium – high). The data from the interviews was analysed to extract both general and group-specific requirements for the development of a novel, personalisable cookie banner.

3.1 Methods

In the following, the study procedure of the interview study, the recruitment and analysis are described.

3.1.1 Study Procedure. Initially, the participants completed an online survey on the platform *Sosci Survey*¹ which included demographic data and the Online Privacy Literacy Scale (OPLIS) questionnaire [36, 59]. The OPLIS questionnaire is well validated and consists of 20 items asking about users' knowledge of their online privacy and data protection. Although the OPLIS scale does not exclusively measure knowledge of cookies, it includes items related to cookies and other similar technologies. In addition, the implications of accepting or rejecting certain cookies are broader than knowing specific cookie definitions, and it is helpful to have a broader understanding of online privacy in general in order to make informed decisions. Answers were given either by choosing *true* or *false* for a statement or choosing the right answer to a question.

The questionnaire was followed by a semi-structured interview on the participants' ideas of an ideal cookie management with five more specific questions and one open question:

- (1) What do you think of the current options for setting cookies?
- (2) How could a different type of cookie banner best support you?
- (3) Do you have any suggestions for the design of an alternative cookie banner?
- (4) What are the minimum requirements for you to use it?
- (5) What would be an absolute nuisance?
- (6) Do you have any other suggestions or things you would like to contribute?

The first two questions served to introduce participants to the topic and to activate the relevant concepts. To approach potential improvements to the cookie banner from different directions, questions two to six focused on identifying requirements. While question two aimed to identify more general improvements for cookie banners, question three focused on specific design ideas. Questions four and five further attempted to decipher the key components of an improved cookie banner. All interviews were conducted via Zoom² and recorded locally for transcription. The questionnaire and subsequent interview lasted about 15 minutes.

3.1.2 Recruitment. Eleven participants were recruited through a convenience sampling method from the local (German) community, which included individuals who were known to the researchers

through personal and professional networks and reached via mailing lists. To further increase the diversity of the sample, an additional eight interview participants were recruited via the crowd-working platform *Prolific*³. The overall aim was to achieve a diverse age and gender distribution. We also made sure to interview participants with low, medium and high levels of education. Finally, although it was not used as a screening criterion, we also assessed the level of self-assessed IT proficiency using one item on a five-point scale ("How would you rate your IT proficiency?"). *Prolific* is a platform specifically designed to provide samples for scientific studies. Several studies have confirmed the reliability of *Prolific* and its ability to collect high quality and diverse data [1, 45]. A total of 19 German-speaking respondents participated in the study. Their ages ranged from 23 to 85, 12 were female and 7 were male, and their self-reported IT proficiency tended to be higher (see Table 7 for detailed demographic information). Respondents were compensated with €3, corresponding to the German minimum wage at the time of the data collection.

3.1.3 Ethics. The study was conducted in accordance with the requirements of the local ethics committee at our university. These requirements include the avoidance of unnecessary stress, the exclusion of risk and harm, and the anonymisation of participants. Personal information collected was limited to age, gender and education. Sensitive data (e.g. ethnicity, religion, health data) was not collected. The data was collected on the platform *SoSci Survey*, whose servers are located in Germany and who store the data in accordance with the GDPR⁴. Participants were not misled but were given transparent information about the procedure and aims of the study. They also had the opportunity to stop the interview at any time without giving a reason. They then gave their informed consent to participate. The data was secured and processed in accordance with the data protection provisions of the GDPR.

3.1.4 User Clustering. A threefold division of the OPLIS results was conducted to separate users into three privacy knowledge groups. Using the norm table for the total German population provided by the OPLIS manual⁵, the user groups were created based on percentile ranks (see Table 5).

3.1.5 Analysis. The interview transcripts were first roughly divided into (1) opinions on current cookie setting options (Q1), (2) feature requests (Q2, Q4, Q6), (3) UI or design ideas (Q3), and (4) nuisances (Q5), with the vast majority of statements falling under (2). In a second step, these passages were grouped by similarity and summarised as concisely and accurately as possible, retaining specific wording where possible. Based on these units, each of these groups of similar units was given an overarching header (e.g. *well-structured information* or *simple language*, see Table 12 for all categories), which served as the coding categories for the desired features of the novel cookie banner. After further iterations using these coding categories to ensure the appropriate mapping of the relevant units across participants, they were categorised by privacy knowledge groups. This resulted in an ordered list of features per user group, which also indicated the priority of a feature based on

¹<https://www.socisurvey.de>

²<https://zoom.us/>

³<https://prolific.co/>

⁴<https://www.socisurvey.de/en/privacy/>

⁵German version available here: https://oplis.de/docs/OPLIS_Manual_deutsch.pdf

the number of similar requests (i.e. of the same coding category) in that group. These coding categories were then compared across all three groups to extract group-specific feature requests as well as overarching features. The coding process was carried out by one researcher.

3.2 Results

Of the 19 participants, ten showed high privacy knowledge according to their OPLIS score, four showed medium privacy knowledge and five showed low privacy knowledge. In line with previous work (see Section 2), participants were generally annoyed and/or overwhelmed by the current cookie setting options, and mostly aligned their behaviour according to convenience.

3.2.1 General Findings. All participants displayed negative reactions to the current options on cookie settings. The most common reaction was a heartfelt “*too complex*” or “*too much effort*” (participants H5, H7, H8, H11, L12, M13, H14, M15). Other participants described the current experience as plain “*annoying!*” (participants H1, H6, L9, H10). Participants particularly complained about the number of clicks necessary and the (deliberately) confusing menus. Additionally, some participants also remarked about the hiding or obscuring of information. Several participants (H2, M3, H11) remarked that they experienced a psychological push towards accepting as many cookies as possible.

Speed and Simplification. When participants were asked how a software product could best support them, their main requests were simplicity, usability and speed. Participants wanted to be presented with simple, understandable options and finish their settings with one click. H7 and M13 remarked that the cookie settings hinder them in their online uses. Most participants stated that while they inherently want to consent to as little cookies as possible, they are not willing to navigate through numerous sub-menus to do so. They frequently choose the quickest route and accept all cookies instead. It was no surprise that the most common suggestion was to have a software product which would simply have all cookie options as one-click options. The options participants listed can be summarised as: (1) allow all cookies, (2) legitimate interest or advertisement cookies, (3) only necessary cookies and (4) no cookies at all. However, simplification is not the same as the rejection of cookies. Several users were aware that certain cookies greatly facilitated their daily lives. They were, e.g., fond of using automatic completion of common forms, such as address forms, in online shopping. Additionally, they wanted to keep their shopping cart contents and their logins for their most frequently used services. When thinking about the product, they suggested extending the current browser functionality to delete all cookies. They wished for a way to except some (useful) cookies from the “delete all” option. H5 remarked that when deleting all cookies, it would be great to keep the “initial cookie” that stores cookie consent information. This way, they could get rid of all cookies acquired by mistake without having to reset their settings on every page.

Better Information. While participants wanted quick one-click buttons for different settings, they also wanted to be able to find out the details of each setting – especially given that different providers use varying labels. Additionally, participants requested

in-depth information about each cookie, what data it stores and which advantages and disadvantages this entails. Participants were also interested in the usage and propagation of their data. They requested to know the type and extent of data stored and also a list of potential receivers and processors of their data.

Participants, especially from the lower privacy knowledge groups, repeatedly remarked that they have difficulties following the currently available information texts. Some guessed that the texts were deliberately difficult to obscure the truth about what happens with their data. Others stated that they need the information to be presented in their mother tongue (here German) and without technical jargon. L9 especially remarked that terms like *browser* and *cookie* confused them and required further explanation.

Nuisances. During the interviews, participants were asked if there was anything that would be an absolute nuisance. The first set of responses referred to bugs and malfunctions. Participants explicitly stated that they would not use the product if it contained bugs or made the process complicated. Furthermore, they stated that the product must be compatible with all websites. Another set of answers related directly to the need for speed, simplicity and usability. Finally, it was frequently mentioned that advertising within the cookie banner would not be acceptable. Participants generally listed the opposite of what they requested as nuisances.

General Findings

- dissatisfaction with the current options for cookie settings
 - core requirements for alternative cookie banner:
 - simplification, usability and speed
 - comprehensive display of information on cookies, data stored and potential receivers and processors
 - short, concise and well-structured display of information
-
-

Table 1: General findings across user groups in the pre-study.

3.2.2 High Privacy Knowledge Group. The following findings are specific to participants belonging to the high privacy knowledge group.

Automation – Set Cookies Once and Apply on Every Page. High knowledge participants spend a lot of time online and thus encounter many different websites and their respective cookie pop-ups daily. So while they insisted on always having the *option* to check, they also requested automation to speed up the process. The general consensus (mentioned by eight out of ten participants) was to set some general settings once and use these automatically on all websites, without disturbing the users’ workflow. An important benefit of this was raised by H2: “*And I think if you only do it once instead of having to do it on every website, you might take a little more time to go through these things and think about them.*” Others requested a pop-up along the lines of *I have set only the necessary cookies as per your request*, with an option to check or agree (H12). H2 suggested a variant which would perform the settings automatically in the background and only acknowledge these with a short message in the status bar. To ensure an option to check and adjust,

the participant suggested to have a small icon to the right of the address bar, similarly to the icon displayed by the AdBlock-Plugin or the downloads information. This icon should include the number of currently active cookies and act as a button to access the settings.

Well-Structured Information and Short Texts. Besides the desire for an automated solution, the present description of cookies was criticised by many interviewees. A common requirement was identified in a clearer description of the different cookie types with a focus on short and concise texts (H2, H3, H4). H4 illustrated this: “*There is such a long text, I have no time to muddle through it. So what do I do? Yes. Accept all.*” Similarly, H3 supported the need for more conciseness: “[...] *that it is described as briefly and concisely as possible, for people who do not know.*”

Temporary Cookie Consent. Most participants in the high group rejected all non-necessary and third-party cookies, some however enjoyed the merits of third-party cookies. The main focus here was on personalised advertisements, but the use of websites that require third-party cookies to function was also mentioned. Several participants stated that they enjoy and actively use personalised advertisements while searching for a specific product and deciding which type or brand to buy. However, once they decided on a product and bought it they were annoyed by the advertisements. One participant just asked why the advertisements do not automatically stop once a purchase has been made. Others suggested introducing an expiry date for cookies. This would allow them to enjoy the merits without being constantly tracked or continuously presented with a product they are no longer interested in. The time frame given varied greatly and lasted from allowing third-party cookies for the duration of the active session only (H10), three days (H7) or up to several weeks (H6).

Compatibility. Participants stated that they work with different operating systems as well as different browsers and wanted the product to be compatible with all of them and to synchronise settings. H2 stated: “*So maybe, it would also be cool if this could somehow be overarching. So that you don’t have to do it again for every browser.*” Participant H10 remarked that the best way to do this would be for the browsers to implement a standardised API to govern the cookie exchange with websites. The product could then also access the API and thus regulate the cookie settings. H10 also stated that they would appreciate it if such a product would, by default, come with the browsers.

Additional Finding: Suspicion Towards Developers. Subjects in the high knowledge group demonstrated heightened skepticism towards the product, indicating concerns about the credibility of the proposed developers and operators of a novel cookie banner. A number of participants emphasised the importance of full compliance with the GDPR as a prerequisite for trustworthiness. Furthermore, they expressed a preference for products originating from reliable sources, such as open-source solutions.

3.2.3 Medium Privacy Knowledge Group. There are clear differences between the highest and both the middle and low group. While participants from the high group moved confidently and effortlessly on the internet, participants from the middle and low

group felt rather uncomfortable and insecure. Generally the accounts from the medium and low group were very similar.

Simplicity and Usability. Participants from the medium group admitted that they were overwhelmed by the current cookie options. Some indicated that they did not understand the commonly used terminology (*browser* or *cookie*). Others were not able to distinguish between cookies and online advertising. M15, e.g., was confused why some websites allow users to reject all cookies and others force them to accept “technically necessary” cookies. Other participants stated that they would want to adjust (more) cookie settings, but were not capable to do so. M16, for example, stated “*Because I find many sites always push you to just accept everything. So a program that simply supports you so you don’t have to accept everything on some sites would be desirable because sometimes I don’t find the option.*” Participant M13 usually agreed to all cookies as they thought that without agreeing they would compromise the (core) functionality of the visited website. Thus, an overarching need for easier operability with clearly marked choice options was a central theme in the interviews here.

Simplicity Through Simple Language. Participants in the medium group also had trouble understanding the explanatory texts provided. They generally attributed this problem to the use of technical jargon in information texts and deliberately confusing setups of cookie pop-ups. To be able to adjust the cookie settings, these participants requested clear, simple language without words borrowed from foreign languages and a neat setup. This is well captured by M3: “*And what I would of course prefer if these terms ... If for these terms there were always the German or the national language terms. [...] Browser, cookie, what do I know. I always have to think about it [...] what is the difference between, between this and that.*”

3.2.4 Low Privacy Knowledge Group. Similar to the medium group, participants in the low group stated that they felt uncomfortable online. They too did neither understand the terminology nor were they familiar with the common controls. Generally, participants from this group hesitated to make suggestions as they felt unqualified to do so.

Simple Design and Well-Structured Simple Information. Participants from the low group stated that they were overwhelmed by the current options and did not know how to adjust cookie settings. Participant L12 admitted to not knowing which setting entails what. They just chose “agree” without further examining the provided details as they were usually confused by the terminology and structuring used there. As a result, participants in the low group consistently expressed a desire for a simpler design and easily comprehensible information (L9, L12, L18, L19): “*... I personally always like it when it’s somehow plain and simple and that you immediately have a good overview.*” (L19)

Additional unique requirements. Participants in the low group reported limited experience with the internet and software. Unique requirements were raised by these participants, such as participant L12’s suggestion for a step-by-step walk-through tutorial to assist with understanding the functionality of control options. Additionally, participant L19 requested a feature such as an audio button accompanying informational texts to provide an alternative means

of explanation. However, since those were mentioned only once, they were not considered a core requirement across individuals.

Group	Core requirements
High (N=10)	<ul style="list-style-type: none"> • automation through pre-set general settings in the background • well-structured information and short texts • temporary cookie consent • compatibility across browsers and operating systems • reliability/open source development
Medium (N=4)	<ul style="list-style-type: none"> • simplicity and usability: simple one-click buttons for main choices • less use of technical terminology • more and well-structured simple information
Low (N=5)	<ul style="list-style-type: none"> • simple design • more and well-structured simple information • less use of technical terminology

Table 2: Overview of core requirements for the three privacy knowledge groups.

3.3 Conclusion

When reviewing the interviews, it was clear that the wishes of the participants were quite similar across all three groups. They demanded usability, simplicity and comprehensible information. All participants were either annoyed or overwhelmed by the current options, and many remarked that these served the interests of the page providers rather than the users. High knowledge participants had a reasonable to great understanding of cookies but were annoyed with the repetitive action of setting cookies. They generally suggested more nuanced features such as automation and temporary cookie storage. However, they also requested generally better structured and consistent and, above all, short textual information. These participants also displayed suspicion towards online products and requested trustworthy developers and open-source-based support. Participants from the medium and low groups were generally overwhelmed and confused. They wanted to get rid of cookies altogether and sometimes did not understand that some of them were necessary. These participants had more difficulties with complex cookie banner designs with several sub-windows, the technical terms used and generally requested more, yet simplified information about the functionality of the cookies they were about to set.

4 DESIGN AND IMPLEMENTATION

Due to mainly homogeneous and rather unspecific interface wishes across all user groups, it was decided to use one interface design for the novel cookie banner. The general requirements across all

groups were incorporated here. The personalisation approach involved tailoring the content, particularly the explanations of the cookie types, to the users' level of online privacy knowledge. This decision was based on the interview findings according to which individuals with medium and low privacy knowledge expressed a need for clear and simple explanations, while those with high privacy knowledge indicated a preference for more concise information. The high knowledge respondents in particular also wanted an automated process without having to explicitly decide for each website. However, this would have meant that we would have had to develop a completely separate automated solution for this group, which would have prevented us from directly comparing similar cookie banner versions in different groups. Furthermore, well-structured information and short texts were also the second most frequently mentioned requirements in the high group. After careful consideration, we therefore chose to prioritise a consistent interface design, while differentiating the explanation text to align with individual privacy knowledge levels. This approach allowed us to draw more nuanced conclusions about the impact of personalisation and avoid the potentially confounding effects of fundamentally different interface designs or fundamentally different technical solutions. By limiting the sources of variation to the personalised explanation text, we were better able to isolate and evaluate the personalisation dimension.

4.1 Cookie Assistant Design

This subsection describes the design principles that were considered for creating a more user-friendly cookie banner.

4.1.1 General Design.

Frame. The position of the cookie banner is potentially significant. A banner that obscures the website would lead to more interaction, provided it is integrated in suitable places (i.e. on the left or bottom of the page) [61]. The pre-study resulted in no clear preferences in this regard. To force an interaction with the cookie banner, it was positioned in the middle (attention grabbing and page obscuring) and utilised an interaction blocking overlay. Thus, users were forced to interact.

Setting Options. The main feature request during the pre-study was to have a *reject all* (or *no cookies at all*) option directly on the main pop-up. Users also wished to reduce effort, mainly the number of clicks, to reach their goal and speed up the process. Studies on *choice proliferation* show that users constantly balance between information and control and feel overwhelmed by too many options [10]. However, the difference between one or three options is negligible [35].

To combine all this, the cookie banner had three options for users to choose from: accept only necessary technical cookies; accept technical and marketing cookies; accept all cookies. To not nudge and influence users' choices, all options had the same design and were reachable with the same number of clicks. As several participants in the pre-study stated that they actively allow marketing cookies, this option was specifically incorporated. Besides, this is also an option present in many of the existing cookie banners.

Number of Pages. A 2020 study on cookie banners [43] stated that second pages are effectively ignored. Our pre-study also revealed that users were confused and annoyed by multi-layered menus. Thus, the designed cookie banner included no second pages. All information was presented clearly and as concisely as possible on the first page.

Clear and Well-Structured Information. Participants of the pre-study also requested clear information on their choices. Depending on the group, the focus was either on simple wording and examples (low group), information on demand (medium group) or a short and precise overview (high group). Participants strongly disapproved of too large amounts of text. So the cookie banner had an explanatory text for each offered choice. These texts were personalised to the group the user had been assigned to (see Section 4.2). For the low and medium groups additional examples were available as drop downs (as proposed by Kulyk et al. [31]). This way, they did not obscure the general options but were available on demand.

Error Prevention. Several participants of the pre-study stated that they are sometimes unfocused and just click something to get on with their task. This led to them choosing cookie settings they did not intend. In addition, some complained that there was no way to go back after clicking. Therefore, a confirmation pop-up has been added to the novel cookie banner. It clearly shows the chosen settings and offers the option to accept or change them.

Colour Scheme and Icon. When designing the optical aspects of the cookie banner, the focus was on a modern but simple design. This fulfilled Nielsen’s [41] eighth heuristic and was achieved by using the preset design of the utilised Ionic 6 framework and following standards used in cookie banner design. Reinheimer et al. found that users prefer a subtle colour scheme over a gray scale design [49]. Having a signature colour also aided users in recognising that they were dealing with the assistant instead of a website pop-up. To meet these criteria, the cookie banner utilised a simple cookie icon as well as a muted blue and petrol colour scheme.

Usability and Privacy Enhancing Design. Participants have repeatedly requested usability. In order to ensure this, the prototype was designed in accordance with the usability heuristics by Nielsen [41]. Additionally, the design followed the suggestions by Terpstra et al. [58]. The cookie banner buttons offered users meaningful controls to set their cookies quickly, yet at their own discretion. Additionally, the confirmation pop-up provided a choice.

Comparison with Existing Cookie Banners. Existing cookie banners tend to be inconsistent in design. They vary in the number of initial choices and often use dark nudges (see section 2.1). They also often have multiple sub-levels and only offer two choices on the first level, ‘Accept all cookies’ and ‘Customise settings’. We provided three options: (1) ‘Accept technically necessary’ and (2) ‘Accept all’ were chosen to allow both extremes, i.e. accept all and none, in principle. In order to provide an additional choice on the first level – since second pages of cookie banners are usually ignored [43] – we took our cue from the pre-study in which several participants indicated that they specifically allow (3) marketing cookies. Thus, our novel cookie banner differed from existing cookie banners mainly in the following dimensions:

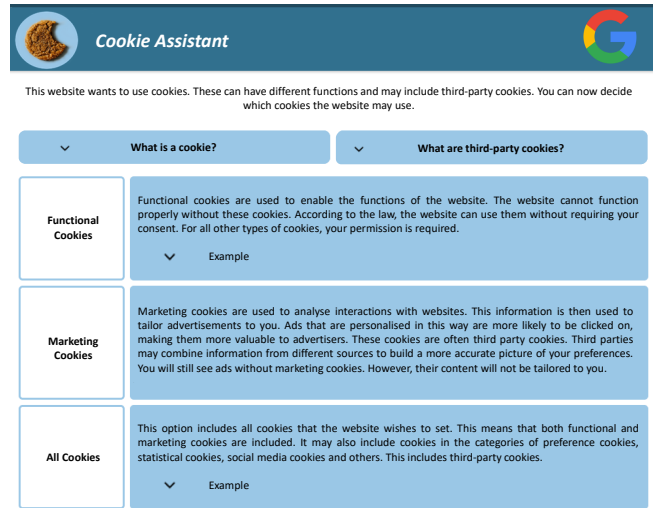


Figure 1: Screenshot of the Cookie Assistant for a user from the medium group (English translation). The versions for the low and high group differed in the explanation texts – see Table 4 for the corresponding texts

- (1) clear explanation of what cookies are, rather than just presenting the purposes for which they are used
- (2) more detailed explanations of cookie types depending on the level of privacy knowledge, e.g. why some are technically necessary
- (3) three choices on the first level
- (4) no nudges, consistent coloring of choices
- (5) choice confirmation button

4.2 Personalisation

Studies exist on the design and perception of cookie banners and on the wording of privacy notices [19, 38]. However, little research has been done on the wording used in cookie banners. While intentionally biased vs. non-biased text in cookie banners had no effect on user choice [4], the effect of personalised and overall neutral wording has not been studied before. Thus, the wording of the information texts was based on the examples given on best-practice websites⁶⁷ and the wishes expressed by the interviewees. Moreover, the texts were written in German and according to specific rules (see Appendix A.1 for details). Among other things, one main rule for creating the texts was that they only differed in explanations and examples related to privacy knowledge so as not to introduce other potentially biasing variables. As different explanatory texts could also act as nudges, we generally focused on creating neutral texts to allow users to form their own opinions and make an informed choice. The inclusion of pros and cons allowed users to form their own opinions and prevent reactance. The texts were iteratively reviewed and refined according to the rules (see Appendix A.1) by

⁶www.cookiebot.com – provider of cookie banners, allows check of GDPR compliance, focus on usability and GDPR compliance

⁷www.gdpr.eu – resource for organisations and individuals researching the GDPR. Includes a library of straightforward and up-to-date information to help organisations achieve GDPR compliance.

two privacy researchers, taking into account the comments of two pre-study participants (H1 and L12).

All groups essentially needed well-structured and clear information. This included simple information without technical terms, especially for the medium and low privacy knowledge groups. The requirements of the medium and low groups were particularly similar. Nevertheless, we decided to offer three different levels of personalisation. The main reason for this was that although they have similar desires, such as well-structured and simple information and less use of terminology, they are similar only at an abstract level. On a more specific level, the understanding of what is considered comprehensible information varies considerably depending on the current level of privacy knowledge. For example, percentile rank 10 (low privacy knowledge) is very different from percentile rank 60 (medium privacy knowledge). Consequently, the different texts have been worded to reflect the different levels of prior knowledge in this area and we have decided to retain the three-way segmentation. Thus, for different types of cookies, different explanations were developed for each of the three groups. In general, for the high group, short and concise explanations were developed as a quick refresher. For the medium group, this information was expanded to include illustrative examples of the use of cookies, including examples of useful cookies, functional cookies, marketing cookies and information on their functionality. The information that third parties set their cookies on multiple pages and use them to track users was added because users in the medium group showed confusion about this process in the pre-study. For the lowest group, the wording was simplified and additional (real-life) examples and more detailed explanations were given, e.g. regarding login, the shopping basket and the need for functional cookies (see Table 4 in the Appendix for the exact texts for all three groups).

4.3 Implementation

For implementation, a browser-based mock-up prototype was developed: Here, access to seven websites was simulated where decisions about browser cookies had to be made.

4.3.1 Technical Details. The mock-up for the study was developed as an Ionic 6 desktop application. This was chosen for its modern interface design and scaffolding support, which accelerated the creation of the various pages required. This meant a combination of Typescript for the back-end logic and SCSS-style HTML for the front-end. For the experiment, the software was wrapped in a Docker container and hosted on the faculty server.

4.3.2 Cookie Choices on Simulated Websites. The access to seven websites was embedded one after the other. A screenshot of the actual website was used as a background image. Depending on the participant group, the view of the website was overlaid with the respective cookie banner. Depending on the choices made in the cookie banner, the number of cookies selected by the user was stored. To estimate the number of cookies accepted, the cookies set for each choice were read from the original cookie banner and categorised (functional, marketing, etc.) using various cookie explanation websites (e.g. *Cookie-Script*⁸). In this context, the pages of the simulated native cookie banner also contained the entire setup

⁸www.cookie-script.com

of the native cookie banner, including navigation between different layers and the number of cookies set with an option. This entailed extensive additional SCSS styling as all native cookie banners followed different themes. Once the user had made their choice, they were redirected to the next simulated website access.

Similar to previous research, we included the most popular websites in orientation to three publicly available lists of the most visited websites in Germany⁹. In this way, a total of 10 websites were selected (see Table 6). These were reduced by three websites because they did not have a cookie pop-up, because they only use functional cookies or because they use a so-called paywall system where users can either accept all cookies or subscribe to the website.

4.3.3 Initial Revisions to the Cookie Banner. In order to validate our study protocol for the evaluation and to initially assess the novel cookie banner, we conducted a small study with N=39 participants. As in the pre-study (see Table 7 for demographic information), the sample was drawn from personal and professional networks from the university context. The study procedure was analogous to that described in Section 5.1.1 and allowed for feedback to be given via open comments. These open comments were coded by one researcher into the two categories (1) neutral/positive and (2) potential for improvement. In a subsequent step, the statements with potential for improvement were categorised into similar statements such as "Too much text" or "Non-appealing design". Here, the majority of respondents were satisfied with the design of the cookie banner and the accompanying explanatory text (see Table 9). The criticisms mentioned were the size of the pop-up window, a non-appealing design and the length of the text. However, this was only mentioned by a minority of respondents (N=2-4). Overall, because the feedback was rather (1) unspecific, (2) positive in general, and the refinement study was mainly used to evaluate the (3) experimental procedure, no fundamental changes were made to the experimental design and the cookie banner. The purpose of the small study was explicitly not to try out different versions of text explanations and then choose the best ones, but rather to validate the overall approach for the evaluation studies and to identify individual improvements. The changes were therefore limited to (1) a slight reduction in the size of the cookie banner, (2) the inclusion of the logo of the visited site in the header (to make the design more appealing), (3) a shortening of the cookie explanations (less than 10 %) and (4) the provision of drop-down menus for more detailed explanations.

5 EVALUATION

This section details the study conducted to test the novel, personalised cookie banner developed based on the results of the pre-study (see Section 3.2). Participants were randomly assigned to three different cookie banner version groups. One control group viewed the (1) *original page cookie banner* (CG), a second control group viewed the (2) *novel cookie banner without personalisation* (PA) and a third experimental group viewed the (3) *novel cookie banner personalised to their privacy knowledge* (PPA). The personalised group

⁹We only chose websites appearing on all three lists on www.semrush.com, www.ugwire.com and www.similarweb.com. Websites that appeared multiple times with different top-level domains, such as google.com and google.de, were combined to only use the German domain.

thereby viewed the novel cookie banner version with explanations tailored to their privacy knowledge (low, medium, high) while the non-personalised group viewed the novel cookie banner with explanations from one of the two versions not congruent with their privacy knowledge.

5.1 Methods

5.1.1 Study Procedure. During the study, participants switched between a questionnaire hosted on *SoSci Survey* and the online experiment website. After providing some demographic information, participants were presented with the OPLIS questionnaire. Upon its completion, the results were automatically scored and the participant was assigned to a low, medium or high privacy knowledge group. Participants were then asked to set cookies on the seven websites. Depending on the participant group and OPLIS group, users were provided with different interfaces and information to set cookies on the seven websites, using either the PPA, PA or standard cookie banner of the respective website. After setting cookies on one site, they were automatically redirected to the next until all sites were completed. Participants were then presented with a follow-up questionnaire in which they were asked to complete the System Usability Scale (SUS) [8].

5.1.2 Sample Size & Recruitment. Based on an expected medium effect size of $d = .5$, an α -level of $.05$ and a desired statistical power of $.8$, we calculated the optimal sample size to be $N=157$ and recruited those participants via *Prolific*. Particular attention was paid to ensuring a diverse sample in terms of age, gender and education. We also collected information about technical affinity using the Affinity for Technology Interaction Short Scale (ATI-S) [63]. Participants' ages ranged from 19 to 72 years, 61 were female, 93 were male and 3 were diverse and their technical affinity tended to be higher (see Table 7 for detailed demographic information). The results of an ordinal regression model predicting the OPLIS scores on the basis of these variables showed that both higher education (i.e. having an academic degree vs. having lower secondary education) and higher ATI-S scores were associated with higher OPLIS scores (see Table 11 for the regression results). The experiment itself was conducted online and participants were paid €3 for their 15 minutes of time.

5.1.3 Ethics. Like the interview study, this evaluation study was conducted according to the guidelines of the local ethics committee at our university, and the same statements apply here (see section 3.1.3).

5.1.4 Statistical Analysis. We hypothesised that the personalised novel cookie banner in particular would be associated with fewer accepted cookies and a better usability, as it would more accurately address user needs through tailored information and a user-centric development process. As the non-personalised version also met the latter criterion, we also hypothesised that it would be superior to the standard cookie banner. Specifically, we assessed differences in the three groups (personalised novel cookie window (PPA), non-personalised novel cookie-window (PA) and standard cookie windows) regarding the dependent variables (1) number of accepted cookies and (2) SUS scores. Due to the violation of statistical assumptions such as normal distribution, we applied

the non-parametric Kruskal-Wallis test to assess differences between the three experimental groups with subsequent Dunn's Post Hoc test with Benjamini-Hochberg correction for multiple comparisons [3]. The dependent variable "number of accepted cookies" was therefore the total number of cookies accepted summed across all seven sites. The lowest possible number (selecting only technically necessary cookies on all sites) was 65, while the highest possible number possible (selecting all cookies on all sites) was 259.

5.2 Results

In general, the OPLIS scores were heavily skewed towards high scores resulting in the majority of participants falling into the high privacy knowledge group. Of the 157 respondents, 113 showed high privacy knowledge, 35 showed medium privacy knowledge and only 9 showed low privacy knowledge. This prevented us from making well-founded direct comparisons between all privacy knowledge groups due to the small sample size, particularly in the low privacy knowledge group (see Table 8). This limitation is discussed in Section 6.4.

5.2.1 Does the novel cookie banner in general, and the personalised version in particular, lead to fewer cookies being accepted than the standard cookie banners? To assess the overall group differences between the PPA group, the PA group and the standard cookie banner control group in terms of the number of cookies accepted, we performed a Kruskal-Wallis test. This test revealed a significant difference in the number of accepted cookies between the three groups ($\chi^2(2) = 43.725, p < 0.001$). Post-hoc pairwise comparisons using Dunn's test were performed to examine the differences in the number of accepted cookies between the groups. The results showed significant differences between all groups at the $\alpha = 0.05$ level. That is, both the PPA group (mean = 76, median = 65) accepted fewer cookies than the control group ($Z = 6.51, p < 0.001$) and the PA group (mean = 99, median = 65) accepted fewer cookies than the control group (mean = 173, median = 231) ($Z = 4.20, p < 0.001$). Importantly, the number of accepted cookies was also significantly lower in the PPA group compared to the PA group ($Z = 2.15, p = 0.03$), thus revealing a specific effect for the personalisation. The effect size here was medium ($r = .5$). An exploratory inspection of interaction effects and group differences between subgroups revealed some statistically significant differences, for example between control group (CG)-high vs. PA-high and CG-high vs. PPA-high (see Table 13 for all comparisons), mainly due to a comparatively high number of accepted cookies in the CG-high group. However, due to the skewed sample sizes in the different groups, these are only considered exploratory results.

5.2.2 Does the novel cookie banner in general, and the personalised version in particular, lead to an increased usability compared to the regular cookie banners? To assess the overall group differences between the PPA group, the PA group and the standard cookie banner control group in terms of perceived cookie banner usability as measured by the SUS, we again ran a Kruskal-Wallis test. This test revealed a significant difference in perceived usability between the three groups ($\chi^2(2) = 29.804, p < 0.001$). Post-hoc pairwise comparisons again showed significant differences between all groups at the $\alpha = 0.05$ level. Both the PPA group (mean = 83, median = 87)

Experimental Group	Privacy Knowledge		
	Low	Med	High
PPA	70 (120)	65 (87)	65 (71)
PA	145 (126)	66 (120)	65 (84)
CG	154 (154)	189 (165)	231 (176)

Table 3: Median (and mean) values of overall accepted cookies per participant. PPA = Personalised Privacy Assistant, PA = Privacy Assistant, CG = Control Group

reported higher SUS scores for the cookie banner than the control group ($Z = -5.43, p < 0.001$) and the PA group (mean = 72, median = 77) reported higher SUS scores for the cookie banner than the control group (mean = 63, median = 67) ($Z = -2.18, p = 0.03$). Again, there was a specific effect for personalisation as the SUS scores were significantly higher in the PPA group compared to the PA group ($Z = -3.12, p < 0.01$). The median SUS scores across the different sub-groups (see Table 10) suggest that although the number of accepted cookies is lower in the low privacy knowledge PPA group compared to the low privacy knowledge PA and CG groups, perceived usability is still lower (63) than in the medium (82) and high privacy knowledge PPA groups (85). However, the decision to accept a certain number of cookies has many determinants, of which perceived usability is only one.

Conclusion. The results indicate that the implementation of the newly developed cookie banner leads to a reduction in the number of accepted cookies (see Table 3) and an improvement in usability (see Fig. 2). Notably, the personalised version of the cookie banner elicits a specific effect, demonstrating both a decrease in the number of accepted cookies and an enhancement in usability compared to the non-personalised variant.




Figure 2: Usability measured via the SUS in the different experimental groups. CG = Control Group, PA = Privacy Assistant, PPA = Personalised Privacy Assistant

6 DISCUSSION

Cookie choices can have a significant impact on a user’s privacy, and allowing more cookies than necessary can compromise it. This

study sheds light on a user-centred development process of a cookie banner tailored to the privacy knowledge of users.

We used semi-structured interviews to explore key attitudes towards cookie banners and requirements for improved, alternative cookie banners. Respondents confirmed a general dissatisfaction with current cookie banners, which are often considered difficult to understand, often use dark patterns to hide more privacy-friendly choices, and are often overly complex with multiple sub-windows. In order to provide personalised support to users, we specifically researched the requirements of three groups of users with low, medium and high levels of privacy knowledge. On this basis, we developed an alternative, personalised cookie banner, taking into account the requirements from the interviews, but mainly consisting of different levels of detail in the descriptions of the different types of cookies and the consequences of accepting them.

Following a small prototype study ($N=39$), which resulted in minor refinements to the cookie banner, we conducted an evaluation study ($N=157$). This showed that the cookie banner we developed resulted in both fewer cookies being accepted and improved perceived usability compared to the standard cookie banners. Importantly, we were able to show a specific effect for the personalised version of the novel cookie banner, which led to significantly fewer accepted cookies and significantly higher perceived usability compared to the non-personalised novel cookie banner. Thus, taken together, our main research question, whether privacy knowledge-based cookie banner personalisation improves users’ cookie choice, can be answered in the affirmative. Although browser cookies have their place and can simplify the user’s online experience (which was indeed sometimes mentioned positively in the interviews, e.g. in the context of online shopping), our results suggest that the current standard cookie banners do not sufficiently inform users with different backgrounds about the consequences of their choices and do not enable informed consent.

Our main contribution therefore lies in (1) providing a user-centred approach to cookie banner design, from which general design implications for cookie banners can be derived. Furthermore, (2) we show the potential of personalisation in the cookie banner domain to improve user privacy, beyond the effect of “just” overcoming dark nudges.

6.1 Countering Dark Nudges

In our study, dark nudges were sometimes used within the regular cookie banners (four out of seven pages). Our aim was not just to explicitly overcome dark nudges, but to compare our alternative cookie banner with the typical internet experience, where users are occasionally faced with dark nudges. Although the study did not specifically examine the impact of the novel cookie banner compared to cookie banners with dark nudges, the results suggest that the banner may have the potential to counteract such nudges. This is evidenced by the reduction in the total number of cookies accepted. It has been shown that dark nudges are very common in browser cookies [29, 37, 43]. Therefore, automatic detection of dark nudges is also potentially promising. However, it remains a challenge to do this accurately, partly due to the large number of types of dark patterns [23, 53]. Furthermore, detecting dark patterns is only the first step, the second step of correcting them or

providing user-centred feedback to explain the dark patterns would still need to be added to truly empower users. Therefore, it remains important to strengthen user-centred cookie banners in general, to learn what users need to interact with them confidently, and to make them immediately understandable. User-centred approaches are therefore urgently needed.

One possible reason for the fewer accepted cookies in the novel cookie banner group compared to the standard cookie banner group was design intent. Previous research has repeatedly shown that design has a strong influence on user behaviour [35, 61]. The difference between the novel cookie banner in this study and the one commonly used lies partly in the designer's intention: While existing cookie banners are sometimes designed to encourage users to allow as many cookies as possible, the novel cookie banner was not. Importantly, the novel cookie banner was designed to be as neutral as possible, so as not to be a nudge in the opposite direction of specifically accepting fewer cookies. Rather, its main aim was to provide users with neutral information, congruent with their privacy knowledge, to enable them to make an informed choice. This was confirmed by the fact that we did not find the same results in both the personalised and non-personalised groups of the novel cookie banner, but that we could isolate a specific effect for the personalised group. Importantly, the improvement of the novel cookie banner was not only reflected in fewer accepted cookies, but also in an increase in perceived usability. Taken together, in line with previous research, this experiment once again demonstrated the importance of design and designer intent.

6.2 The Benefits of Personalisation

While most user interventions are one-size-fits-all approaches, they are often not as effective as desired. The main reason for this is that users and contexts of use are highly diverse, leading to individual intrinsic privacy requirements. It is therefore crucial to take into account inter-individual differences and different groups of end-users. The present study demonstrates the potential of personalisation to enhance users' privacy by enabling them to better understand the consequences of privacy issues, to inform themselves more efficiently and to make easier decisions (at the first level of choice) when choosing cookies. Specifically, a personalised version of the novel cookie banner resulted in fewer cookies being accepted than the non-personalised version as well as a higher usability of the cookie banner. A sub-analysis revealed differences between the three experimental groups, especially in the group with high privacy knowledge. This suggests that the effect of personalisation may not be equally strong in all privacy knowledge groups. However, due to the uneven sample size, with particularly small numbers in the low privacy knowledge group, the power to detect such differences was not equal in all groups. Therefore, these results can only be considered exploratory and the detailed dynamics remain to be explored in future research.

This study generally joins other studies demonstrating the benefits of personalised support rather than one-size-fits-all solutions. Similar to studies in the area of mobile app permissions with privacy implications, the present study shows a positive effect of personalisation in the area of browser cookie choice support. For example, the present results are consistent with the results of a field study in

which personalised recommendations led to high acceptance of the recommended actions and thus to high congruence with desired privacy profiles using the PPA [33]. The present results are also consistent with a study showing the beneficial effect of personalised privacy notifications when users' decisions did not match their previously stated attitudes – thus personalisation enabled individuals to act more in line with their privacy preferences [25]. A similar positive potential for improving privacy-related outcomes, into which the present results fit, has been shown in the area of IoT devices [10, 11]. In the area of browser cookies, the approach of providing personalised support tailored to privacy knowledge has not been studied so far. However, a related experiment on the general improvement of users' cookie handling has been conducted by Kulyk et al. [31]. Both this and the present study present a concept for a privacy-friendly cookie setting interface that would assist users in configuring their cookie settings. The system by Kulyk et al. [31] determined the desired configuration for the user after asking a series of questions. The alternative cookie banner developed here instead provided personalised information directly to the user and encouraged them to make their own decision. In addition, Kulyk et al. approached the problem of cookie settings from the browser's point of view, while the present results were designed to replace the cookie banner that appears on every page regardless of the browser settings.

6.3 Theoretical Considerations

It has previously been suggested that personalised privacy enhancing technologies may shed light on explanations for the privacy paradox and help to overcome it. Research on the privacy paradox suggests that it is not necessarily a paradox, but a phenomenon that can be explained by several factors, including cognitive boundaries and a perceived imbalance of costs and benefits [15, 27, 28, 47, 54]. PPAs have already been proposed as specific means of overcoming the privacy paradox. For example, explicitly pointing out the paradox to individuals by showing the inconsistency of their actions with their attitudes led to a reduction in the paradox [25].

Our results also shed light on how the privacy paradox can be understood as a lack of competence or tools to actually act in accordance with one's attitudes. This is because our results show clear differences and less privacy disclosure behaviour (i.e. fewer accepted cookies) in the personalised cookie banner group compared to the standard cookie banner group. Thus, if users were generally provided with more comprehensible information about the privacy choices they need to make online, and if they were empowered through personalised support, they might actually be less likely to disclose their data – and their attitudes and behaviours might be more aligned.

The more personalised support for users becomes the norm in protecting their privacy, however, the more the personalisation-privacy paradox [48] has to be considered. In general, the personalisation of online services has the obvious risk of violating privacy, as it is based on the collection of personal data. Privacy and personalisation thus appear to be mutually exclusive. Improving one is often done at the expense of the other [48]. This issue is of particular interest when personalisation is used to enhance user privacy, as this does not magically solve the problem that personalisation needs to

be based on potentially sensitive information about individuals' attitudes and behaviours. The practical question here is how websites and consent management providers could choose which personalised banner to show to which user. One possibility is to explicitly use short questionnaires, integrated into the website's registration or login process to determine the appropriate cookie banner to show for new users. However, this would only be expedient if this classification could be used across several websites. Another possibility is to base the clustering of users on the collection of behavioural data on knowledge and preferences [11]. Machine learning algorithms that take into account various user attributes would, in theory, be able to make a prediction based on this data. However, this would have its own serious privacy implications if based on a large database and algorithmic prediction of individual privacy preferences and behaviours. This has the potential to undermine the whole approach of enabling informed privacy choices. These approaches would therefore need to be implemented by a trusted entity using privacy-preserving techniques. Finally, another possibility would be to allow the user to choose the level of cookie banner explanation. Regardless of the exact implementation, individuals' privacy knowledge may change over time, so the practical implementation would need to allow for re-categorisation. In general, as PPAs becomes more widespread in a variety of areas such as app permissions, IoT, app permissions, nudging and browser cookie management, as in this study, the relevance of the personalisation-privacy paradox and its implications should be carefully considered. Future work should therefore focus in particular on assessing the optimal way to provide such personalised support to the end-user, taking into account these privacy implications.

6.4 Limitations and Future Work

Some limitations of the current study need to be considered and potentially addressed in future studies. First, the coding process in the interviews was carried out by one researcher. Therefore, some subjective bias cannot be ruled out and we do not report on inter-rater reliability. However, the questions were quite specific and focused on eliciting explicit requirements, so the responses were considered to be quite unambiguous. Furthermore, no distraction task was used in the evaluation. Therefore, we cannot rule out the possibility that some social desirability bias may have influenced the results to some extent. However, this is unlikely to be a decisive factor, given the differentiated results in the three study groups. Furthermore, our personalised cookie banner did not address all the requirements raised in the interviews. In particular, the group of users with a high level of privacy knowledge preferred an automated solution. However, the chosen approach of varying the cookie explanations allowed us to draw clear conclusions about the impact of personalisation and to avoid potentially confounding effects of fundamentally different interface designs or technical solutions. Regarding the different cookie explanations, interviews were conducted with all three knowledge groups. However, while one participant from the low privacy knowledge group and one from the high privacy knowledge group made specific comments on the explanatory texts, none of the participants from the medium privacy knowledge group made specific comments during the creation of the texts.

Furthermore, the privacy knowledge groups were not evenly distributed within our study. In particular, the low privacy knowledge group was quite small. More recently, online prolific workers have been shown to be significantly more knowledgeable about privacy and security issues than the US population as a whole [57]. At the same time, as a potential alternative, MTurk's data quality has deteriorated in recent years and is also not easily generalisable [57]. Therefore, a trade-off would have to be made between high data quality in general and obtaining specific participants. However, the phenomenon of a skew towards high privacy knowledge based on the OPLIS is not unique to our Prolific-based study. For example, in the study by Sindermann et al. [52], the mean OPLIS score in their sample would have fallen into the high privacy knowledge group, just as in the work by Ortloff et al. [44]. Our results also showed an association between higher education levels and higher OPLIS scores. In anticipation of this, we took particular care to ensure that we had a diverse sample outside the university context and that all levels of education were sufficiently represented. To some extent, the self-selection on a crowdworking platform such as Prolific goes hand in hand with a certain IT affinity - which may ultimately also go hand in hand with higher privacy knowledge. In general, it seems that a basic knowledge of privacy, and at least some affinity with and regular use of the internet is enough to get a higher score in the OPLIS questionnaire. Moreover, a user group that almost never uses the internet may not be the target group for such an everyday support tool. However, it is all the more important to recognise that there may be a vulnerable population group that requires a different approach than the less vulnerable.

Future work should consider these limitations and, among other things, could include extending these results to provide automated translations of cookie banners into a form that is understandable to end-user groups. Furthermore, while the medium and low privacy knowledge groups in particular had very similar requirements, there may be potential for future research to explore and implement broader dimensions of personalisation. Based on an initial assessment of privacy knowledge, a privacy assistant could, among other things, either automate desired cookie choices and/or inform users about the implications of the choices made according to their privacy knowledge.

7 CONCLUSION

This study employed a user-centred approach to develop an alternative cookie banner with explanations tailored to individuals with low, medium and high privacy knowledge. The results suggest that the novel cookie banner in general, and the personalised version in particular, resulted in fewer accepted cookies and increased usability. The results of this study therefore have significant implications for the design of cookie banners and the communication of privacy information to users. The tailored approach used in the design of the novel cookie banner appears to be an effective strategy to empower users to interact with privacy information in a sovereign manner and, as a result, to act more in line with their actual attitudes. Overall, this study provides important insights for designers and developers who want to create effective and user-friendly privacy tools for website users.

ACKNOWLEDGMENTS

This research work has been funded by the German Federal Ministry of Education and Research and the Hessian Ministry of Higher Education, Research, Science and the Arts within their joint support of the National Research Center for Applied Cybersecurity ATHENE. Furthermore, it has been funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – SFB 1119 (CROSSING) – 236615297

REFERENCES

- [1] Troy L Adams, Yuanxia Li, and Hao Liu. 2020. A Replication of Beyond the Turk: Alternative Platforms for Crowdsourcing Behavioral Research – Sometimes Preferable to Student Groups. *AIS Transactions on Replication Research* 6 (10 2020), 15. Issue 1. <https://doi.org/10.17705/1atrr.00058>
- [2] Natã M. Barbosa, Joon S. Park, Yaxing Yao, and Yang Wang. 2019. “What if?” Predicting Individual Users’ Smart Home Privacy Preferences and Their Changes. *Proceedings on Privacy Enhancing Technologies* 2019, 4 (oct 2019), 211–231. <https://doi.org/10.2478/popets-2019-0066>
- [3] Yoav Benjamini and Yoel Hochberg. 1995. Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing. *Journal of the Royal Statistical Society. Series B (Methodological)* 57, 1 (1995), 289–300. <http://www.jstor.org/stable/2346101>
- [4] Benjamin Maximilian Berens, Heike Dietmann, Chiara Krisam, Oksana Kulyk, and Melanie Volkamer. 2022. Cookie Disclaimers: Impact of Design and Users’ Attitude. In *Proceedings of the 17th International Conference on Availability, Reliability and Security (Vienna, Austria) (ARES ’22)*. Association for Computing Machinery, New York, NY, USA, Article 12, 20 pages. <https://doi.org/10.1145/3538969.3539008>
- [5] Tom Biselli, Enno Steinbrink, Franziska Herbert, Gina M Schmidbauer-Wolf, and Christian Reuter. 2022. On the Challenges of Developing a Concise Questionnaire to Identify Privacy Personas. *Proceedings on Privacy Enhancing Technologies* 4 (2022), 645–669.
- [6] Christoph Bösch, Benjamin Erb, Frank Kargl, Henning Kopp, and Stefan Pfattheicher. 2016. Tales from the Dark Side: Privacy Dark Strategies and Privacy Dark Patterns. *Proc. Priv. Enhancing Technol.* 2016, 4 (2016), 237–254.
- [7] Harry Brignull, Marc Miquel, Jeremy Rosenberg, and James Offer. 2010. *Types of Deceptive Design*. Deceptive Design. <https://www.deceptive.design/types>
- [8] John Brooke et al. 1996. SUS-A quick and dirty usability scale. *Usability evaluation in industry* 189, 194 (1996), 4–7.
- [9] Aaron Cahn, Scott Alfeld, Paul Barford, and Shanmugavelayutham Muthukrishnan. 2016. An Empirical Study of Web Cookies. In *Proceedings of the 25th International Conference on World Wide Web (Montréal, Québec, Canada) (WWW ’16)*. International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, CHE, 891–901. <https://doi.org/10.1145/2872427.2882991>
- [10] Jessica Colnago, Yuanyuan Feng, Tharangini Palanivel, Sarah Pearman, Megan Ung, Alessandro Acquisti, Lorrie Faith Cranor, and Norman Sadeh. 2020. Informing the Design of a Personalized Privacy Assistant for the Internet of Things. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (Honolulu, HI, USA) (CHI ’20)*. Association for Computing Machinery, New York, NY, USA, 1–13. <https://doi.org/10.1145/3313831.3376389>
- [11] Anupam Das, Martin Degeling, Daniel Smullen, and Norman Sadeh. 2018. Personalized privacy assistants for the internet of things: Providing users with notice and choice. *IEEE Pervasive Computing* 17, 3 (2018), 35–46.
- [12] 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 *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (Florence, Italy) (CHI ’08)*. Association for Computing Machinery, New York, NY, USA, 1065–1074. <https://doi.org/10.1145/1357054.1357219>
- [13] Serge Egelman and Eyal Peer. 2015. The Myth of the Average User: Improving Privacy and Security Systems through Individualization. In *Proceedings of the 2015 New Security Paradigms Workshop (Twente, Netherlands) (NSPW ’15)*. Association for Computing Machinery, New York, NY, USA, 16–28. <https://doi.org/10.1145/2841113.2841115>
- [14] Tom Funk. 2008. Successful Online Business Models for Web 2.0 and Beyond. In *Web 2.0 and Beyond: Understanding the New Online Business Models, Trends, and Technologies: Understanding the New Online Business Models, Trends, and Technologies*. Praeger Publishers, Westport, Connecticut/London, Chapter 5, 79–92.
- [15] Nina Gerber, Paul Gerber, and Melanie Volkamer. 2018. Explaining the privacy paradox: A systematic review of literature investigating privacy attitude and behavior. *Computers & security* 77 (2018), 226–261.
- [16] Nina Gerber, Benjamin Reinheimer, and Melanie Volkamer. 2019. Investigating People’s Privacy Risk Perception. *Proceedings on Privacy Enhancing Technologies* 2019, 3 (jul 2019), 267–288. <https://doi.org/10.2478/POPETS-2019-0047>
- [17] Colin M. Gray, Yubo Kou, Bryan Battles, Joseph Hoggatt, and Austin L. Toombs. 2018. The Dark (Patterns) Side of UX Design. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (Montreal QC, Canada) (CHI ’18)*. Association for Computing Machinery, New York, NY, USA, 1–14. <https://doi.org/10.1145/3173574.3174108>
- [18] Paul Grafl, Hanna Schraffenberger, Frederik Zuiderveen Borgesius, and Moniek Buijzen. 2021. Dark and Bright Patterns in Cookie Consent Requests. *Journal of Digital Social Research* 3, 1 (2021), 1–38. <https://doi.org/10.33621/jdsr.v3i1.54>
- [19] Elias Grünwald and Frank Pallas. 2021. TILT: A GDPR-Aligned Transparency Information Language and Toolkit for Practical Privacy Engineering. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (Virtual Event, Canada) (FAccT ’21)*. Association for Computing Machinery, New York, NY, USA, 636–646. <https://doi.org/10.1145/3442188.3445925>
- [20] Tristan Harris. 2016. *How Technology is Hijacking Your Mind – from a Magician and Google Design Ethicist*. Medium Magazine. <https://medium.com/thrive-global/how-technology-hijacks-peoples-minds-from-a-magician-and-google-s-design-ethicist-56d62ef5edf3>
- [21] Katrin Hartwig and Christian Reuter. 2021. Nudging Users Towards Better Security Decisions in Password Creation Using Whitebox-based Multi-dimensional Visualizations. *Behaviour & Information Technology (BIT)* (2021). <https://doi.org/10.1080/0144929X.2021.1876167>
- [22] Woodrow Hartzog. 2018. The case against idealising control. *Eur. Data Prot. L. Rev.* 4 (2018), 423.
- [23] Philip Hausner and Michael Gertz. 2021. Dark Patterns in the Interaction with Cookie Banners. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems (Position Paper at the Workshop “What Can CHI Do About Dark Patterns?”)*. ACM, Yokohama, Japan, 5 pages. <https://doi.org/10.48550/ARXIV.2103.14956>
- [24] Jaap-Henk Hoepman. 2014. Privacy Design Strategies. In *ICT Systems Security and Privacy Protection (IFIP Advances in Information and Communication Technology book series (IFIPAICT, volume 428))*, Nora Cuppens-Boulahia, Frédéric Cuppens, Sushil Jajodia, Anas Abou El Kalam, and Thierry Sans (Eds.). Springer, Berlin, Heidelberg, 446–459.
- [25] Corey Brian Jackson and Yang Wang. 2018. Addressing the privacy paradox through personalized privacy notifications. *Proceedings of the ACM on interactive, mobile, wearable and ubiquitous technologies* 2, 2 (2018), 1–25.
- [26] Georgios Kampanos and Siamak F. Shahandashti. 2021. Accept All: The Landscape of Cookie Banners in Greece and the UK. In *ICT Systems Security and Privacy Protection (IFIP Advances in Information and Communication Technology book series (IFIPAICT, volume 625))*, Audun Jøsang, Lynn Fletcher, and Janne Hagen (Eds.). Springer International Publishing, Cham, 213–227.
- [27] Bart P. Knijnenburg. 2017. Privacy? I Can’t Even! Making a Case for User-Tailored Privacy. *IEEE Security Privacy* 15, 4 (2017), 62–67. <https://doi.org/10.1109/MSP.2017.3151331>
- [28] Spyros Kokolakis. 2017. Privacy attitudes and privacy behaviour: A review of current research on the privacy paradox phenomenon. *Computers & security* 64 (2017), 122–134.
- [29] Chiara Krisam, Heike Dietmann, Melanie Volkamer, and Oksana Kulyk. 2021. Dark Patterns in the Wild: Review of Cookie Disclaimer Designs on Top 500 German Websites. In *Proceedings of the 2021 European Symposium on Usable Security (Karlsruhe, Germany) (EuroUSEC ’21)*. Association for Computing Machinery, New York, NY, USA, 1–8. <https://doi.org/10.1145/3481357.3481516>
- [30] David M. Kristol. 2001. HTTP Cookies: Standards, privacy, and politics. *ACM Transactions on Internet Technology (TOIT)* 1, 2 (2001), 151–198.
- [31] Oksana Kulyk, Peter Mayer, Melanie Volkamer, and Oliver Käfer. 2018. A Concept and Evaluation of Usable and Fine-Grained Privacy-Friendly Cookie Settings Interface. In *2018 17th IEEE International Conference on Trust, Security and Privacy in Computing and Communications/12th IEEE International Conference on Big Data Science and Engineering (TrustCom/BigDataSE)*. IEEE, New York, NY, USA, 1058–1063. <https://doi.org/10.1109/TrustCom/BigDataSE.2018.00148>
- [32] Sebastian Linsner, Franz Kuntke, Enno Steinbrink, Jonas Franken, and Christian Reuter. 2021. The Role of Privacy in Digitalization – Analysing the German Farmers’ Perspective. *Proceedings on Privacy Enhancing Technologies (PoPETS)* 2021, 3 (2021), 334–350. <https://doi.org/10.2478/popets-2021-0050>
- [33] Bin Liu, Mads Schaarup Andersen, Florian Schaub, Hazim Almuhammedi, Shikun Zhang, Norman Sadeh, Alessandro Acquisti, and Yuvraj Agarwal. 2016. Follow My Recommendations: A Personalized Privacy Assistant for Mobile App Permissions. In *Proceedings of the Twelfth USENIX Conference on Usable Privacy and Security (Denver, CO, USA) (SOUPS ’16)*. USENIX Association, USA, 27–41.
- [34] Bin Liu, Jialiu Lin, and Norman Sadeh. 2014. Reconciling Mobile App Privacy and Usability on Smartphones: Could User Privacy Profiles Help?. In *Proceedings of the 23rd International Conference on World Wide Web (Seoul, Korea) (WWW ’14)*. Association for Computing Machinery, New York, NY, USA, 201–212. <https://doi.org/10.1145/2566486.2568035>
- [35] Dominique Machuletz and Rainer Böhme. 2020. Multiple Purposes, Multiple Problems: A User Study of Consent Dialogs after GDPR. In *Proceedings on Privacy Enhancing Technologies Symposium*. Scienco, Montreal, Canada, 481–498. <https://doi.org/10.2478/popets-2020-0037>

- [36] Philipp K. Masur, Doris Teutsch, and Sabine Trepte. 2017. Entwicklung und Validierung der Online-Privatheitskompetenzskala (OPLIS). *Diagnostica* 63, 4 (2017), 256–268. <https://doi.org/10.1026/0012-1924/a000179>
- [37] Arunesh Mathur, Gunes Acar, Michael J Friedman, Elena Lucherini, Jonathan Mayer, Marshini Chetty, and Arvind Narayanan. 2019. Dark patterns at scale: Findings from a crawl of 11K shopping websites. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019), 1–32.
- [38] Yannic Meier, Johanna Schäwel, and Nicole C Krämer. 2020. The shorter the better? Effects of privacy policy length on online privacy decision-making. *Media and Communication* 8, 2 (2020), 291–301.
- [39] William R. Miller and Stephen Rollnick. 2012. *Motivational interviewing: Helping people change* (3 ed.). The Guilford Press, New York NY.
- [40] Karsten Neß. 2021. *Wie Cookies genutzt werden*. Privacy-Handbuch. https://www.privacy-handbuch.de/handbuch_21b.htm
- [41] Jakob Nielsen. 2020. *10 Usability Heuristics for User Interface Design*. Nielsen Norman Group. <http://www.nngroup.com/articles/ten-usability-heuristics/>
- [42] Faryal Nosheen and Usman Qamar. 2015. Flexibility and privacy control by cookie management. In *2015 Third International Conference on Digital Information, Networking, and Wireless Communications (DINWC)*. IEEE, Moscow, Russia, 94–98. <https://doi.org/10.1109/DINWC.2015.7054224>
- [43] Midas Nouwens, Ilaria Liccardi, Michael Veale, David Karger, and Lalana Kagal. 2020. Dark Patterns after the GDPR: Scraping Consent Pop-Ups and Demonstrating Their Influence. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '20). Association for Computing Machinery, New York, NY, USA, 1–13. <https://doi.org/10.1145/3313831.3376321>
- [44] Anna-Marie Ortloff, Steven Zimmerman, David Elswiler, and Niels Henze. 2021. The Effect of Nudges and Boosts on Browsing Privacy in a Naturalistic Environment. In *Proceedings of the 2021 Conference on Human Information Interaction and Retrieval* (Canberra ACT, Australia) (CHIIR '21). Association for Computing Machinery, New York, NY, USA, 63–73. <https://doi.org/10.1145/3406522.3446014>
- [45] Eyal Peer, Laura Brandimarte, Sonam Samat, and Alessandro Acquisti. 2017. Beyond the Turk: Alternative platforms for crowdsourcing behavioral research. *Journal of Experimental Social Psychology* 70 (2017), 153–163. <https://doi.org/10.1016/j.jesp.2017.01.006>
- [46] Eyal Peer, Serge Egelman, Marian Harbach, Nathan Malkin, Arunesh Mathur, and Alisa Frik. 2020. Nudge me right: Personalizing online security nudges to people’s decision-making styles. *Computers in Human Behavior* 109 (2020), 106347. <https://doi.org/10.1016/j.chb.2020.106347>
- [47] Stefanie Pöttsch. 2009. Privacy Awareness: A Means to Solve the Privacy Paradox?. In *The Future of Identity in the Information Society (IFIP Advances in Information and Communication Technology book series (IFIPACT, volume 298))*, Vashek Matyáš, Simone Fischer-Hübner, Daniel Cvrček, and Petr Švenda (Eds.). Springer, Berlin, Heidelberg, 226–236.
- [48] Shubhadip Ray, Tharangini Palanivel, Norbert Herman, and Yixuan Li. 2021. Dynamics in Data Privacy and Sharing Economics. *IEEE Transactions on Technology and Society* 2, 3 (2021), 114–115. <https://doi.org/10.1109/TTS.2021.3077534>
- [49] Benjamin Maximilian Reinheimer, Kristoffer Braun, and Melanie Volkamer. 2016. Entwicklung eines Interfaces zur privacy-friendly Cookie-Einstellung. In *Mensch und Computer 2016 – MuC-Workshopband “Usable Security and Privacy”*, Benjamin Weyers and Anke Dittmar (Eds.). Gesellschaft für Informatik e.V., Aachen, 8 pages. <https://doi.org/10.18420/muc2016-ws03-0001>
- [50] Antonio Ruiz-Martínez. 2012. A survey on solutions and main free tools for privacy enhancing Web communications. *Journal of network and computer applications* 35, 5 (2012), 1473–1492.
- [51] Umesh Shankar and Chris Karlof. 2006. Doppelgänger: Better Browser Privacy without the Bother. In *Proceedings of the 13th ACM Conference on Computer and Communications Security* (Alexandria, Virginia, USA) (CCS '06). Association for Computing Machinery, New York, NY, USA, 154–167. <https://doi.org/10.1145/1180405.1180426>
- [52] Cornelia Sindermann, Helena Sophia Schmitt, Frank Kargl, Cornelia Herbert, and Christian Montag. 2021. Online Privacy Literacy and Online Privacy Behavior – The Role of Crystallized Intelligence and Personality. *International Journal of Human-Computer Interaction* 37, 15 (2021), 1455–1466. <https://doi.org/10.1080/10447318.2021.1894799> arXiv:<https://doi.org/10.1080/10447318.2021.1894799>
- [53] Than Htut Soe, Cristiana Teixeira Santos, and Marija Slavkovik. 2022. Automated detection of dark patterns in cookie banners: how to do it poorly and why it is hard to do it any other way. <https://doi.org/10.48550/ARXIV.2204.11836>
- [54] Daniel J. Solove. 2021. The Myth of the Privacy Paradox. *George Washington Law Review* 89, 1 (2021), 51 pages. <https://doi.org/10.2139/ssrn.3536265>
- [55] Enno Steinbrink, Lilian Reichert, Michelle Mende, and Christian Reuter. 2021. Digital Privacy Perceptions of Asylum Seekers in Germany - An Empirical Study about Smartphone Usage during the Flight. *Proceedings of the ACM: Human Computer Interaction (PACM): Computer-Supported Cooperative Work and Social Computing* 5, CSCW2 (2021). <https://doi.org/10.1145/3479526>
- [56] Alina Stöver, Sara Hahn, Felix Kretschmer, and Nina Gerber. 2023. Investigating how Users Imagine their Personal Privacy Assistant. *Proceedings on Privacy Enhancing Technologies* 2023 (2023), 384–402.
- [57] Jenny Tang, Eleanor Birrell, and Ada Lerner. 2022. Replication: How Well Do My Results Generalize Now? The External Validity of Online Privacy and Security Surveys. In *Proceedings of the Eighteenth USENIX Conference on Usable Privacy and Security* (Boston, MA, USA) (SOUPS'22). USENIX Association, USA, 1326–1343.
- [58] Arnout Terpstra, Alexander P. Schouten, Alwin de Rooij, and Ronald E. Leenes. 2019. Improving privacy choice through design: How designing for reflection could support privacy self-management. *First Monday* 24, 7 (2019). <https://doi.org/10.5210/fm.v24i7.9358>
- [59] Sabine Trepte, Doris Teutsch, Philipp K Masur, Carolin Eicher, Mona Fischer, Alisa Hennhöfer, and Fabienne Lind. 2015. Do people know about privacy and data protection strategies? Towards the “Online Privacy Literacy Scale”(OPLIS). *Reforming European data protection law* (2015), 333–365.
- [60] Martino Trevisan, Stefano Traverso, Eleonora Bassi, and Marco Mellia. 2019. 4 Years of EU Cookie Law: Results and Lessons Learned. *Proc. Priv. Enhancing Technol.* 2019, 2 (2019), 126–145.
- [61] Christine Utz, Martin Degeling, Sascha Fahl, Florian Schaub, and Thorsten Holz. 2019. (Un)Informed Consent: Studying GDPR Consent Notices in the Field. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security* (London, United Kingdom) (CCS '19). Association for Computing Machinery, New York, NY, USA, 973–990. <https://doi.org/10.1145/3319535.3354212>
- [62] Logan Warber, Alessandro Acquisti, and Douglas Sicker. 2019. Can Privacy Nudges Be Tailored to Individuals’ Decision Making and Personality Traits?. In *Proceedings of the 18th ACM Workshop on Privacy in the Electronic Society* (London, United Kingdom) (WPES'19). Association for Computing Machinery, New York, NY, USA, 175–197. <https://doi.org/10.1145/3338498.3358656>
- [63] Daniel Wessel, Christiane Attig, and Thomas Franke. 2019. ATI-S - An Ultra-Short Scale for Assessing Affinity for Technology Interaction in User Studies. In *Proceedings of Mensch Und Computer 2019* (Hamburg, Germany) (MuC'19). Association for Computing Machinery, New York, NY, USA, 147–154. <https://doi.org/10.1145/3340764.3340766>
- [64] Chuan Yue, Mengjun Xie, and Haining Wang. 2007. Automatic Cookie Usage Setting with Cookiepicker. In *37th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN'07)*. IEEE, Edinburgh, UK, 460–470. <https://doi.org/10.1109/DSN.2007.21>
- [65] Chuan Yue, Mengjun Xie, and Haining Wang. 2010. An automatic HTTP cookie management system. *Computer Networks* 54, 13 (2010), 2182–2198.

A APPENDIX

A.1 Personalised Cookie Explanations

The cookie texts were written according to the following rules. After that they were revised with another privacy researcher and two participants from the pre-study, H1 and L12, who read them and gave feedback.

Rules for Creating Cookie Explanations:

- Texts resembled real world examples as this was what users expected.
- Texts for all clusters differed only in explanations and examples related to privacy knowledge. This way there were no additional variables influencing the study.
- According to pre-study participants, texts had to be short and concise, not a wall of text but providing plenty of information.
- The used terms were explained in the beginning to encourage comprehension [49].
- The options were supported by illustrative examples, to support users with less privacy knowledge [49].
- The text were phrased in a neutral tone, not nudging the users towards either choice.
- The texts included pros and cons, which allowed users to form their own opinion and prevented reactance [39].
- Users focused more on information if said information prompted them to think for themselves instead of simply providing courses of action [12]. This has prompted the line “you can now decide which cookies the website may use”.

In the following table, the cookie texts are listed for high, medium and low privacy knowledge, respectively. While the texts in the study were presented in German, they are translated into English below.

Category	Low Cluster	Medium Cluster	High Cluster
Main header	This website wants to use cookies. Cookies are text files which the website stores on your device and reads on demand, e.g. the next visit to the website. These can have different functions and may include third-party cookies. You can now decide which cookies the website may use.	This website wants to use cookies. These can have different functions and may include third-party cookies. You can now decide which cookies the website may use.	This website wants to use cookies. These can have different functions and may include third-party cookies. You can now decide which cookies the website may use.
Explanation of cookies	Cookies can be thought of as a slip of paper on which the website writes something and then stores it with you. When revisiting the website, it will ask for that piece of paper and use the information it has written on it. Cookies usages include, but are not limited to, storing logins (on pages with user accounts such as Facebook), to store the contents of the shopping basket at an online shop. But also to collect general information about the user e.g. to improve the efficiency of advertisements.	Cookies are text files stored on your device by a website and read as needed e.g. on the next visit to the website. Cookies usages include but are not limited to identify users, to store logins or the contents of the shopping basket. But also to collect general information about the user e.g. to improve the efficiency of advertisements.	Cookies are text files stored on your device by a website and read as needed e.g. on the next visit to the website.
Explanation of third-party cookies	These cookies are managed by third parties who work with various websites. These parties are e.g. service providers or agencies. The agency sets the cookie on a website and if you then visit another site that also works with the agency, it can read the cookie and continue to use it. This allows the agency to collect and combine your information from all participating sites to create the most accurate picture of your preferences. They may also use information from other sources to do so.	These cookies are managed by third parties who work with various websites. The third parties can set and read their cookie on several pages. This allows these parties to collect and combine your information from all participating sites to create the most accurate picture of your preferences. They may also use information from other sources to do so.	These cookies are managed by third parties who work with various websites. This allows these parties to collect and combine your information from all participating sites to create the most accurate picture of your preferences.
Information on functional cookies	Functional cookies are used to enable the functions of the website. The website cannot function properly without these cookies. These cookies only relate to the website visited and are not passed on to third parties. Therefore, according to the law, the website can use them without requiring your consent. For all other types of cookies, your permission is required. Example: For example, when shopping online, cookies are often used to store the products in your shopping basket. This allows the website to remember and display your basket, even if you have left the page.	Functional cookies are used to enable the functions of the website. The website cannot function properly without these cookies. According to the law, the website can use them without requiring your consent. For all other types of cookies, your permission is required. Example: Functional cookies enable basic functions such as page navigation and access to secure areas of the website. These are not third-party cookies.	Functional cookies are used to enable the functions of the website. By law, no consent is required for their use.

Table 4 continued from previous page

<p>Information on marketing cookies</p>	<p>Marketing cookies are used to analyse interactions with websites. This information is then used to target advertisements specifically to the user in the hope that they will be more influenced by them and therefore buy more. These personalised ads are therefore more valuable to advertisers. These cookies are often third party cookies, i.e. they pass information to partners of the website. These partners may combine this information with other information that you have provided to them or that they have collected as part of your use of the services to build a more accurate picture of your preferences. You will still see ads without marketing cookies. However, their content will not be tailored to you. Example: For example, you have a pet and are looking for vets or toys online. You also visit the website of a major pet food company. The ad cookie remembers all this and shows you ads for pet supplies.</p>	<p>Marketing cookies are used to analyse interactions with websites. This information is then used to tailor advertisements to you. Ads that are personalised in this way are more likely to be clicked on, making them more valuable to advertisers. These cookies are often third party cookies. Third parties may combine information from different sources to build a more accurate picture of your preferences. You will still see ads without marketing cookies. However, their content will not be tailored to you.</p>	<p>Marketing cookies are used to show targeted, appealing advertisements. This includes third-party cookies. You will still see ads without marketing cookies. However, their content will not be tailored to you.</p>
<p>Information on all cookies (This information was presented to all clusters with the same wording. The only difference was in the examples provided to the low and medium clusters. No additional examples were provided to the high cluster, as it was assumed that users in this cluster were familiar with the use of the different cookie types.): This option includes all cookies that the website wishes to set. This means that both functional and marketing cookies are included. It may also include cookies in the categories of preference cookies, statistical cookies, social media cookies and others. This includes third-party cookies.</p>	<p>Example for low cluster:</p> <ul style="list-style-type: none"> • Preference cookies: Remember your settings on the site, e.g. if you have changed the font size or the website offers different positions for a menu and you have decided on one of them. • Statistics cookies: Analyse how you interact with the website in order to gradually make the website more efficient. For example, if the website notices that many users click on a button, then immediately go back and select another one, it can be concluded that these buttons need to be optimised. • Social media cookies: Display content from social media (e.g. Facebook, Instagram or Twitter) and share information with social media. For example, tweets from politicians or well-known personalities are often displayed on news sites to match the content of an article. If you interact with it, these interactions are forwarded to the social media providers. <p>All cookie types can also be third-party cookies.</p>	<p>Example for medium cluster:</p> <ul style="list-style-type: none"> • Preference cookies: Remember your settings on the site. • Statistics cookies: Analyse how you interact with the website in order to gradually make the website more efficient. • Social media cookies: Display content from social media (e.g. Facebook, Instagram or Twitter) and share information with social media. <p>All cookie types can also be third-party cookies.</p>	<p>No additional examples were provided to the high cluster.</p>

Table 4: Personalised Information Texts – English Version

	Privacy Knowledge		
	Low	Medium	High
Percentile rank	0-33	34-66	67-100
Raw value	0-9	10-13	14-20

Table 5: Privacy knowledge groups created based on OPLIS scores

	Website	Type
1	google.de	research
2	youtube.de	entertainment
3	facebook.de	entertainment
4	amazon.de	commerce
removed	wikipedia.org	research
removed	bild.de	news
removed	spiegel.de	news
5	gmx.de	mail
6	web.de	mail
7	ebay.de	commerce

Table 6: List of popular websites selected for the study, including their type

	Interview Study		Refinement Study		Evaluation Study	
	N	%	N	%	N	%
Age Group						
18-24	1	5	3	8	24	15
25-34	8	42	22	56	39	25
35-44	2	11	4	10	54	35
45-54	0	0	3	8	19	12
55-64	5	26	2	5	13	8
>64	3	16	5	13	8	5
Gender						
female	12	63	20	51	61	39
male	7	37	18	46	93	59
diverse	0	0	1	3	3	2
Education						
lower secondary education	2	11	3	8	27	17
middle or high school	5	26	6	15	73	47
academic degree	12	63	30	77	57	36
Total N	19		39		157	
IT Affinity/Proficiency	Mean	SD	Mean	SD	Mean	SD
Self-assessed on 5 point scale	3.73	1.03	3.79	0.89		
ATI-S on 6 point scale					4.10	1.00

Table 7: Participant information of the interview, refinement and evaluation study

	Privacy Knowledge		
	low	medium	high
Control Group (standard cookie banner)	2	10	41
Privacy Assistant	4	16	29
Personalised Privacy Assistant	3	9	43

Table 8: Distribution of privacy knowledge in the different experimental groups

Cookie Banner Design	Cookie Texts
Overall neutral or positive evaluation (N=15), e.g.: "the design was appealing" "clear" "kept simple, good and understandable"	Overall neutral or positive evaluation (N=13), e.g.: "good comprehensiveness" "good", "explanations were very good and understandable"
Non-appealing design (N=2), e.g.: "a bit old-fashioned" "there are more appealing designs"	Only skimmed over (N=5), e.g.: "only skimmed the text because I always choose functional cookies" "did not read"
Too big / too much text (N=4), e.g.: "a little too big, but easy to use" "a little too full, but otherwise appropriate"	Too much text (N=3), e.g.: "good, but too much visible at once" "difficult to read"

Table 9: Themes of feedback in the refinement of the study protocol (N=39)

Experimental Group	Privacy Knowledge		
	Low	Med	High
PPA	63	82	85
PA	69	72	74
CG	64	66	62

Table 10: SUS scores in the different experimental groups. PPA = Personalised Privacy Assistant, PA = Privacy Assistant, CG = Control Group

Variable	Estimate	Std. Error	z value	Pr(> z)
Age	0.01	0.01	0.82	0.41
Gender: female	-0.46	1.23	-0.37	0.71
Gender: male	0.15	1.21	0.13	0.90
Education: academic degree	1.46	0.45	3.24	< 0.01
Education: middle or high school	0.82	0.42	1.95	0.051
ATI-S	0.56	0.15	3.69	< 0.001

Table 11: Regression model for the OPLIS scores as outcome. The categorical variables are dummy coded, with gender (diverse) and education (lower secondary education) as reference.

Categories	Example
Automation	set cookies once and apply on every website
Temporary cookie consent	set expiry date for certain cookies
Compatibility	synchronize settings across browsers
Reliability	consideration of credibility of provider / developer
Well-structured information	clear & concise information, better structured on first level
Simplicity and usability	simple one-click buttons for main choices on first level
Simple language	less use of technical jargon

Table 12: Coding categories for feature requests in the pre-study

Comparison	Z	P.unadj	P.adj
CG.high - CG.low	-0.18668657	8.519064e-01	8.762466e-01
CG.high - CG.medium	0.42245025	6.726964e-01	8.072357e-01
CG.low - CG.medium	0.36687826	7.137098e-01	8.288243e-01
CG.high - PA.high	4.07232113	4.654694e-05	8.378450e-04
CG.low - PA.high	1.53646915	1.244234e-01	3.732701e-01
CG.medium - PA.high	2.28814005	2.212937e-02	1.327762e-01
CG.high - PA.low	0.68908176	4.907718e-01	7.681646e-01
CG.low - PA.low	0.57289933	5.667129e-01	7.035056e-01
CG.medium - PA.low	0.35828235	7.201320e-01	8.101485e-01
PA.high - PA.low	-1.17581117	2.396703e-01	5.392583e-01
CG.high - PA.medium	2.49934407	1.244234e-02	8.958488e-02
CG.low - PA.medium	1.16256550	2.450058e-01	5.188359e-01
CG.medium - PA.medium	1.45800709	1.448386e-01	4.010915e-01
PA.high - PA.medium	-0.80714897	4.195807e-01	7.192811e-01
PA.low - PA.medium	0.67221346	5.014478e-01	7.220849e-01
CG.high - PPA.high	6.35192512	2.126368e-10	7.654926e-09
CG.low - PPA.high	2.10362274	3.541137e-02	1.593512e-01
CG.medium - PPA.high	3.52486487	4.236989e-04	5.084386e-03
PA.high - PPA.high	1.65800076	9.731731e-02	3.184930e-01
PA.low - PPA.high	1.96185962	4.977883e-02	1.991153e-01
PA.medium - PPA.high	2.21882000	2.649897e-02	1.362804e-01
CG.high - PPA.low	0.99782146	3.183659e-01	5.730587e-01
CG.low - PPA.low	0.80185014	4.226397e-01	6.915922e-01
CG.medium - PPA.low	0.68026160	4.963388e-01	7.445083e-01
PA.high - PPA.low	-0.64520219	5.187962e-01	7.183332e-01
PA.low - PPA.low	0.30878745	7.574832e-01	8.263453e-01
PA.medium - PPA.low	-0.22242359	8.239842e-01	8.724538e-01
PPA.high - PPA.low	-1.32244600	1.860197e-01	4.464472e-01
CG.high - PPA.medium	2.67047120	7.574487e-03	6.817038e-02
CG.low - PPA.medium	1.43040974	1.525995e-01	3.923986e-01
CG.medium - PPA.medium	1.81518522	6.949545e-02	2.501836e-01
PA.high - PPA.medium	-0.01332443	9.893690e-01	9.893690e-01
PA.low - PPA.medium	1.03516601	3.005914e-01	5.695416e-01
PA.medium - PPA.medium	0.59106865	5.544744e-01	7.392992e-01
PPA.high - PPA.medium	-1.10072480	2.710165e-01	5.420329e-01
PPA.low - PPA.medium	0.57932557	5.623695e-01	7.230465e-01

Table 13: Exploration of interaction effects regarding the number of accepted cookies based on a Kruskal-Wallis test and subsequent multiple comparisons using Dunn's test with Benjamini-Hochberg correction for multiple comparisons