The Effect of Nudges and Boosts on Browsing Privacy in a Naturalistic Environment

Anna-Marie Ortloff
anna-marie.ortloff@student.ur.de
Universität Regensburg

David Elweiler
david.elweiler@ur.de
Universität Regensburg

Steven Zimmerman
szimme@essex.ac.uk
University of Essex

Niels Henze
niels.henze@ur.de
Universität Regensburg

ABSTRACT

During everyday web browsing and search users reveal many pieces of private information to third parties. Even though people report being concerned about their privacy online, they often do not take steps to protect it. This is known as the ‘privacy paradox’ in the literature. In this work we study two well-known strategies based on theories from the behavioral sciences, nudging and boosting, which encourage users to browse in a way that their private data are less exposed. First, an online survey (N=127) tested the comprehensibility and efficacy of various facts (boosts), before the most effective of these were evaluated against ‘nudge’ interventions previously shown to be efficacious in lab-studies. A three week naturalistic study (N=68) using a browser extension revealed that both nudges and boosts improve browsing privacy, as approximated by different measures. Boosts are also shown to improve user knowledge about privacy in the short term, but the benefit weakens over time.

CCS CONCEPTS

• **Security and privacy** → **Human and societal aspects of security and privacy**; • **Human-centered computing** → *Field studies*; *Empirical studies in HCI*.

KEYWORDS

Nudges, Boosts, Privacy, Web Browsing, Human Information Behavior

ACM Reference Format:

Anna-Marie Ortloff, Steven Zimmerman, David Elweiler, and Niels Henze. 2021. The Effect of Nudges and Boosts on Browsing Privacy in a Naturalistic Environment. In *Proceedings of the 2021 ACM SIGIR Conference on Human Information Interaction and Retrieval (CHIIR '21), March 14–19, 2021, Canberra, ACT, Australia*. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3406522.3446014>

This work was supported by the Economic and Social Research Council [grant number ES/M010236/1].

This work is licensed under a Creative Commons Attribution International 4.0 License.

CHIIR '21, March 14–19, 2021, Canberra, Australia.

© 2021 Copyright is held by the owner/author(s).

ACM ISBN 978-1-4503-8055-3/21/03. <https://doi.org/10.1145/3406522.3446014>

1 INTRODUCTION

World wide, users spend almost 7 hours per day online [44], e.g. searching for information or browsing the web. This can expose private data, which can be harvested and used for many different purposes [31], leading to, e.g. search or price discrimination [59], targeted advertising [20] or even identity theft [12]. Despite many users expressing concern about their privacy online, they seldom take actions to protect it [45]. This is known as the privacy paradox [61]. Often privacy breaches could be reduced with small changes to user behaviour [32, 46]. If users could be persuaded to change their actions, therefore, they, themselves, may provide a solution.

Nudges and boosts are competing strategies from the behavioral sciences, which aim to change behavior [35] and have been studied for diverse reasons in the interactive IR community [2, 7, 38, 91]. Nudges exploit human cognitive biases to subconsciously influence behaviour towards favourable outcomes, whereas boosts empower users by providing more information or a better environment to support their decision making process [35]. Previous work has shown that carefully designed nudges and boosts can reduce privacy impacts [91] and cookie acceptance [17], increase password strength [84], get users to pay a premium for privacy during online shopping [22], encourage reflection on privacy behavior [62] or otherwise promote behavior beneficial to privacy [1]. Such interventions have, however, mainly been studied in highly controlled lab studies. It is unclear, to what extent nudge and boost interventions can reduce privacy invasion in real situations during every day internet usage.

To learn if nudges and boosts can help users to better preserve their privacy, we first conducted a between-groups online survey to evaluate 6 different boosts. Afterwards, we deployed both interventions over three 1-week-long study phases using a browser extension in a mixed design with the two interventions and a control condition. Boosts improved participants’ privacy knowledge in the short term, but knowledge gained is not retained over longer time periods. Since knowledge influences behavioral intentions, this could lead to a reduced effectiveness of boosts over time in terms of behavioral change [39]. Changes to cookies were reduced in the nudge group and third party requests in the boost group during and after the intervention deployment. This suggests that nudges and boosts could lead to higher browsing privacy in users.

2 RELATED WORK

We describe related work concerning online privacy and summarize how web tracking compromises users’ privacy. We then outline different technological and behavioral counter-interventions.

2.1 Current State of Online Privacy

Privacy is a complex construct, representing: "... control over transactions between person(s) and other(s), the ultimate aim of which is to enhance autonomy and/or to minimize vulnerability." [53, p.10]. Online privacy is relevant for users' everyday lives and awareness of this has increased since the General Data Protection Regulation (GDPR) became active in Europe on May 25 2018¹. On one hand, this can be a tool for users to control their data online [72], but on the other the requests for consent may just be a source of annoyance for users [86]. Consent notices can also use design elements to nudge users to disclose more [51], thus counteracting their purpose. Possibly in response to this, German legislation recently banned pre-selected privacy invasive choices in such notices [78]. This is an example of an attempt to protect online privacy through law.

Even though there are other threats to online privacy, such as insecure data storage, we focus on web tracking, since it can be operationalized to measure browsing privacy. Web tracking means that personal information is collected about users' online behavior [12]. Users cannot completely control this and in some cases are not aware of it [79], contrary to the way they control disclosure of, e.g., their contact data on social media. However, it is possible for them to influence web tracking, unlike e.g. data storage security.

Tracking methods developed from early session-only methods [12] to the various cross-session methods of today, such as cookies and browser fingerprinting [24]. There are many personal and societal harms that result from such practices: Web tracking can lead to search or price discrimination [36, 59] or targeted advertising, which can also be discriminatory [20]. The large amount of data collected on individuals online makes it easier to steal identities using both data actively disclosed on the web and accumulated by web tracking [12]. In a broader context, web tracking can be used to sway elections, as was investigated for the 2016 US election [42].

In general, internet users are concerned about their privacy online [6, 25, 74], but that does not mean that they actively protect it [45]. This discrepancy - the privacy paradox [61]- has been extensively studied in previous work, both providing evidence supporting [e.g. 15, 77], and contesting the notion [e.g. 49, 82]. In case of a privacy paradox, users do not take actions to protect their privacy, e.g. disclose less information online [77] or pay more for less privacy invasive services [9], even though they are concerned about online privacy. When the privacy paradox is absent, more concerned users act on their concerns and try to protect their privacy more than less concerned users, e.g. by changing their privacy settings [11].

2.2 Technological Interventions

The development of technology to reduce privacy invasion is like a cat-and-mouse game with technology being countered by newer methods of tracking. Tracking protection often makes use of black-lists [12], but it is difficult to correctly identify trackers [27]. The use of heuristics to identify trackers, complete blocking of javascript [57], filtering or changing requests, e.g. through query obfuscation [33, 63], or hiding the IP address, e.g. using anonymous proxy servers or VPN services are all different methods to protect online privacy [12]. All of these, however, require a certain degree

of technological knowledge. Using privacy-focused search engines or private browsing mode [12] are more accessible strategies. Contrary to popular belief, the latter does not prevent tracking, but only prevents attackers with access to the user's computer from seeing the user's browsing history [30]. Other techniques which preserve privacy include omitting data to make individual users more similar [50, 76] or using communities to hide individual users' data [71]. However, individual users cannot easily use these techniques, since they have to be implemented by software providers.

2.3 Behavioral Interventions

Implementing technological measures requires action from users, e.g. installing a browser extension to block tracking, or actively changing browser settings, but lack of knowledge, both about their existence and their application hinders their adoption [69, 81]. Increasing privacy during web browsing, therefore, requires users to change their behavior. Changing user behavior such that it aligns better with their privacy concerns can prevent user regret [88]. In this work, we examine two strategies from the cognitive and behavioral sciences (boosting and nudging) which aim to encourage behavioral change to reduce risk of harm to the individual and more broadly to society. Other strategies apply to communities as a whole, e.g. by influencing social norms [67], however, since privacy preferences vary, we focus on the two most popular interventions targeted at individuals, which we introduce in the following.

Nudges exploit systematic cognitive biases to influence behavior [80]. These biases are seen as the source of humans' failure to make completely rational decisions due to their lack of resources, termed bounded rationality [70]. Nudges are "any aspect of the choice architecture that alters people's behavior in a predictable way without forbidding any options or significantly changing their economic incentives" [80] (p. 6). A classic nudge is one that manipulates the environment in a non-transparent way and sets a default opt-out mechanism which does not prevent the individual from other options (e.g. in some countries, being an organ donor is the default and one must opt-out). 'Educative' nudges, such as warning labels [22, 91], are a more transparent approach which address some ethical concerns raised with nudging [35, 39], but they do not empower individuals with skills to overcome their cognitive biases.

Boosts address the issues raised with nudging [35, 39] and also promote behavioral change. Boosts assume that bounded rationality stems from humans' use of heuristics to make not optimal but satisfying decisions [34]. Contrary to nudges, boosts explicitly attempt to heighten users' knowledge and skills to better cope with the environment in which they make decisions [40]. As boosts intervene in a transparent manner, they are considered ethically sound [35, 39]. An example of a boost is teaching users heuristics to help them make better decisions, e.g. showing them which types of web sites generally contain fewer cookies and third party trackers, so they can protect their privacy by reducing the frequency with which they visit these sites. Boosts can give access to new abilities applicable in a wide range of contexts or focused on a specific situation [40]. Educative nudges overlap somewhat with short-term boosts, which promote a competence useful in a specific situation [40].

Both nudges and boosts have been used to promote online privacy. Nudges were implemented as privacy indicators in search

¹The GDPR is an European regulation focused on privacy and data protection, see also <https://gdpr-info.eu/>.

engines [22, 82, 90] and app stores [43], and by reranking or filtering search results based on privacy metrics [91]. Social nudges – interventions explicitly designed to exploit humans’ social nature and their tendency to be influenced by their peers [80]– can reduce cookie acceptance among Amazon Mechanical Turk users [17]. They compare users’ current behavior with their peers’ more preferable behavior and have also been effectively applied in other domains, such as energy conservation [3] or reducing food waste [21].

Boosts have also been applied to online privacy. They were used to make privacy policies more accessible, for example, by using a plugin to warn if a website’s privacy policy does not correspond with a user’s expectation [19], or by presenting concise summaries of privacy policy content in users’ current context of use [26, 62]. Another application was promoting secure actions on the internet, including actionable tips [87]².

2.4 Summary

While browsing the internet, it is challenging for users to control their personal information, since many parties have an interest in it [12, 24]. Technological interventions can help users protect their data [12], but even concerned users often do not use these adequately [45]. Behavioral interventions, such as nudges or boosts, can promote privacy, although the evidence for this has come from simulated laboratory studies [17, 90]. While longitudinal naturalistic studies are a well-known tool in the IR and HIB community, e.g. [8], to our knowledge, no such study has yet been conducted to directly compare nudges and boosts with regards to browsing privacy. Consequently, we do not know how effective nudges and boosts for privacy are in practice.

3 DESIGNING BOOSTS

Since boosts aim to educate, comprehension is crucial to their effectiveness [35]. We evaluated various boosts with an online survey to then use the easiest to understand and most effective boosts in a naturalistic study. Nudges differ in that they aim to influence users subconsciously [80]. As such, it is not necessary to evaluate them similarly. The boosts were inspired by previous work, from which we extracted information pertaining to level of privacy invasion of different types of websites [23, 85], as well as simple strategies to preserve privacy, which do not require much technical knowledge, such as using private browsing [83], disabling third party cookies [23] and using adblockers [32]. We selected facts relating to these aspects for boosts based on browser-independent applicability, and accessibility even for non technically adept users. The survey was conducted on the crowd sourcing platform Prolific³ in a between-subjects design with seven conditions. There were six manipulation conditions, one for each possible boost from related work, and one control condition, where participants were not exposed to a boost.

3.1 Procedure

The survey was pretested with eight subjects to determine its duration and ensure sufficient payment of participants. One boost, which caused misunderstanding in the pretest, was excluded.

In the final version, informed consent was obtained before participants were randomly allocated to one of seven conditions. In the manipulation conditions, they saw one boost and were asked two questions referring to their subjective comprehension of this information. Examples of the evaluated boosts can be found in Table 1, while all evaluated boosts are on Github⁴. Following recommendations from Prolific, an attention check question was included [64]. On the next page, participants first had to imagine they were browsing the internet and then answer questions, which required them to use the knowledge from the boost. We utilized several broad types of questions, whereby participants had to select specific URLs from different website types (Which of the following websites share data with the [most/least] third parties?), or where they either had to select correct statements or estimate numbers of third parties (Imagine that on a website, there are normally 10 third parties. How many third parties are there in the given circumstances? Please state a value between 0 and 10) pertaining to situations depicted in screenshots of Firefox and Chrome browsers⁵. There were questions pertaining to each type of boost evaluated in the survey to verify knowledge gain through that boost.

In the control condition, participants did not see a boost, but then had to answer comprehension questions pertaining to all the boosts. To prevent their workload from being much higher than in the manipulation conditions, we reduced the number of questions asked per boost and selected these question subsets randomly.

All participants also rated the helpfulness of each boost on a seven-point likert scale from not at all helpful to very helpful (concerning making their browsing more private). Finally, demographic information was collected and the participants were debriefed.

3.2 Participants

A total of 127 participants took part via Prolific, taking on average 6.56 minutes to complete the study. The conditions received roughly the same number of participants, with most of the conditions being assigned to 18 participants, except the adblocking (17) and entertainment (19) conditions. 70 of the participants identified as male, 56 as female, and 1 as diverse. They were between 18 and 67 years old ($M = 30.79$, $SD = 10.77$) and most were either working (66) or students (43), but 17 also reported currently not being employed, and one person was retired. The participants had a relatively high level of education, with 67 having at least a bachelor’s degree, 14 a finished vocational training, 33 the entrance qualification to higher education and 19 had graduated from other levels of schooling.

3.3 Results

All data analyses in this paper were conducted using the Gnu R software [66] and assumed significance at $p = .05$. To examine whether boosts were successful in conveying knowledge, participants’ answers in a boost condition were compared to answers to the same type of question of participants in the control condition.

For the website category boosts (news, entertainment, education) and the private browsing boost, the percentage of correct answers was calculated for both the manipulation and the control conditions.

²Bavel et al refer to this approach as nudging, but according to the definitions above, we consider it a boost.

³<https://www.prolific.co/>

⁴<https://github.com/blueCat11/chiir-2021/tree/main/boosts>

⁵Exact phrasing of these questions on Github: <https://github.com/blueCat11/chiir-2021/tree/main/questionnaires>

Group	Intervention
Control	no intervention
Nudge	Once per day one of: • average of total privacy points per visited site • number of privacy points achieved through browser settings (average per visited site) • number of privacy points achieved through website visits (average per visited site) • average number of third party requests per visited site • average number of added cookies per visited site
Boost	Once per day one of: • Entertainment websites have more cookies and third parties • Education websites have less cookies and third parties per page than most other kinds of websites • Blocking third party cookies in the browser settings leads to a reduction of the number of third parties per page by about 30%. • By using private browsing (Firefox) or incognito mode (Chrome), cookies are deleted automatically after the browser is closed. • By using an adblocker, the number of third parties per page is reduced by 40%, even without changing the blocker settings.

Table 1: Interventions for each experimental group during the 2nd week of the naturalistic study, translated from German by first author. Only includes the interventions used in the naturalistic study; does not contain the news boost.

Due to the small number of questions, possible percentages were few, and thus the data cannot strictly be considered metric, so a Wilcoxon rank sum test for ordinal data was conducted. Comprehension questions for third party reduction and adblocking boosts asked participants how many third parties are present on average on a certain site with different browser settings. Assuming the presence of ten third parties with default settings, the correct answer would have been seven in case third party cookies are blocked, and six when an adblocker is used. To measure the deviation from this correct value, the absolute value of deviation from the correct value was calculated and averaged over all the relevant questions answered by the specific participant. Since normality was violated, Wilcoxon rank sum tests were also conducted for these two boosts. Results are in Table 2.

The results were used to select the boosts for the naturalistic study, see Table 1. All boosts, except the third party reduction boost, were effective: Participants in the manipulation conditions provided more correct answers or answers that were closer to the correct answer than participants in the control condition. This boost was the only one where we could directly measure whether participants heeded the advice it contained, so we retained it in the naturalistic study, since it also did not significantly worsen participants' knowledge. However, the boost about news websites was removed for

	control		manipulation		W	p
	M	SD	M	SD		
boost						
news	0.13	0.34	0.47	0.47	85	.017*
entertainment	0	0	0.97	0.11	0	<.001***
education	0.06	0.24	0.89	0.21	11	<.001***
private browsing	0.67	0.24	0.85	0.21	99	.028*
adblocking	2.72	1.04	1.50	1.75	348	.002**
third party reduction	3.86	2.07	3.67	3.86	638	.914

Table 2: Results of Wilcoxon rank sum tests to detect differences between boost and manipulation conditions

ethical reasons, since we did not want to promote filter bubbles by encouraging participants to limit their news information sources.

4 NATURALISTIC STUDY

To determine the influence of nudges and boosts on browsing privacy, a three-week study was conducted in a naturalistic environment on the participants' own devices, using a browser extension.

4.1 Study Design

We utilized a mixed design, with the study phase (pre-intervention, intervention, post-intervention) being a within group and the condition, e.g. the intervention to which the participants were exposed (control, nudge, boost) a between group independent variable.

Since browsing privacy is not a concept that can be directly measured, we used previously applied proxies as dependent variables: the number of new cookies stored after visiting a given site [56] and the number of requests to third parties⁶ [e.g. 23, 29]. Although neither metric is exclusively related to tracking (cookies can be used to provide functionality, such as remembering that a user was already logged in, and third party requests may also load media), they are deemed appropriate proxies given their strong association with tracking [57]. We consequently assume that higher levels of browsing privacy occur with lower numbers of third party requests and cookies changes.

Behavioral features, such as which types of websites were visited, and browser settings were also examined together with participants' self-reported data on different concepts related to privacy. We measured affinity for technology interaction (ATI) [28] to further specify the sample. Privacy concerns were measured with a German translation [37] of the Internet Users' Information Privacy Concerns questionnaire (IUIPC) [52]. Together with general knowledge about privacy, as measured with the Online Privacy Literacy Scale (OPLIS) [55, 81] and self-reported privacy behavior, these measures were used to study the relationship between reported concerns and actual behaviour. The questionnaire for privacy behavior was adapted from Zimmerman et al [90], and translated to German by the first author. Boost knowledge was also measured, using the items used in the control condition of the boost evaluation.

⁶A third party is any site, other than the one which a user is currently visiting, which receives data about them while they are interacting with the original site [65]

(a) Screenshot of a nudge using privacy points

(b) Screenshot of the entertainment boost

Figure 1: Intervention examples, translated by first author

4.2 Interventions

Examples of the interventions, as users saw them during the study, are depicted in Figure 1. We used the facts evaluated during the online survey as five different boosts (see Table 1).

Nudges were implemented as social nudges, using the proxies for browsing privacy. This kind of nudge was shown to be effective both in the privacy domain [17] and in other contexts [21]. The same number of variants were used for nudges and boosts to make the learning effect through multiple exposure to the same intervention comparable for participants in different conditions. To achieve this we used an additional measure for browsing privacy in nudges, which we called privacy points. They were assigned to users when they exhibited behavior protecting their privacy, and detracted when they encountered third party requests or cookies. The different kinds of nudges used in the study are summarized in Table 1. A participant's average value of browsing privacy was compared to all of the participants in the study. We combined the textual representation of a participant's rank within the group with a visualization of privacy on a scale from red, for comparably privacy invasive, to green, for comparably not as privacy invasive.

4.3 Apparatus

The study was implemented using a browser extension, a PostgreSQL database backend and an Application Programming Interface (API), with which the extension communicated. The browser extension served three main purposes: It collected data on participants' browsing behavior related to our dependent variable

browsing privacy, while preserving their privacy. For example, we did not collect the URL of the site they visited, but categorized websites through the use of a domain categorizing API⁷. Using this API each website was assigned up to three categories, whereby the first assigned category can be considered the main category. In case of sensitive categories, e.g. adult content, the categories were anonymized to *uncategorized*, which was also assigned when a categorization was not possible through the API. The API was chosen by comparing multiple available APIs to a ground truth established by categorizing a set of websites by hand, and then choosing the API with the best agreement with the ground truth. Information on four different privacy related browser settings (Do Not Track, private browsing, cookie blocking policy, and WebRTC IP handling policy) was additionally collected for each website visit, and we used these to calculate the privacy points measure used in nudges.

Nudges or boosts were shown to participants once per day during the intervention phase via a modal dialog. This interruption of participant's normal browsing provided a small measure of control, since missing the intervention completely was not possible in this setup. Participants were also able to access study information, such as the lead author's contact data, and depending on the study phase, additional information, through the extension's icon in the browser tool bar. At the end of the study, a participation code was made available to the participants through this interface, which they used to confirm participation to obtain compensation in a way that did not connect the data collected by the extension with their identity.

The extension was implemented for Firefox and Chrome browsers, as they are currently the most used browsers on desktop devices in Germany, where the study took place [75], and they enable extensions to be implemented using a unified API [58]. The extension was based on previous work by [62] and [68], and we provide the adapted code on Github⁸. The extension was extensively tested prior to the study.

4.4 Procedure

Recruited participants received instructions per e-mail and provided informed consent two days before the study began. Participants received illustrated instructions to install the extension on the browser of their choice. They were encouraged to ask questions or report any problems to the author. When the study extension was correctly installed, it requested an anonymised participant id and assigned the participant to one of the three conditions. Participants filled out pre-study questionnaires and were prompted to use their browser normally during the course of the next three weeks. Participants were instructed only to install the extension on the device and browser, which they use most often.

Each study phase lasted one week. The first week was a control phase to measure browsing behavior (see above) for all participants without the influence of an intervention. Since behavior change could occur after simply contemplating privacy issues due to participating in privacy related questionnaires [54], this pre-intervention phase also served to allow behavior to normalize before the interventions. During the second week, participants in the nudge and

⁷<https://docs.webshrinker.com/v3/website-category-api.html>

⁸<https://github.com/blueCat11/chiir-2021/tree/main/extension>

boost conditions were exposed to appropriate interventions once per day. The presented intervention (either a nudge or a boost) was chosen randomly from the set of previously not yet displayed interventions, and two interventions were selected randomly to be displayed twice. The nudges and boosts presented during the study phases are summarized in Table 1. During the third week, the interventions for the nudge and boost groups stopped. Nudges and boosts are often claimed to differ in the way their effects last when the nudge or boost itself is not present anymore [39]. The final week captured behavior after nudges and boosts were no longer in effect.

On completing the study, participants filled out the post-study boost knowledge and demographic questionnaires. They received debriefing information about the study, explaining the between-groups design, nudges and boosts. A participation code was provided which allowed them to sign up for course credit compensation or a voucher lottery, and a means to access all possible nudges and boosts, as related to their own behavior. Thus, in the end, all participants received the same information, so that control group participants were also able to benefit from any positive effects that the nudges or boosts may have had. Participants were thanked and provided with instructions on how to uninstall the extension.

4.5 Participants

76 German speaking participants who use Chrome or Firefox on desktop or laptop computers were recruited via university and social media platforms and a snowball sampling procedure. Of these, 68 participants filled out the questionnaires at the beginning of the study, resulting in a response rate of 89%. At the end of the study, 60 participants filled out the questionnaires (12% drop-outs) and there was one participant (p 122), who did not fill out any questionnaires, but who provided browsing data. The data of the drop-outs and p 122 was nevertheless used in analyses where appropriate, since installing the browser extension was considered informed consent. Participants were compensated with course credits if they were students, and two 30€ vouchers were additionally raffled among the participants, after the study was finished.

Of those who completed the study, 30 identified as female, 29 as male and 1 as diverse. The participants were aged between 19 and 60 ($M = 25.52$, $SD = 8.8$), with 17 with at least Bachelor's degree and 38 with an entrance qualification to higher education. The remaining five named vocational training or graduation from other levels of schooling as their highest level of education.

During the study, 36 of the participants used Chrome, and 24 used Firefox. Windows was used by 47 participants, while 12 used MacOS, and 1 person used Linux. Of the participants who finishing the study, 21 were in the nudge condition, 18 in the control condition and 21 in the boost condition. Of all those where browsing data was available, the conditions were not quite as equally distributed, with 25 people in the nudge and boost conditions, but only 19 in the control condition.

4.6 Results

In our analysis of the naturalistic study, we investigated several aspects of our data. First, we further characterized our sample by

looking at their device and browser usage and their browsing behavior and were able to verify the information concerning website categories and privacy, which was used in boosts. Second, we used the self-reported and behavioral data to determine if there was evidence for the privacy paradox. Third we examined whether boosts succeeded in conveying knowledge, similar to our online survey. Finally, we investigated the influence of nudges and boosts on browsing privacy during different phases of the study.

4.6.1 Privacy Attitudes and Browsing Behavior. The sample ($N=68$) can be characterized as being more knowledgeable about privacy than the general German population ($M=74.6$ th percentile, $SD=20.4$), according to their OPLIS-score. Participants' privacy concern, as measured by UIIPC [52] was moderate to high ($M = 5.59$, $SD = 0.71$) and they used between 0 and 7 ($M = 2.67$, $SD = 1.91$) different measures to protect their privacy. Some also used browsers or search engines, such as Tor or Duckduckgo, which protect privacy more than the most common variants.

Although we instructed participants to use mostly the device and browser with the extension, multiple device and browser usage was self-reported by participants at the end of the study. However, we only recorded behavior on the main device and browser. In most cases of multiple device or multiple browser usage, the participants used the study device and/or study browser more or equally often compared to other equivalent devices and/or browsers.

During the study, 115,545 website visits from 30 different main categories were recorded. The number of visits per website category varied from a minimum of 3 (*Jobs & Careers* category) to a maximum of 28322 (*Entertainment* category), but the mean number of visits per category was 3852 ($SD=7601$). In summary, with some exceptions, the participants in this study visited similar types of websites across conditions and phases of the study. These exceptions were often due to certain categories of websites being visited by few distinct participants, so there is not enough data to justify a trend.

At the start of the study, the number of active users of the extension increased over the first two days, because some participants did not begin on the designated first day. For all conditions, the number of daily internet users declined slightly towards the end of the study, when only 36 participants in total were recorded as browsing on the last full day of the study. Overall, participants visited at least one website on between 1 and 22 days ($M = 15.1$, $SD = 5.7$). In general, the number of daily website visits for the three conditions are relatively similar, and relatively stable during the whole duration of the study.

4.6.2 Confirming Boost Information. Descriptive analyses allowed us to determine whether the advice provided in the boost conditions relating to the privacy risks and web site categories was true. The following results consider the main category assigned to a website. We consider entertainment, education and news sites, while all other website visits fall into the *other* category.

In our data, both entertainment ($M = 30.4$, $SD = 163.3$) and news sites ($M = 11.5$, $SD = 21.9$) had more third party requests than other sites ($M = 7.9$, $SD = 91.6$), and education websites ($M = 3.00$, $SD = 19.0$) had fewer, which replicates previous work [23, 85]. However, concerning cookie changes, this number was smaller for entertainment sites ($M = -0.12$, $SD = 6.73$) than other sites ($M = -0.06$, $SD = 19.2$), and larger for education sites ($M =$

predictor	estimate	SE estimate	robust CI		z	p
			lower	upper		
intercept	8.66	3.27	2.03	15.3	2.65	.008**
OPLIS	-0.62	0.22	-1.08	-0.16	-2.85	.004**
IUPC	-1.49	0.58	-2.62	-0.35	-2.57	.010*
OPLIS:IUPC	0.12	0.04	0.04	0.20	3.10	.002**

Table 3: Parameter estimates w/ robust 95% CI for Poisson regression model

0.20, $SD = 8.67$), even though news sites ($M = 1.44$, $SD = 3.63$) still had the largest number of cookie changes of these four categories. As in previous work [85], our metric observes the change in cookies and not the number of cookies so this may be an explanation for the difference.

4.6.3 Investigating the Privacy Paradox. To determine whether the privacy paradox applies for our study participants, the influence of privacy knowledge and privacy concerns as independent variables on self-reported privacy related actions as the dependent variable was investigated using Poisson regression. It might seem intuitive, that more protective actions are undertaken with heightened privacy concern, however, according to literature supporting the privacy paradox, this is not the case [45, 77]. Without knowledge about privacy, it is hard to judge threats, and know which actions to take [13], so privacy knowledge is included as a predictor. Since the number of privacy actions was measured as a theoretically unconstrained count variable, a Poisson regression was used [18]. Further assumptions were tested using a poissonness plot [41] and tests for over and under-dispersion, and were judged to be accurate. Stepwise regression was used adding privacy knowledge, privacy concern, and their interaction into the model one after the other, as each significantly improved the model fit. The parameter estimates for the final model are in Table 3.

Case-wise diagnostics show the analysis to be reliable. The significant interaction was plotted using the *interact.plot*-function from the *interactions* package [47], see Figure 2. It shows that when privacy concern is above average (+1 SD), the number of privacy protective actions becomes larger with growing privacy knowledge. With average privacy concern, the slope of the regression line is much less steep, but the predicted number of protective actions still becomes larger with growing privacy knowledge. On the contrary, when privacy concern is below average (-1 SD), the number of privacy protective actions does not become larger with growing privacy knowledge, but decreases.

However, self-reported privacy behavior differs from actual behavior [45], so this was additionally investigated. Proxies for browsing privacy can be considered somewhat related to privacy behavior, even if they are not quite the same. It is assumed that they reflect the outcomes of privacy behavior so that if a participant behaves in a way to achieve high browsing privacy, then the proxies for browsing privacy, as utilized in this study, will be low.

For each of the two proxies for browsing privacy and each participant, an average value was calculated including all the website

Figure 2: Interaction of final model, different color regression lines for M and M ± 1 SD of IUPC values, w/ 95% CI

visits of the participants during the first week. The first week established a baseline of behavior for all participants when none of them had been exposed to any intervention yet. One participant was excluded from this analysis, because they visited only three websites on the final day of the first week.

Privacy concern and privacy knowledge did not predict a significant amount of the variance in the average number of third party requests during the first week of the study, $F(3, 61) = 0.81$, $p = .81$, $R^2 = 0.02$, $R^2_{Adjusted} = -0.03$. Likewise, they also did not predict a significant amount of the variance in the average number of changes in cookies during that time period, $F(3, 61) = 0.24$, $p = .87$, $R^2 = 0.01$, $R^2_{Adjusted} = -0.04$. In general both models do not seem to be a good fit for the data, as multiple cases cause concern, both among potential outliers and the total sample. Violated assumptions, such as non-normal residuals, suggest that these two models do not generalize well to the population. Thus, despite some participants being concerned about privacy and reporting that they take steps to address this, in practice there is little evidence of improved outcomes as we measure them.

4.6.4 Boost Knowledge. To establish if experiencing boost interventions resulted in increased knowledge, we compared responses of participants from the boost condition to those in the other two conditions, who were not exposed to boosts. All participants answered the same questions as the control-group participants in the online survey, both at the beginning, and at the end of the three week study. A one-way ANOVA was conducted using the difference between pre-study and post-study boost knowledge as the dependent variable and condition (boost, other) as the independent variable. Boost knowledge was calculated as the sum of the correct questions, where there was only yes or no as an answer possibility. In cases where participants had to give a freetext answer, they were assigned fractions of points depending on how far away their answer was from the correct one. The difference was measured by subtracting the pre-study boost knowledge from post-study boost knowledge. The effect of condition on the difference between pre-study and post-study boost knowledge was not significant, $F(1, 58) = 0.52$, $p = .67$, $\omega^2 = -0.01$.

Dependent variable	Effect	F	df	p
average amount of third party requests	condition	0.18	2	<.001***
	study phase	12.8	2	<.001***
	condition:study phase	0.38	4	.490
average cookie change	condition	44.2	2	<.001***
	study phase	46.2	2	<.001***
	condition:study phase	1.43	4	.230

Table 4: Results of the factorial ART for average amount of third party requests and average cookie change

4.6.5 The Effect of Nudges and Boosts on Browsing Privacy. Distributional assumptions did not hold for multilevel modeling, and since the two predictor variables of the most interest, condition, and study phase, were categorical variables, an aligned rank transform (ART) was performed [89]. It can be used as a non-parametric alternative to repeated-measures ANOVA⁹. The dependent variables were normalized over the number of site visits per day.

The first website visit for each participant id was excluded from the analysis, as the change in cookies per website was measured as the difference in cookies from the previous to the current website visit. Consequently, for the first website visit, all cookies collected on the participant’s device previous to participation in the study were assigned to this website visit, which is not accurate. The results of this analysis for both the dependent variables investigated above are in Table 4. Post-hoc tests using the Tukey method to adjust p-values were conducted to analyze the highly significant main effects of condition and study phase on the proxies for browsing privacy. The results were averaged over the levels of the variable currently not under examination. According to these, the average amount of third party requests in the first week differed significantly from those in the second week $t(128) = 3.42, p = .002$ and the third week $t(106) = 5.48, p < .001$. However, there was no significant difference between the second and the third week with respect to the average amount of third party requests, $t(107) = 1.88, p = .15$. The average amount of change in cookies in the first week differed significantly from both those in the second week $t(141) = 5.65, p < .001$ and those in the third week $t(105) = 5.89, p < .001$, but there was no significant difference in changes to cookies between the third and second week, $t(107) = -0.14, p = .99$. No significant differences were found between the conditions and the proxies for browsing privacy. This finding suggests a continued effect of the interventions even after they were removed again in week 3.

These results are visualized in Figure 3 for third party requests, and in Figure 4 for the change in number of cookies. They show a drop in the number of third party requests for the boost condition during the intervention phase, and this lower number stays stable for the third week of the study. For the number of changes to cookies, this is not the case. Instead, the levels for boosts are similar to the control group, and browsing privacy for the nudge group starts out lower (depicted by higher levels of the proxy) and then increases during the intervention and post-intervention phases. Confidence intervals for the first week are generally larger than

Figure 3: Effect of condition and study phase for average number of third party requests (w/ 95% CI)

Figure 4: Effect of condition and study phase for average change in cookies (w/ 95% CI)

for the other two weeks, maybe because due to some participants starting the study late, they had fewer website visits overall.

5 DISCUSSION

Besides largely confirming previous work on the privacy invasiveness of different types of websites, we investigated three facets of naturalistic privacy research. First, we evaluated whether being exposed to boosts changed users’ knowledge about privacy. Participants in the online survey were significantly more knowledgeable when they were queried directly after having been exposed to boosts. In the naturalistic study, knowledge was not improved for boost group participants a week after their exposure. It could be that contrary to what is claimed about boosts [39], their effect weakens after the intervention, or that the knowledge conveyed in the boost is not retained. Second we examined the widely known privacy paradox phenomenon [45, 61]. Using self-reported behavior, our analyses suggest that the privacy paradox does not apply to our data, however using log data (e.g. the proxies for browsing privacy), the opposite is the case. It must be taken into account that the behavioral measures used in this analysis were only partially under the users’ control. While they were able to influence their browsing privacy by visiting certain types of websites, or using certain browser settings, outside factors may have required them to visit websites that add many cookies or send many requests to third parties. Additionally, these facets of browsing privacy were not visible to participants. Consequently, the behavioral measures might not reflect participants’ agency concerning privacy. Third, we conducted analyses using the proxies for browsing privacy to determine whether participants’ behavior changed to be more privacy preserving under the influence of nudges and boosts. There

⁹This procedure was implemented in the *ARTool* package [89].

were two significant main effects for both dependent variables. In post-hoc analyses, only differences between the pre-intervention and intervention, and pre-intervention and post-intervention study phases were significant, and there were no significant differences between the conditions. Plots (with un-normed means) indicate that browsing privacy was worst in the first week for the nudge group and change in number of cookies (Fig. 4) and for the boost participants with respect to third party requests (Fig. 3).

This study had some limitations, such as the control used and participant sample. Furthermore, even with thorough testing before the study, we encountered technical challenges with the browser extension. We were not able to retroactively change the extension during the study, but after reports from participants, the final questionnaire was adjusted to include questions about problems. Analyses were performed with and without these participants, but our conclusions did not change and therefore include all data.

Naturalistic work is valuable, because it has a high ecological validity, which is especially important in the domain of privacy [48]. In the main study, participants used their own devices, and, in most cases, the browser they were used to, over a period of multiple weeks. Laboratory studies have shown that nudges and boosts can influence behavior [e.g. 87, 91]. However, in this work, the naturalistic nature of the data also meant that noise was introduced because there was little control. There was no control over how long or how focused participants interacted with the nudge and boost interventions and over when, how often and how participants used their devices to browse the internet. This may be another reason why boost knowledge did not significantly improve over the course of the naturalistic study, since interaction with interventions was not enforced. Nevertheless, this reflects realistic conditions when deploying nudges and boosts in the wild, so it is important to do research under such conditions.

Our sample of participants is limited in size and heterogeneity. For both analyses comparing nudges and boosts, confidence intervals were wide, indicating uncertainty in the data. During the process of fitting models, there were multiple instances when models did not converge, or other issues, such as overparametrization, or singular fit, arose. This could be due to too little data. Since participants were manually recruited for this study, and required to install additional software themselves, the sample size was limited. Web browsing data was available for 69 participants, but not for all of the days in the study, and some participants only used the study extension for a few days. This may be an explanation for why the effectiveness of nudges and boosts as identified in laboratory studies [17, 87, 91] is not as clear in this work.

Our sample was also limited in diversity. Snowball sampling was used to acquire a more diverse sample, and succeeded to some extent, but even so most of the participants were students, and had a higher than average level of education. A common reason why non-students did not participate was that participation required installing software on a laptop device which was frequently used, and people from the working force declined participation e.g. because they were not allowed to install external software on their work laptop. Earlier work comparing different populations' responses to SSL warnings suggested that students' responses may not necessarily differ from responses from a more diverse sample

[73], so we conclude that our results are valuable regardless the high percentage of student participants.

The participants were advised to use the device and browser with the study extension for their internet usage during the study. Since it was difficult to recruit large numbers of participants for a longitudinal study, participants who had multiple devices they occasionally used, or who used different browsers, were also approved for participation. They were advised to use the study device and browser as much as possible, and the final questionnaire included questions on device and browser usage. Multiple device and browser usage was fairly common among participants, but in almost all cases, the study extension and browser were used more or equally often. Nevertheless, future work should examine cross-device privacy behavior.

6 CONCLUSION

This paper described an online survey examining the comprehensibility of boosts, and a three-week naturalistic experiment, whereby the effect of nudge and boost interventions on web browsing privacy was studied. It contributes to research on online privacy by comparing nudges and boosts, as two paradigms to induce behavioral change, using behavioral log data. We provide the source code for an extension which can be used to explore users' normal browsing habits while presenting interventions in a real browsing environment. Privacy is generally an interesting domain to compare nudges and boosts, since it spans the whole web and is not focused on certain sites or contexts, such as health search. As such, acquiring more naturalistic data to get a broader picture in the domain of privacy is easier than in more narrowly focused topics. Nudges and boosts for privacy have potential to help users act more privacy protectively [87, 91], but more work is needed to find out how to best present and include them naturally in everyday browsing.

Future work could consider a more detailed examination of information interaction during browsing. In this study, we aggregated the proxies for browsing privacy per day, but our data enables distinguishing website visits. This makes it possible to retrace the progression of website visits over time, including possible changes to settings, to identify possible reasons for such changes. Aside from nudge or boost interventions, behavior changes could occur dependent on the website type (e.g. dependent on the task type or website, users may want more security [10, 82]). Settings could also be changed because a website breaks due to measures intended to protect privacy, but this cannot be derived from our data.

Finally, the use of mobile devices is on the rise, and an increasing amount of people mainly access the internet through their mobile phone [16]. However, mobile devices are more vulnerable to privacy breaches, and do not offer as many privacy protective measures [60]. Most research related to privacy on mobile devices focuses on apps [e.g. 4, 5], even though browsing and searching the internet is the activity on which the second largest amount of time is spent on mobile devices [14]. Thus, future work should be extended beyond desktop and laptop computers. Multiple device usage occurred in this study, and it would be interesting to investigate how users manage their privacy on multiple devices. Studying such effects further can ultimately help align people's privacy related behavior to the level of privacy they wish to achieve, on any device.

REFERENCES

- [1] Alessandro Acquisti, Idris Adjerid, Rebecca Balebako, Laura Brandimarte, Lorie Cranor, Saranga Komanduri, Pedro Giovanni Leon, Norman Sadeh, Florian Schaub, Many Sleeper, Yang Wang, and Shomir Wilson. 2017. Nudges for Privacy and Security: Understanding and Assisting Users' Choices Online. *ACM Computing Surveys (CSUR)* 50, 3, Article 44 (2017), 41 pages.
- [2] Johannes Aigner, Amelie Durwardt, Thiemo Kersting, Markus Kattenbeck, and David Elswiler. 2017. Manipulating the perception of credibility in refugee related social media posts. In *Proceedings of the 2017 Conference on Conference Human Information Interaction and Retrieval*. 297–300.
- [3] Hunt Allcott. 2011. Social norms and energy conservation. *Journal of Public Economics* 95, 9 (2011), 1082 – 1095. Special Issue: The Role of Firms in Tax Systems.
- [4] Hazim Almuhammedi, Florian Schaub, Norman Sadeh, Idris Adjerid, Alessandro Acquisti, Joshua Gluck, Lorrie Cranor, and Yuvraj Agarwal. 2015. Your Location Has Been Shared 5,398 Times! A Field Study on Mobile App Privacy Nudging. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (Seoul, Republic of Korea) (CHI '15). ACM, New York, 787–796.
- [5] Manar Alohaly and Hassan Takabi. 2016. Better Privacy Indicators: A New Approach to Quantification of Privacy Policies. In *Twelfth Symposium on Usable Privacy and Security (SOUPS 2016)*. USENIX Association, Denver, CO.
- [6] Brooke Auxier, Lee Rainie, Monica Andersen, Andrew Perrin, Madhu Kumar, and Erica Turner. 2019. *Americans and Privacy: Concerned, Confused and Feeling Lack of Control Over Their Personal Information*. Technical Report. Pew Research Center.
- [7] Scott Bateman, Jaime Teevan, and Ryen W. White. 2012. The Search Dashboard: How Reflection and Comparison Impact Search Behavior. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '12)*. ACM, 1785–1794.
- [8] Scott Bateman, Jaime Teevan, and Ryen W. White. 2012. The Search Dashboard: How Reflection and Comparison Impact Search Behavior. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '12)*. ACM, 1785–1794.
- [9] Alastair R. Beresford, Dorothea Kübler, and Sören Preibusch. 2012. Unwillingness to pay for privacy: A field experiment. *Economics Letters* 117, 1 (2012), 25 – 27.
- [10] Annika Bergström. 2015. Online privacy concerns: A broad approach to understanding the concerns of different groups for different uses. *Computers in Human Behavior* 53 (2015), 419 – 426.
- [11] Grant Blank, Gillian Bolsover, and Elizabeth Dubois. 2014. *A New Privacy Paradox: Young People and Privacy on Social Network Sites*. Technical Report. University of Oxford, Global Cyber Security Capacity Centre. <https://ssrn.com/abstract=2479938>
- [12] Tomasz Bujlow, Valentin Carrel-Español, Josep Solé-Pareta, and Pere Barlet-Ros. 2017. A Survey on Web Tracking: Mechanisms, Implications, and Defenses. *Proc. IEEE* 105, 8 (Aug 2017), 1476–1510.
- [13] Moritz Büchi, Natascha Just, and Michael Latzer. 2017. Caring is not enough: the importance of Internet skills for online privacy protection. *Information, Communication & Society* 20, 8 (2017), 1261–1278.
- [14] Juan Pablo Carrascal and Karen Church. 2015. An In-Situ Study of Mobile App & Mobile Search Interactions. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (Seoul, Republic of Korea) (CHI '15). ACM, New York, 2739–2748.
- [15] Juan Pablo Carrascal, Christopher Riederer, Vijay Erramilli, Mauro Cherubini, and Rodrigo de Oliveira. 2013. Your Browsing Behavior for a Big Mac: Economics of Personal Information Online. In *Proceedings of the 22nd International Conference on World Wide Web (Rio de Janeiro, Brazil) (WWW '13)*. ACM, New York, 189–200.
- [16] comscore. 2015. Number of Mobile-Only Internet Users Now Exceeds Desktop-Only in the U.S. <https://www.comscore.com/Insights/Blog/Number-of-Mobile-Only-Internet-Users-Now-Exceeds-Desktop-Only-in-the-U.S>. Last accessed: 2020-07-13, archived at <https://archive.st/elq5>.
- [17] Lynne M. Coventry, Debora Jeske, John M. Blythe, James Turland, and Pam Briggs. 2016. Personality and Social Framing in Privacy Decision-Making: A Study on Cookie Acceptance. *Frontiers in Psychology* 7 (2016), 1341.
- [18] Stefany Cox, Stephen G. West, and Leona S. Aiken. 2009. The Analysis of Count Data: A Gentle Introduction to Poisson Regression and Its Alternatives. *Journal of Personality Assessment* 91, 2 (2009), 121–136.
- [19] Lorrie Cranor, Praveen Guduru, and Manjula Arjula. 2006. User Interfaces for Privacy Agents. *ACM Transactions on Computer-Human Interaction* 13, 2 (June 2006), 135–178.
- [20] Amit Datta, Michael Carl Tschantz, and Anupam Datta. 2015. Automated experiments on ad privacy settings. *Proceedings on privacy enhancing technologies* 2015, 1 (2015), 92–112.
- [21] Anna de Visser-Amundson and Mirella Kleijnen. 2020. *Nudging in Food Waste Management: Where Sustainability Meets Cost-Effectiveness*. Springer International Publishing, Cham, 57–87.
- [22] Serge Egelman, Janice Tsai, Lorrie Cranor, and Alessandro Acquisti. 2009. Timing is Everything? The Effects of Timing and Placement of Online Privacy Indicators. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Boston, MA, USA) (CHI '09). ACM, New York, 319–328.
- [23] Steven Englehardt and Arvind Narayanan. 2016. Online Tracking: A 1-Million-Site Measurement and Analysis. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security (Vienna, Austria) (CCS '16)*. ACM, New York, 1388–1401.
- [24] Tatiana Ermakova, Benjamin Fabian, Benedict Bender, and Kerstin Klimek. 2018. Web Tracking - A Literature Review on the State of Research. In *Proceedings of the 51st Hawaii International Conference on System Sciences* (Hilton Waikoloa Village, Hawaii, USA) (HICSS'18). 4732–4741. <http://hdl.handle.net/10125/50485>
- [25] European Commission. 2011. *Attitudes on Data Protection and Electronic Identity in the European Union*. Technical Report. Special Eurobarometer 359. <https://ec.europa.eu/commfrontoffice/publicopinion/index.cfm/Survey/getSurveyDetail/instruments/SPECIAL/surveyKy/864>
- [26] Denis Feth. 2017. Transparency through Contextual Privacy Statements. In *Mensch und Computer 2017-Workshopband*, Manuel Burghardt, Raphael Wimmer, Christian Wolff, and Christa Womser-Hacker (Eds.).
- [27] Imane Fouad, Nataliia Bielova, Arnaud Legout, and Natasa Sarafijanovic-Djukic. 2020. Missed by Filter Lists: Detecting Unknown Third-Party Trackers with Invisible Pixels. *Proceedings on Privacy Enhancing Technologies* 2020, 2 (2020), 499 – 518.
- [28] Thomas Franke, Christiane Attig, and Daniel Wessel. 2019. A Personal Resource for Technology Interaction: Development and Validation of the Affinity for Technology Interaction (ATT) Scale. *International Journal of Human-Computer Interaction* 35, 6 (2019), 456–467.
- [29] Nathaniel Fruchter, Hsin Miao, Scott Stevenson, and Rebecca Balebako. 2015. Variations in Tracking in Relation to Geographic Location. In *Proceedings of the 9th Workshop on Web 2.0 Security and Privacy (W2SP) 2015*.
- [30] Xianyi Gao, Yulong Yang, Huiqing Fu, Janne Lindqvist, and Yang Wang. 2014. Private Browsing: An Inquiry on Usability and Privacy Protection. In *Proceedings of the 13th Workshop on Privacy in the Electronic Society* (Scottsdale, Arizona, USA) (WPES '14). ACM, New York, 97–106.
- [31] Michelle Geronimo. 2017. Online browsing: Can, should, and may companies combine online and offline data to learn about you. *Hastings Science and Technology Law Journal* 9, 2 (2017).
- [32] Arthur Gervais, Alexandros Filios, Vincent Lenders, and Srdjan Capkun. 2017. Quantifying Web Adblocker Privacy. In *Computer Security – ESORICS 2017*, Simon N. Foley, Dieter Gollmann, and Einar Snekkenes (Eds.). Springer International Publishing, Cham, 21–42.
- [33] Arthur Gervais, Reza Shokri, Adish Singla, Srdjan Capkun, and Vincent Lenders. 2014. Quantifying Web-Search Privacy. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security* (Scottsdale, Arizona, USA) (CCS '14). ACM, New York, 966–977.
- [34] Gerd Gigerenzer. 2006. *Heuristics*. MIT Press, Cambridge, MA, 17–44.
- [35] Till Grüne-Yanoff and Ralph Hertwig. 2016. Nudge Versus Boost: How Coherent are Policy and Theory? *Minds and Machines* 26, 1 (01 Mar 2016), 149–183.
- [36] Aniko Hannak, Gary Soeller, David Lazer, Alan Mislove, and Christo Wilson. 2014. Measuring Price Discrimination and Steering on E-Commerce Web Sites. In *Proceedings of the 2014 Conference on Internet Measurement Conference* (Vancouver, BC, Canada) (IMC '14). ACM, New York, 305–318.
- [37] David Harborth and Sebastian Pape. 2019. How Privacy Concerns and Trust and Risk Beliefs Influence Users' Intentions to Use Privacy-Enhancing Technologies - The Case of Tor. In *Proceedings of the 52nd Hawaii International Conference on System Sciences*. 10.
- [38] Morgan Harvey, Claudia Hauff, and David Elswiler. 2015. Learning by Example: training users with high-quality query suggestions. In *Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 133–142.
- [39] Ralph Hertwig. 2017. When to consider boosting: some rules for policy-makers. *Behavioural Public Policy* 1, 2 (2017), 143–161.
- [40] Ralph Hertwig and Till Grüne-Yanoff. 2017. Nudging and Boosting: Steering or Empowering Good Decisions. *Perspectives on Psychological Science* 12, 6 (2017), 973–986.
- [41] David C. Hoaglin. 1980. A Poissonness Plot. *The American Statistician* 34, 3 (1980), 146–149.
- [42] Jim Isaak and Mina J. Hanna. 2018. User Data Privacy: Facebook, Cambridge Analytica, and Privacy Protection. *Computer* 51, 8 (August 2018), 56–59.
- [43] Patrick Gage Kelley, Lorrie Cranor, and Norman Sadeh. 2013. Privacy as Part of the App Decision-Making Process. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Paris, France) (CHI '13). ACM, New York, 3393–3402.
- [44] Simon Kemp. 2019. Digital trends 2019: Every single stat you need to know about the internet. <https://thenextweb.com/contributors/2019/01/30/digital-trends-2019-every-single-stat-you-need-to-know-about-the-internet/>. Last accessed: 2020-07-13, archived at <https://archive.st/hp33>.
- [45] Spyros Kokolakis. 2017. Privacy attitudes and privacy behaviour: A review of current research on the privacy paradox phenomenon. *Computers & Security* 64 (2017), 122 – 134.

- [46] Georgios Kontaxis and Monica Chew. 2015. Tracking Protection in Firefox For Privacy and Performance. *Computing Research Repository (CoRR)* abs/1506.04104 (2015). arXiv:1506.04104
- [47] Jacob A. Long. 2019. *interactions: Comprehensive, User-Friendly Toolkit for Probing Interactions*. <https://cran.r-project.org/package=interactions> R package version 1.1.0.
- [48] Paul Benjamin Lowry, Tamara Dinev, and Robert Willison. 2017. Why security and privacy research lies at the centre of the information systems (IS) artefact: proposing a bold research agenda. *European Journal of Information Systems* 26, 6 (2017), 546–563.
- [49] Christoph Lutz and Pepe Strathoff. 2014. Privacy concerns and online behavior – Not so paradoxical after all? Viewing the privacy paradox through different theoretical lenses.
- [50] Ashwin Machanavajjhala, Daniel Kifer, Johannes Gehrke, and Muthuramakrishnan Venkatasubramanian. 2007. L-Diversity: Privacy beyond k-Anonymity. *ACM Trans. Knowl. Discov. Data* 1, 1 (March 2007), 3–es.
- [51] Dominique Machuletz and Rainer Böhm. 2020. Multiple Purposes, Multiple Problems: A User Study of Consent Dialogs after GDPR. *Proceedings on Privacy Enhancing Technologies* 2020, 2 (2020).
- [52] Naresh K. Malhotra, Sung S. Kim, and James Agarwal. 2004. Internet Users' Information Privacy Concerns (IUPC): The Construct, the Scale, and a Causal Model. *Information Systems Research* 15, 4 (2004), 336–355.
- [53] Stephen T. Margulis. 1977. Conceptions of Privacy: Current Status and Next Steps. *Journal of Social Issues* 33, 3 (1977), 5–21.
- [54] Helia Marreiros, Mirco Tonin, Michael Vlassopoulos, and M.C. Schraefel. 2017. "Now that you mention it": A survey experiment on information, inattention and online privacy. *Journal of Economic Behavior & Organization* 140 (2017), 1 – 17.
- [55] Philipp K. Masur, Doris Teutsch, and Sabine Trepte. 2017. Entwicklung und Validierung der Online-Privatheitskompetenzskala (OPLIS). *Diagnostica* 63, 4 (2017), 256–268.
- [56] Jonathan R. Mayer and John C. Mitchell. 2012. Third-Party Web Tracking: Policy and Technology. In *Proceedings of IEEE Symposium on Security and Privacy*. 413–427.
- [57] Johan Mazel, Richard Garnier, and Kensuke Fukuda. 2019. A comparison of web privacy protection techniques. *Computer Communications* 144 (2019), 162 – 174.
- [58] MDN contributors. 2020. Browser Extensions. <https://developer.mozilla.org/en-US/docs/Mozilla/Add-ons/WebExtensions>. Last accessed: 2020-06-02.
- [59] Jakub Mikians, László Gyarmati, Vijay Erramilli, and Nikolaos Laoutaris. 2012. Detecting Price and Search Discrimination on the Internet. In *Proceedings of the 11th ACM Workshop on Hot Topics in Networks* (Redmond, Washington) (*HotNets-XI*). ACM, New York, 79–84.
- [60] Alexios Mylonas, Nikolaos Tsalis, and Dimitris Gritzalis. 2013. Evaluating the Manageability of Web Browsers Controls. In *Security and Trust Management*, Rafael Accorsi and Silvio Ranise (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 82–98.
- [61] Patricia A. Norberg, Daniel R. Horne, and David A. Horne. 2007. The Privacy Paradox: Personal Information Disclosure Intentions versus Behaviors. *Journal of Consumer Affairs* 41, 1 (2007), 100–126.
- [62] Anna-Marie Orloff, Maximiliane Windl, Valentin Schwind, and Niels Henze. 2020. Implementation and In Situ Assessment of Contextual Privacy Policies. In *Proceedings of the 2020 Conference on Designing Interactive Systems* (Eindhoven, Netherlands) (*DIS '20*). ACM.
- [63] Sai Teja Peddinti and Nitesh Saxena. 2010. On the Privacy of Web Search Based on Query Obfuscation: A Case Study of TrackMeNot. In *Privacy Enhancing Technologies*, Mikhail J. Atallah and Nicholas J. Hopper (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 19–37.
- [64] Prolific Team. 2018. Using attention checks as a measure of data quality. <https://researcher-help.prolific.co/hc/en-gb/articles/360009223553-Using-attention-checks-as-a-measure-of-data-quality>. Last accessed: 2020-06-01.
- [65] Gaston Pugliese. 2015. Web tracking: Overview and applicability in digital investigations. *it - Information Technology* 57, 6 (2015), 366 – 375.
- [66] R Core Team. 2019. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. <https://www.R-project.org/>
- [67] Samuli Reijula, Jaakko Kuorikoski, Timo Ehrig, Konstantinos Katsikopoulos, and Shyam Sunder. 2018. Nudge, Boost, or Design? Limitations of behaviorally informed policy under social interaction. *Journal of Behavioral Economics for Policy (SABE)* 2 (03 2018), 99–105.
- [68] Princiya Marina Sequeira. 2019. *lightbeam-we*. <https://github.com/princiya/lightbeam-we>. Last accessed: 2020-07-13, archived at <https://archive.st/12cc>.
- [69] Martin Shelton, Lee Rainie, and Mary Madden. 2015. *Americans' Privacy Strategies Post-Snowden*. Technical Report. Pew Research Center. <http://www.pewinternet.org/2015/03/16/Americans-Privacy-Strategies-Post-Snowden/>
- [70] Herbert A. Simon. 1955. A Behavioral Model of Rational Choice. *The Quarterly Journal of Economics* 69, 1 (1955), 99–118.
- [71] Barry Smyth, Jill Freyne, Maurice Coyle, Peter Briggs, and Evelyn Balfe. 2004. I-SPY – Anonymous, Community-Based Personalization by Collaborative Meta-Search. In *Research and Development in Intelligent Systems XX*, Frans Coenen, Alun Preece, and Ann Macintosh (Eds.). Springer London, London, 367–380.
- [72] Maciej Sobolewski, Joanna Mazur, and Michał Paliński. 2017. GDPR: A Step Towards a User-centric Internet? *Interconomics* 52, 4 (01 Jul 2017), 207–213.
- [73] Andreas Sotirakopoulos, Kirstie Hawkey, and Konstantin Beznosov. 2011. On the Challenges in Usable Security Lab Studies: Lessons Learned from Replicating a Study on SSL Warnings. In *Proceedings of the Seventh Symposium on Usable Privacy and Security* (Pittsburgh, Pennsylvania) (*SOUPS '11*). ACM, New York, Article 3, 18 pages.
- [74] Tola St. Matthew and Daniel Thatcher. 2016. TRUSTe/NCSA Consumer Privacy Index Reveals Rising Consumer Concerns and a Significant Awareness Deficit; Businesses Pay as Privacy Concerns Discourage Consumers. <https://trustarc.com/study-finds-more-americans-concerned-about-data-privacy-than-losing-their-income/> Last accessed: Oct 10, 2020.
- [75] Statcounter - GlobalStats. 2020. Desktop Browser Market Share Germany. April 2019 - May 2020. <https://gs.statcounter.com/browser-market-share/desktop/germany>. Last accessed: 2020-06-02.
- [76] Latanya Sweeney. 2002. k-anonymity: A model for protecting privacy. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems* 10, 05 (2002), 557–570.
- [77] Monika Taddicken. 2014. The 'Privacy Paradox' in the Social Web: The Impact of Privacy Concerns, Individual Characteristics, and the Perceived Social Relevance on Different Forms of Self-Disclosure. *Journal of Computer-Mediated Communication* 19, 2 (01 2014), 248–273.
- [78] Tagesschau. 2020. Aktives Ja zu Cookies muss sein. <https://www.tagesschau.de/inland/cookies-bundesgerichtshof-101.html>.
- [79] Yuuki Takano, Satoshi Ohta, Takeshi Takahashi, Ruo Ando, and Tomoya Inoue. 2014. MindYourPrivacy: Design and implementation of a visualization system for third-party Web tracking. In *2014 Twelfth Annual International Conference on Privacy, Security and Trust* (Toronto, Ontario, Canada). IEEE, 48–56.
- [80] Richard H. Thaler and Cass R. Sunstein. 2008. *Nudge: Improving Decisions about Health, Wealth, and Happiness*. Yale University Press, New Haven & London.
- [81] 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)*. Springer Netherlands, Dordrecht, 333–365.
- [82] Janice Y. Tsai, Serge Egelman, Lorrie Cranor, and Alessandro Acquisti. 2011. The Effect of Online Privacy Information on Purchasing Behavior: An Experimental Study. *Information Systems Research* 22, 2 (2011), 254–268.
- [83] Nikolaos Tsalis, Alexios Mylonas, Antonia Nisioti, Dimitris Gritzalis, and Vasilios Katos. 2017. Exploring the protection of private browsing in desktop browsers. *Computers & Security* 67 (2017), 181 – 197.
- [84] Blase Ur, Patrick Gage Kelley, Saranga Komanduri, Joel Lee, Michael Maass, Michelle L. Mazurek, Timothy Passaro, Richard Shay, Timothy Vidas, Lujo Bauer, Nicolas Christin, and Lorrie Cranor. 2012. How Does Your Password Measure Up? The Effect of Strength Meters on Password Creation. In *21st USENIX Security Symposium (USENIX Security 12)*. USENIX Association, Bellevue, WA, 65–80.
- [85] Tobias Urban, Martin Degeling, Thorsten Holz, and Norbert Pohlmann. 2020. Beyond the Front Page: Measuring Third Party Dynamics in the Field. In *Proceedings of The Web Conference 2020* (Taipei, Taiwan) (*WWW '20*). ACM, New York, 1275–1286.
- [86] 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*). ACM, New York, 973–990.
- [87] René van Bavel, Nuria Rodríguez-Priego, José Vila, and Pam Briggs. 2019. Using protection motivation theory in the design of nudges to improve online security behavior. *International Journal of Human-Computer Studies* 123 (2019), 29 – 39.
- [88] Yang Wang, Gregory Norcie, Saranga Komanduri, Alessandro Acquisti, Pedro Giovanni Leon, and Lorrie Cranor. 2011. "I Regretted the Minute I Pressed Share": A Qualitative Study of Regrets on Facebook. In *Proceedings of the Seventh Symposium on Usable Privacy and Security* (Pittsburgh, Pennsylvania) (*SOUPS '11*). ACM, New York, Article 10, 16 pages.
- [89] Jacob O. Wobbrock, Leah Findlater, Darren Gergle, and James J. Higgins. 2011. The Aligned Rank Transform for Nonparametric Factorial Analyses Using Only ANOVA Procedures. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI '11)*. ACM Press, New York, 143–146.
- [90] Steven Zimmerman, Alistair Thorpe, Chris Fox, and Udo Kruschwitz. 2019. Investigating the Interplay Between Searchers' Privacy Concerns and Their Search Behavior. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval* (Paris) (*SIGIR '19*). ACM, New York, 953–956.
- [91] Steven Zimmerman, Alistair Thorpe, Chris Fox, and Udo Kruschwitz. 2019. Privacy Nudging in Search: Investigating Potential Impacts. In *Proceedings of the 2019 Conference on Human Information Interaction and Retrieval* (Glasgow, Scotland UK) (*CHIIR '19*). ACM, New York, 283–287.