



Is Personalization Worth It? Notifying Blogs about a Privacy Issue Resulting from Poorly Implemented Consent Banners

Theresa Kriecherbauer
Ludwig Maximilian University
Munich, Germany
t.kriecherbauer@web.de

Richard Schwank
Technical University of Munich
Munich, Germany
richard.schwank@tum.de

Adrian Krauss
Technical University of Munich
Munich, Germany
adrian.krauss@tum.de

Konstantin Neureither
Technical University of Munich
Munich, Germany
k.neureither@tum.de

Lian Remme
Heinrich Heine University
Düsseldorf, Germany
lian.remme@hhu.de

Melanie Volkamer
Karlsruhe Institute of Technology
Karlsruhe, Germany
melanie.volkamer@kit.edu

Dominik Herrmann
University of Bamberg
Bamberg, Germany
dominik.herrmann@uni-bamberg.de

ABSTRACT

Several websites integrate trackers without users' consent. Previous research studied whether notifying responsible website operators about such issues is an effective measure, often with limited success. Insights from marketing research suggest that personalizing notification emails may be an effective means to improve remediation rates, with previous research pointing in both directions. We studied this approach using a sample of 119 German fitness and sports blogs employing Google Analytics (GA) without user consent: In a first step, we compare the fix rate of blog operators that received a personalized notification tailored to their blog with the fix rate of operators that received a generic notification. We find that personalized notifications do neither increase remediation rate nor operators' response behavior. In a second step, we analyzed the reasons not to fix mentioned in (A) the email responses and (B) a survey sent to the blog operators. We find that they mostly center around (I) denial that a data leak exists, (II) a lack of resources to remedy the issue and (III) claims of specifically requiring GA. We hypothesize that an additional reason not to fix could be the so-called *moral credentials* phenomenon and sketch how to study that in future work.

CCS CONCEPTS

• **Security and privacy** → **Social aspects of security and privacy**; *Usability in security and privacy*; Privacy protections.

KEYWORDS

Human Factors in Privacy, Notification Study, Personalization, GDPR



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

ARES 2024, July 30–August 02, 2024, Vienna, Austria
© 2024 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-1718-5/24/07
<https://doi.org/10.1145/3664476.3664499>

ACM Reference Format:

Theresa Kriecherbauer, Richard Schwank, Adrian Krauss, Konstantin Neureither, Lian Remme, Melanie Volkamer, and Dominik Herrmann. 2024. Is Personalization Worth It? Notifying Blogs about a Privacy Issue Resulting from Poorly Implemented Consent Banners. In *The 19th International Conference on Availability, Reliability and Security (ARES 2024)*, July 30–August 02, 2024, Vienna, Austria. ACM, New York, NY, USA, 7 pages. <https://doi.org/10.1145/3664476.3664499>

1 INTRODUCTION

Cookie banners present themselves as guards to users' privacy. However, technical implementations have been found to be lackluster, resulting in personal information being logged by the third parties prior to the users' consent [12]. One way to improve that situation is to confront website operators with the fact that their site does not follow common practices to protect personal privacy (e. g. data minimisation, as mandated by Art. 5 § 1 (c) GDPR). The question how to best notify website operators to correct privacy or security faults (e. g., [5, 21, 28]) is investigated in so-called notification studies. So far, these studies have varied channel, context, message content, and sample size.

Drawing inspiration from marketing, personalization could be a tool to make notifications to website operators more convincing. Indeed, psychology research suggests that personalization leads to heightened attention and a more favorable reception of the email [23, 24]. However, some studies point into the opposite direction [17]. Therefore, we pose the following two qualitative *research questions*: (RQ1) To what extent can personalized notification emails make website operators remedy data privacy leaks? And, to better understand the effect of personalization studied in this first step: (RQ2) Which reasons not to fix do the notified subjects mention?

To answer the first research question, we conducted a covert between-subjects study with 119 blogs and informed 79 blogs about their use of Google Analytics (GA) prior to the users' consent. We used an ethical argument to create a coherent setting for personalization. 40 blogs received a *personalized* email and the other 39 a *generic* one. For personalization, we extracted information about

the bloggers manually from the blogs. In the 60 days following the initial notification, we monitored the configuration of GA on the blogs and the responses to our emails. A control group of 40 blogs received no notification for reference.

To answer the second research question, we extracted reasons not to fix the misconfiguration from any email responses we received from subjects during the study. We also offered a voluntary post-debrief survey in which subjects were questioned about their considerations (not) to fix.

In this paper, we address our two research questions in the context of fitness and sports blogs that implement GA without asking for users' consent first. We have chosen blogs, since they provide uniform meta-information (e. g., author name and publication date) that both enable personalization and ensure that personalization is comparable among websites. To further homogenize the sample, we have settled on the fitness and sport realm.

For the first research question, we find that the effort of personalization does not seem to translate into higher fix rates. Overall, 15 of the 79 notified blogs (19%) correct their GA misconfiguration – with no relevant difference between the personalized and generic email groups. For the second research question, we find that the reasons not to fix mostly center around (I) denial that a data leak exists, (II) a lack of resources to remedy the issue and (III) claims of specifically requiring GA, as expressed in the response emails. In the survey, respondents expressed similar reasons not to fix. No clear differences in reasons mentioned between the two groups are evident for most categories. However, it is interesting to note that operators who did not fix the misconfiguration replied to our messages more frequently. This leaves room for further research, e. g. on whether having demonstrated their “good will” to remedy the problem decreases their motivation to do so (*moral credentials* phenomenon [25]). One possibility would be to use no-reply email addresses to eliminate the operators' opportunity to verbally demonstrate their “good will” without acting upon it.

In summary, our study contributes by providing indications that the effort of personalization does (A) not improve remediation rate or (B) change the response behavior, with some exceptions.

2 RELATED WORK

In the work of Maass et al. [21], 4754 websites with misconfigured GA settings have been notified, evaluating the effects of different contact media, senders, and message framings. Maass et al. examined the effects of message framings with respect to privacy concerns, GDPR violations, and potential fines. They focused on reaching a large number of websites using their contact information. In contrast, our study explored the use of additional information about the website owner and content to personalize the notification. To do so, we aimed for a much smaller number of websites. In addition, our study has a slightly different focus, i. e. we focus on GA cookies before the user's explicit consent via a cookie banner, whereas Maass et al. focused on compliance violations resulting from failure to enable IP anonymization with GA.

There are many more large-scale notification studies [4–6, 18, 19, 21, 27, 28, 31, 32]. Like Maas et al., the studies investigated the effectiveness of various properties of notifications, including sender identity [6, 21], contact media (e. g., letter and email) [21], source of

contact address (e. g., WHOIS entry, CERT entry) [5, 18, 28], chosen contact channel (e. g., Google Webmaster Console, emails, landing page of a walled garden environment) [4, 19], message language [18, 32], and message framing (e. g., focusing on legal consequences, privacy concerns, or security risks) [5, 21, 31, 32]. None of these studies focused on blogs only and besides Maas et al. [21], none checked misconfigured GA settings.

In summary, the primary differentiating aspect of our research is that we evaluated the effect of personalization for issues related to GA. We used manually extracted context-specific information, e. g. the title and content of the blog's last published article.

3 BACKGROUND

In the following, we elaborate technical aspects of GA and the psychological principles of personalization.

Technical Background. GA is a web service provided by Google, designed to track and analyze user behavior on websites [11]. By integrating the JavaScript library *analytics.js* or *gtag.js*, owners can collect data about user interactions via tracking objects and the setting of cookies [13]. The user interactions are captured by tracking objects which are then transmitted to Google's servers e. g. in the US for analysis and reporting in the GA dashboard [1].

As of March 2024, the current version of GA is GA 4, which succeeded the previous version Universal Analytics. Universal Analytics ceased processing data on July 1, 2023. Automatic migration to GA 4 took place for Universal Analytics users from March 1, 2023 if they did not opt out of this service [10]. This changeover took place during our measuring time, but did not influence our research, as (i) Google Analytics 4 still logs personal data of website visitors and (ii) our scanner detects both GA versions.

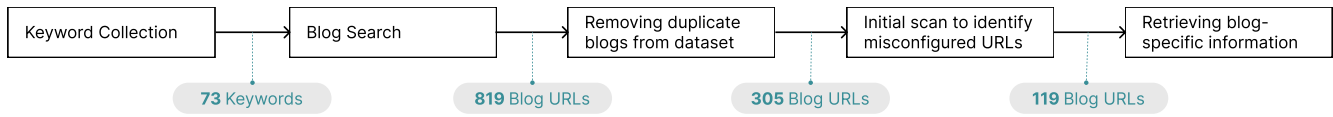
Psychological Background. Personalization of a message is the incorporation of individualizing elements without altering the factual content [8, 14]. Previous works have found that personalized advertising leads to heightened attention [24] and a more favorable reception of the message [23]. It is an open question, however, whether or not personalized messages yield more favorable effects compared to generic ones, with some studies suggesting enhanced effectiveness and others indicating the opposite [17]. This leads to the question of how personalization impacts privacy notifications, which we aim to address in our study.

Our study follows the framework of personalization laid out by Hawkins et al. [14]: (1) *Raising expectation of personalization*: Self-referencing inducing statements like “This is *just for you!*”. (2) *Identification*: Addressing the person distinctly, e. g. by name. (3) *Contextualization*: Framing the message in such a way that the recipient perceives it as more personally relevant. As Maslowska et al. [24] found that combining all three personalization methods worked best, our personalized emails also used all three means.

4 METHODS

In this section, we describe the collection process, including how we checked whether websites use GA, and the notification study. Figure 1 shows an overview of our research.

COLLECTION OF WEBSITES



TREATMENT

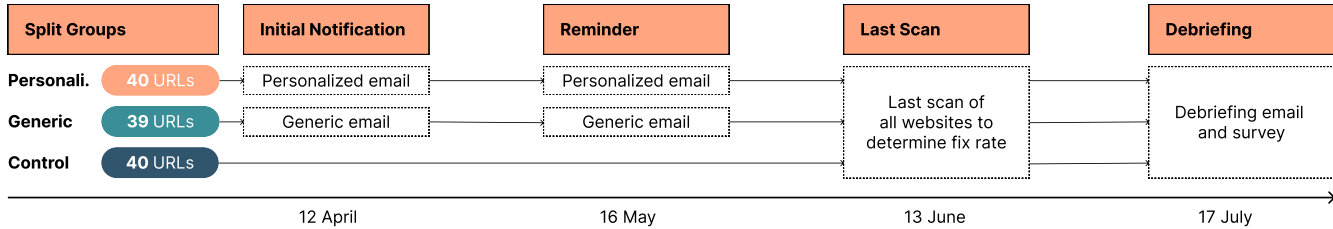


Figure 1: Overview of our research

4.1 Collection of Websites

Keyword collection. 73 keywords were extracted from lists of the most popular sports and fitness categories [15, 26]. The keywords we used can be found in the supplementary material [16].

Blog Search. For each keyword, we searched for “fitness blog <keyword>” or “sport blog <keyword>”, respectively, using the Google search engine and Google Chrome in Incognito mode. The URLs of the first ten search results were recorded. Ads were skipped, and it was also noted whether it was a blog. If a search result was a collection of multiple blogs, all of those websites were recorded. Each search query was assigned to two researchers. Any differences in the collected blogs were discussed until agreement was reached. This resulted in 819 entries. We removed duplicates, and ended up with 305 unique websites that we considered to be a blog.

Scanning for GA Issue. We scanned these 305 blogs to determine those using GA before interacting with the consent dialogues. To scan the websites, we used a *compliance scanner* which was also used by Maass et al [21]. This tool helped us to find websites on which GA was misconfigured, i. e., was loaded without the explicit consent of the user. The compliance scanner loads a website with an instrumented Chromium browser. Since the scanner does not interact with websites, any detected data collection requests sent to GA were performed before users would have consented to that. We ended up with 119 blogs that we could scan automatically and that used GA before the user explicitly agreed to it.

Information Collection. For these 119 blogs, two authors collected the following blog-specific pieces of information independently and discussed any differences afterwards to ensure the accuracy of the collected information:

- the author of the latest article or, if none exists, the name found in an “About” section, or, if also none exists, the name given in the imprint,
- the title and date of the latest article,
- whether the website offers a shop,

- whether the responsible party in the imprint is a person or a company, and
- whether the blog was active in the last 6 months or not.

The name and title are used in the personalized notification email. All other collected pieces of information are used to ensure stratification of the sample and as control variables during the analysis.

4.2 Notification Study Setup

The lower part of Figure 1 titled “Treatment” describes our procedure in this study which will be explained in more detail in the following paragraphs.

Group Assignment. We randomly assigned the 119 websites to three groups: 40 in the control, 39 in the generic, and 40 in the personalized group.

Notification Mails. In terms of content, both emails were framed as a request from a German blog reader who argued being unhappy about the blog transferring personal data to the United States without consent, explaining that sensitive data is less protected in the U.S. ¹. This choice was made to (I) reduce possible biases due to the naming of certain academic institutions and (II) because [21] found no relevant differences between notifications from private individuals and academics. In addition, the emails listed alternatives to GA and concluded with a call to action.

This ethical argument was used for two reasons. First, it had to be plausible with the message framing as coming from an average blog reader. It would be unusual for such a person to show familiarity with data privacy law and the relevant court cases. Second, the ethical argument was used to reduce damage resulting from fear of legal prosecution as observed in a notification study by Princeton

¹One example why blog readers could be unsatisfied with the protection of their personal data in the US (during our study period) is the possibility for U.S. authorities to access the data of EU citizens if it is transferred to the U.S based on Section 702 of the FISA Amendment Act [30]. On 10 July 2023, the EU enacted the Data Privacy Framework [29] which allows companies to legally transfer data from the EU into the US under certain conditions. The data of the EU citizens must equally be protected as in the EU. This new framework did not influence our result, as we stopped scanning websites on 13 June 2023.

University and Radboud University. They employed a legal argument in combination with a deadline to fix the misconfiguration which made some website owners seek costly legal advice [22].

The emails sent to the personalized group were personalized using three techniques (Sect. 3): (1) For *identification*, the first name of the identified author or blog owner to address the recipient and the URL of the blog was used. (2) To *raise expectations of personalization*, the list of alternatives to GA was presented to be hand-picked for the blog. (3) *Contextualization* was achieved by summarizing the last blog post in the email. The summary was limited to keep the emails comparable in overall length. Both notification texts are provided in the supplementary material [16].

We generated two email accounts for the fictitious persons Laura and Tobias Busch² at *mailbox.org* to send the notification emails. For every blog, one of the two persons was selected at random to account for possible gender effects. The decision to pose as private individuals was made to facilitate personalization and was supported by the fact that an earlier study by Maass et. al. [21] found no significant differences between emails from private individuals and many other groups.

Deception. We deceived the recipients in our email, i. e., we did not include any information about the true purpose and our study. Instead, it looked as if the email had been written by an interested blog reader in their spare time. This deception was necessary to avoid biased reactions (observer effects) [2]. We followed best practices for deception in such a notification study [20]. A debriefing with the possibility of opting out was conducted. Ethical considerations will be discussed in more detail in Sect. 4.6.

Reminder Email. Our reminder asked if the subject had already had time to fix the issue as GA was still configured incorrectly. The personalized reminder only differed in form of address from the generic one. Both reminder texts are part of the supplementary material [16].

4.3 Notification Study Execution

The *initial notification emails* were sent on 12 April 2023 to the 79 blogs (to each blog from the personalized and generic group, no email was sent to the control group).

On 16 May 2023, 34 days after the initial notification, we scanned the blogs again and sent *reminder emails* to all blogs that were still setting GA cookies and had not responded to the notification email.

On 13 June, 62 days after the initial notification, we scanned the blogs the last time. These results determined the fix rate presented in Section 5.1. Since we deceived participants in our emails, we sent out *debriefing emails* on 17 July 2023 to all 119 blogs to inform them that they were part of this study. The blogs that were not part of the control group additionally received a link to the *survey*.

4.4 Communicating with Website Owners

We received several replies to our notification and reminder emails. We determined in advance how we would respond to certain type of answers, and outlined five cases and corresponding email answers to ensure comparability. Our prepared cases and responses are included in the supplementary material [16].

²The names were randomly chosen from a list of common German names [9].

4.5 Qualitative Analysis

To analyze the email and survey responses for our second research question, we documented the content of the answers.

For the content of the responses, we used open coding [7]. First, we created codes for the reasons mentioned not to fix, e. g. a code for stating that GA is already correctly configured on their website. Our codebook consisted of five codes (see [16]). Two researchers assigned independently the corresponding codes to the answers. Any difference in the coding of a conversation was then discussed until agreement was reached. For the open coding of the emails, we considered the complete conversation with a blog and not every single email. We additionally documented the author, the date, and the blog’s name of the email responses.

4.6 Ethical Considerations

None of the involved universities required IRB approval for this kind of study. Since we interacted with humans and partially deceived them about our true purposes, we designed the study following the principles described in the Menlo Report [3].

We *showed respect to our participants* by debriefing at the end of the study, clearly offering an opportunity to opt-out and by only publishing data in an anonymized form. Like in other notification studies, we deemed deception necessary to minimize social desirability biases and other observer effects [2].

Secondly, we considered *beneficence* by refraining from imposing time limits or threatening legal consequences (as opposed to e. g. [22]) to minimize the risk of (e. g. financial) harm to our subjects. We hope, our study can benefit the greater public by improving the personal data handling practices.

Thirdly, subjects were selected according to a predefined procedure with the aim of selecting the most visited websites for the respective search query, regardless of sensitive characteristics of the blog authors. In this way, the principle of *justice* is upheld.

Finally, we take into account the principle of *respect for the law and the public interest*. We have consulted legal experts on data protection before we conducted the study and we clearly describe our methods to ensure transparency and reproducibility.

5 RESULTS

As per request by one individual in the control group, we removed their blog from the study. Thus the control group shrunk from 40 to 39 blogs.

5.1 Effect of Personalization on Fixes (RQ1)

To analyze the effect of personalization, we consider the number of fixes in the group that received a personalized email. As Table 1 shows, seven out of 40 (about 18 %) fixed their configurations, which is a lower fix rate than for the generic email (21 %). Given the small sample sizes, we refrain from further statistical analyses.

In the context of notification campaigns about data protection issues, this suggests that personalization is neither useful nor necessary to increase the effectiveness of notifications.

For both groups, the fix rate is comparable to ones in previous studies where an ethical argument was used (e.g. [21]).

To rule out that the fix rate is affected by inherent characteristics of the websites rather than our notification, we referred to the




Figure 2: Distribution of coded reasons not to fix per treatment group. Given the small number of 14 responses in each group, no differences can be inferred for most categories. An interesting six-vs-zero difference can be observed in the “Claims that website does respect all privacy requirements” category.

Table 1: Fix rate and absolute number of fixes versus treatment. Both notified groups (Personalized and Generic) exhibit similar fix rates ($\approx 20\%$), which is higher than 10% in the Control group (no notification).

	Group Size	No. of Fixes	Fix Rate
Notified in total	79	15	19%
Personalized	40	7	18%
Generic	39	8	21%
Control	39	4	10%

additionally collected control variables per website (see Section 4.1): the presence of a webshop, the responsible party for the website (natural vs legal person) and activity in the past six months. Further, we alternated the sender’s gender in the emails. We could not find any patterns suggesting that any of these four characteristics had a sizable effect on the fix rate. This raises our confidence that our observations can indeed be attributed to our notification.

5.2 Communication and Survey Responses (RQ2)

To answer our second research question (“Which reasons not to fix do the notified subjects mention?”), we drew on two sources: Subjects’ responses during the study and a post-debrief survey.

Subjects’ responses to our emails were analyzed to extract reasons the website owners named not to fix the misconfiguration during the study. Out of the 40 blog operators receiving a personalized notification, 14 replied via email during the study. In the generic notification group, 14 out of 39 answered.

The reasons mentioned by the subjects in their responses were coded and compared (as described in Section 4.5) between the treatment groups in Figure 2. The reasons can be classified into three categories: denial of a data leak, a lack of resources to remedy the issue, and claims of specifically requiring GA.

No clear differences in reasons mentioned are evident for the categories with one exception: Six generically notified subjects denied the issue versus zero in the personalized group.

The short survey sent to subjects in the debriefing email included two optional questions. One free-text question asked for the reasons for not having fixed the misconfiguration, adapted to the action taken (questions in supplementary material [16]). The other free-text question asked for improvement ideas and further comments. The answers were coded as described in Section 4.5.

There were six fully completed survey responses from the personalized group and three from the generically notified group. Due to the small number of participants, especially from the generically notified, the survey responses will only be mentioned anecdotally.

Among the group of respondents that didn’t fix the misconfiguration, a lack of time (2 respondents), requiring further explanations (1 respondent) and that GA’s features were essential to monetarization (1 respondent) were given as reasons across both groups. One respondent in the personalized group suggested to further improve the email by adapting the salutation to the number of bloggers (different expressions in the German language) and by removing the mid-text appellation.

We stress that the survey expresses individuals’ voices that may not generalize and that could be biased by the observer effect.

5.3 Further Findings

The analysis in Section 5.2 focused on the reasons that subjects mentioned why they did not fix the misconfiguration. When coding subjects’ email communication, we noticed that subjects frequently made claims. We verified their correctness and made the following three interesting observations:

- Out of the six subjects announcing a fix (i. e. stating they want to remove GA or claiming they have already removed it) across both groups, only four followed up on that promise.
- Five subjects explicitly asserted that GA only became active upon visitor consent to the cookie dialogue. Our response dispelled this misconception and contained a link to the GA checker [21]. Yet, this did not result in a single fix.
- Finally, the data show an interesting relationship between remediation and response behavior. Out of the 15 subjects who fixed the misconfiguration, four replied via email at some point (27%). In contrast, out of the 64 subjects who

did not fix the misconfiguration, 24 answered (38%). In connection with the fact that few answers mentioned fixing the misconfiguration (Fig. 2), this suggests subjects rather responded than taking action.

6 DISCUSSION

In this section, we discuss our findings, further insights, and limitations of the study.

6.1 Effect of Personalization (RQ1)

Our first research question asks whether personalization can improve the effectiveness of notification studies for sport and fitness blogs. For the setup we examined, we can answer this question with *no*. Our personalization efforts did not result in higher fix rates (Sect. 5.1), nor did it have any obvious beneficial effects on the response rates (Sect. 5.2). Since collecting information for personalizing notifications is costly, given our results we cannot recommend employing personalization when notifying blogs about privacy leaks.

6.2 Communication and Survey (RQ2)

Secondly, we wanted to understand why subjects didn't fix the misconfiguration and whether different reasons were communicated to us by the two groups. The fact that the reasons raised across the two groups were quite similar (Sect. 5.2) supports our first finding. It's interesting, however, that we received six denials of the issue from the generic versus zero from the personalized group, which could indicate a higher level of trust into the personalized message.

We further found that claims of respecting privacy standards and promises to fix were true for a surprisingly small number of websites. We also observed that response rates were generally higher for the group of subjects who did not fix (Sect. 5.3). However, in view of the response topics, we cannot deduce that these higher response rates in the non-fix group can be attributed to technical difficulties or questions. Instead, our first observation of empty promises and rebuttals rather points in a different direction, as research by Monin and Miller suggests [25]. They have found that individuals are less likely to act upon certain moral standards after having demonstrated them in previous behavior, a phenomenon they refer to as "moral credentials". Based on their results, we can suppose that our notified subjects might be less likely to fix the misconfiguration after having demonstrated their willingness to take the sender seriously in their reply.

Our observation, therefore, suggests that the *opportunity to respond to the notification* may subconsciously demotivate some subjects from fixing the misconfiguration. On the other hand, an argument in favor of allowing responses is that clarification and guidance can be provided. However, recall that the clarification we provided did not result in a single fix (Sect. 5.3). In sum, our findings question whether the opportunity for subjects to reply to the notification has a net positive effect on the fix rate.

It would be interesting for future work to notify subjects without giving them the opportunity to respond. A basic approach would be to send notifications from a no-reply email address.

6.3 Limitations

A major limitation of this study is the relatively small sample size of roughly 120 blogs (including control group), which prohibits statistical analysis. However, this is very much by design, since personalization requires quite a lot of publicly available information in order to credibly present more familiar information to the recipient than simply their name. This and the fact that personalization needs to stay comparable between subjects for scientific analysis necessarily restricts the sample to quite a homogeneous group. It is one of our learnings that finding a meaningful one-fits-all strategy for personalization is not so easy.

Secondly, personalization was not always guaranteed, given that the addressed person may not be the one reading the email. Some responses show, that the notification can still have an impact in that case: The addressed person was recognized as a team member, leading to some identification of the recipient with the email.

7 CONCLUSION

This study investigated whether contacting operators of websites in a personalized way can help in convincing them to remedy user tracking prior to the consent banner. To this aim, we notified 79 operators of fitness blogs with such a GA configuration via email. 40 of the blogs received an email personalized with information from their website.

We find that personalizing the email did neither have a significant effect on the remediation rate nor an effect on the reasons not to fix as reported in the responses and survey.

While analyzing the responses, we made two further interesting observations. Firstly, we observed that the responses mostly center around false claims and empty promises to fix the misconfiguration. Secondly, the group of individuals that did not fix the misconfiguration had a higher response rate, which motivates further research exploring the role of the *moral credentials* phenomenon for fix rates.

To summarize, we find that personalization has no measurable impact and we identify the possibility to reply as a possibly disadvantageous factor through analyzing subjects' response behavior.

DATA AVAILABILITY

We released the texts of our sent emails, the answers to our survey, the open coding of the received emails (without the email content) and the control variables and final fix status for each blog [16]. Every personal information, such as blog names was removed and replaced with an individual code.

ACKNOWLEDGEMENTS

We would like to thank the Studienstiftung des deutschen Volkes which brought this research team together and supported and funded this research through their academic education and research program. This work was also funded by the Topic Engineering Secure Systems, subtopic 46.23.01 Methods for Engineering Secure Systems, of the Helmholtz Association (HGF) and supported by KASTEL Security Research Labs, Karlsruhe. Further, it was supported by the DAAD program Konrad Zuse Schools of Excellence in Artificial Intelligence, sponsored by the Federal Ministry of Education and Research.

REFERENCES

- [1] Austrian Data Protection Authority. [n. d.]. Datenschutzbeschwerde (Art. 77 Abs. 1 DSGVO) - Teilbeschleidspruch. https://noyb.eu/sites/default/files/2022-01/EDSB%20-%20Google%20Analytics_DE_bk_0.pdf. Accessed: 2024-03-11.
- [2] K. Baclawski. 2018. The Observer Effect. In *2018 IEEE Conference on Cognitive and Computational Aspects of Situation Management (CogSIMA)*. 83–89. <https://doi.org/10.1109/COGSIMA.2018.8423983>
- [3] M. Bailey, E. Kenneally, D. Maughan, and D. Dittrich. 2012. The Menlo Report. *IEEE Security & Privacy* 10, 02 (2012), 71–75. <https://www.computer.org/csdl/magazine/sp/2012/02/msp2012020071/13rRUXNEqNZ>
- [4] Orçun Çetin, Carlos Gañán, Lisette Altena, Samaneh Tajalizadehkhoob, and Michel Van Eeten. 2019. Tell me you fixed it: Evaluating vulnerability notifications via quarantine networks. In *2019 IEEE European Symposium on Security and Privacy (EuroS&P)*. IEEE, 326–339.
- [5] Orçun Cetin, Carlos Ganan, Maciej Korczynski, and Michel van Eeten. 2017. Make notifications great again: learning how to notify in the age of large-scale vulnerability scanning. In *Workshop on the Economics of Information Security (WEIS)*.
- [6] Orçun Cetin, Mohammad Hanif Jhaveri, Carlos Gañán, Michel van Eeten, and Tyler Moore. 2016. Understanding the role of sender reputation in abuse reporting and cleanup. *Journal of Cybersecurity* 2, 1 (2016), 83–98.
- [7] Juliet M. Corbin and Anselm Strauss. 1990. Grounded theory research: Procedures, canons, and evaluative criteria. *Qualitative Sociology* 13, 1 (1990), 3–21.
- [8] Arie Dijkstra. 2005. Working mechanisms of computer-tailored health education: evidence from smoking cessation. *Health education research* 20, 5 (2005), 527–539.
- [9] Gesellschaft für deutsche Sprache e.V. [n. d.]. Die beliebtesten Vornamen. <https://gfd.s.de/vornamen/beliebteste-vornamen/>. Accessed: 2023-09-24.
- [10] Google Ireland Limited. [n. d.]. [UA] Google Analytics 4 has replaced Universal Analytics. <https://support.google.com/analytics/answer/11583528>. Accessed: 2023-09-19.
- [11] Google Ireland Limited. 2023. Analytics Tools & Solutions for Your Buisness - Google Analytics. <https://marketingplatform.google.com/about/analytics/>. Accessed: 2023-09-19.
- [12] Google Ireland Limited. 2023. [GA4] Data collection Understand what Analytics collects through the default implementation. <https://support.google.com/analytics/answer/11593727?hl=en&sjid=7600686896757650253-EU>. Accessed: 2024-03-10.
- [13] Google Ireland Limited. 2024. [GA4] Cookie usage on websites. <https://support.google.com/analytics/answer/11397207>. Accessed: 2024-02-24.
- [14] Robert P Hawkins, Matthew Kreuter, Kenneth Resnicow, Martin Fishbein, and Arie Dijkstra. 2008. Understanding tailoring in communicating about health. *Health education research* 23, 3 (2008), 454–466.
- [15] IfD Allensbach. [n. d.]. Beliebteste Sportarten in Deutschland nach Interesse der Bevölkerung an dem Sport in den Jahren 2017 bis 2023. <https://de.statista.com/statistik/daten/studie/171072/umfrage/sportarten-fuer-die-besonderes-interesse-besteht/>. Accessed: 2023-09-24.
- [16] Theresa Kriecherbauer, Richard Schwank, Adrian Krauss, Konstantin Neureither, Lian Remme, Melanie Volkamer, and Dominik Herrmann. 2024. *Supplementary Material for "Is Personalization Worth It? Notifying Blogs about a Privacy Issue Resulting from Poorly Implemented Consent Banners"*. <https://doi.org/10.5281/zenodo.11165333>
- [17] Cong Li and Jiangmeng Liu. 2017. A name alone is not enough: A reexamination of web-based personalization effect. *Computers in Human Behavior* 72 (2017), 132–139.
- [18] Frank Li, Zakir Durumeric, Jakub Czym, Mohammad Karami, Michael Bailey, Damon McCoy, Stefan Savage, and Vern Paxson. 2016. You’ve got vulnerability: Exploring effective vulnerability notifications. In *25th USENIX Security Symposium (USENIX Security 16)*. 1033–1050.
- [19] Frank Li, Grant Ho, Eric Kuan, Yuan Niu, Lucas Ballard, Kurt Thomas, Elie Bursztein, and Vern Paxson. 2016. Remediating web hijacking: Notification effectiveness and webmaster comprehension. In *Proceedings of the 25th International Conference on World Wide Web*. 1009–1019.
- [20] Max Maaß, Henning Pridöhl, Dominik Herrmann, and Matthias Hollick. 2021. Best Practices for Notification Studies for Security and Privacy Issues on the Internet. In *Proceedings of the 16th International Conference on Availability, Reliability and Security*. ACM. <https://doi.org/10.1145/3465481.3470081>
- [21] Max Maass, Alina Stöver, Henning Pridöhl, Sebastian Bretthauer, Dominik Herrmann, Matthias Hollick, and Indra Spiecker. 2021. Effective Notification Campaigns on the Web: A Matter of Trust, Framing, and Support. In *30th USENIX Security Symposium (USENIX Security 21)*. USENIX Association, 2489–2506. <https://www.usenix.org/conference/usenixsecurity21/presentation/maass>
- [22] Mark Frauenfelder. [n. d.]. Princeton study tricked small websites into thinking they could be sued by a Russian organization. <https://boingboing.net/2021/12/23/princeton-study-tricked-small-websites-into-thinking-they-were-about-to-be-sued-by-a-russian-organization.html>. Accessed: 2023-09-20.
- [23] Ewa Maslowska, Bas van den Putte, and Edith G Smit. 2011. The effectiveness of personalized e-mail newsletters and the role of personal characteristics. *Cyberpsychology, Behavior, and Social Networking* 14, 12 (2011), 765–770.
- [24] Ewa Maslowska, Edith G Smit, and Bas Van den Putte. 2016. It is all in the name: A study of consumers’ responses to personalized communication. *Journal of Interactive Advertising* 16, 1 (2016), 74–85.
- [25] Benoit Monin and Dale T Miller. 2001. Moral credentials and the expression of prejudice. *Journal of personality and social psychology* 81, 1 (2001), 33.
- [26] Michael Mutz, Johannes Müller, and Anne K. Reimers. 2021. Use of Digital Media for Home-Based Sports Activities during the COVID-19 Pandemic: Results from the German SPOVID Survey. *International Journal of Environmental Research and Public Health* 18, 9 (2021). <https://www.mdpi.com/1660-4601/18/9/4409>
- [27] Ben Stock, Giancarlo Pellegrino, Frank Li, Michael Backes, and Christian Rossow. 2018. Didn’t you hear me?—Towards more successful web vulnerability notifications. (2018).
- [28] Ben Stock, Giancarlo Pellegrino, Christian Rossow, Martin Johns, and Michael Backes. 2016. Hey, you have a problem: On the feasibility of {Large-Scale} web vulnerability notification. In *25th USENIX Security Symposium (USENIX Security 16)*. 1015–1032.
- [29] U.S. Department of Commerce. 2023. Data Privacy Framework (DPF) Program Overview. <https://www.dataprivacyframework.gov/Program-Overview>. Accessed: 2024-02-16.
- [30] U.S. Government. [n. d.]. Foreign Intelligence Surveillance Act of 1978 Amendments Act of 2008. <https://www.govinfo.gov/content/pkg/BILLS-110hr6304pcs/html/BILLS-110hr6304pcs.htm>. Accessed: 2024-03-11.
- [31] Christine Utz, Matthias Michels, Martin Degeling, Ninja Marnau, and Ben Stock. 2023. Comparing Large-Scale Privacy and Security Notifications. *Proc. Priv. Enhancing Technol.* 2023 (2023), 173–193. <https://api.semanticscholar.org/CorpusID:258727012>
- [32] Eric Zeng, Frank Li, Emily Stark, Adrienne Porter Felt, and Parisa Tabriz. 2019. Fixing HTTPS misconfigurations at scale: An experiment with security notifications. (2019).