Comparing Large-Scale Privacy and Security Notifications

Christine Utz
CISPA Helmholtz Center for Information Security
Saarbrücken, Germany
christine.utz@cispa.de

Matthias Michels
Saarland University
Saarbrücken, Germany
s8mamich@stud.uni-saarland.de

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

Ninja Marnau
CISPA Helmholtz Center for Information Security
Saarbrücken, Germany
marnau@cispa.de

Ben Stock
CISPA Helmholtz Center for Information Security
Saarbrücken, Germany
stock@cispa.de

ABSTRACT

Over the last decade, web security research has used notification campaigns as a tool to help web operators fix security problems or stop infrastructure abuse. First attempts at applying this approach to privacy issues focused on single services or vendors. Hence, little is known if notifications can also raise awareness and encourage remediation of more complex, vendor-independent violations of privacy legislation at scale, such as informed consent to cookie usage under the EU's ePrivacy Directive or the General Data Protection Regulation's requirement for a privacy policy. It is also unclear how privacy notifications perform and are perceived compared to those about security vulnerabilities. To fill this research gap, we conduct a large-scale, automated email notification study with more than 115K websites we notify about lack of a privacy policy, use of third-party cookies without or before informed consent, and input forms for personal data that do not use HTTPS. We investigate the impact of warnings about fines and compare the results with security notifications to more than 40K domains about openly accessible Git repositories. Based on our measurements and interactions with operators through email and a survey, we find that notifications about privacy issues are not as well received as security notifications. They result in lower fix rates, less incentive to take immediate action, and more negative feedback. Specific reasons include a lack of awareness and knowledge of privacy laws' applicability, difficulties to pinpoint the problem, and limited intrinsic motivation.

KEYWORDS

web privacy, notification study, GDPR, ePrivacy, cookie consent

1 INTRODUCTION

Websites' availability and security depend on operators following best practices, update their systems, and stay alert of new threats. Over the last few years, compliance with privacy regulations has become another important task to ensure that services operate within the legal boundaries and protect user privacy. Recent years have seen the creation of new privacy laws across the globe to better tackle the 21st century's reality of ubiquitous processing of personal data by digital services. New privacy regulations including the European Union's General Data Protection Regulation (GDPR) and the California Consumer Privacy Act (CCPA) have introduced extensive requirements for the processing of personal data and granted people individual rights to learn about and control the use of their personal information. As research has shown, the adoption of new regulations and guidelines in practice is often slow [10, 58]. While the last few years have seen increasing numbers of GDPR-induced fines [8], a lack of monetary and human resources continues to pose a problem in large-scale enforcement of privacy laws [25, 51]. Given that web privacy research has been identifying privacy issues on websites at scale for years, this raises the question if and how the scientific community could aid regulators in identifying and remediating website behavior that violates privacy law.

One promising means for privacy and security research to help boost GDPR compliance on the Web are large-scale email notification campaigns. Informing the operators of affected websites could help raise awareness of practices that do not comply with privacy law and encourage operators to fix the issue before they are subject to GDPR-mandated fines. Such notifications have been repeatedly used by security research to raise awareness and motivate fixes of diverse issues including Heartbleed [12], XSS [49, 50], DDoS amplifiers [34, 61] and potential information leaks [35, 36]. If this approach turns out to also be viable for notifications about privacy violations, this could take the burden off data protection authorities (DPAs) in enforcing laws and help website owners to better protect user privacy. While both researchers and NGOs have conducted notification campaigns about privacy issues before, they have focused on selected Consent Management Platforms (CMPs) [44] or restricted their scope to a single vendor and locale [38], typically because of the manual verification involved. It is also unknown how notifications about privacy problems compare to those about security vulnerabilities in terms of remediation rates and timing.

In this work we explore the feasibility of large-scale, automated email notification campaigns for vendor-independent violations of privacy laws, namely the GDPR's transparency requirement, its mandate to use state-of-the-art data protection mechanisms, and consent to the use of not strictly necessary cookies under the ePrivacy Directive. To identify how notifications about privacy issues perform compared to those about security vulnerabilities,

Proceedings on Privacy Enhancing Technologies YYYY(X), 1–21

https://doi.org/XXXXXXX.XXXXXXX

we also notify websites about publicly accessible Git repositories that may leak sensitive information. We compare fix rates between issues, investigate the impact of mentioning potential fines, and conduct qualitative analyses of feedback from emails and a survey to learn how privacy notifications are perceived by recipients and what could be done to help them address privacy issues in the future. More concretely, we make the following contributions:

- We conduct the first large-scale, automated email notification study with 115K websites that investigates the feasibility of this approach for complex, vendor-independent privacy issues. Our notifications have significant impact on remediation for lack of a privacy policy or consent notice. We find no significant impact of warnings about potential fines.
- We compare the effect of notifications about privacy issues with those about a security vulnerability. For privacy, fix rates are lower than for the security vulnerability, which is also addressed more quickly. Recipients also perceive emails about a privacy compliance issue more negatively, partially due to a lack of intrinsic motivation to fix it or (incorrect) assumption of inapplicability of the relevant laws.
- As for the persisting reachability problem, we investigate if email addresses extracted from websites are an efficient and scalable alternative to prior approaches, manually collected addresses or email generics. Our results confirm this, with 87.8 % and 33.8 % of successful handovers to recipients' mail servers for extracted addresses and generics, respectively.

2 RELATED WORK

Security Notifications. In Web security research, large-scale notification campaigns were first used to alert server operators about abuse of their infrastructure for unintended purposes, including distribution of malicious downloads [6, 59]. This approach was subsequently also applied to raise awareness and motivate fixes of security vulnerabilities including Heartbleed [12], DNS zone poisoning [7], XSS [50], HTTPS misconfigurations [61], DDoS amplifiers [34, 35], misconfigured IPv6 firewalls [35], and leaks of information whose public accessibility could pose a security risk, including industrial control systems [35], Git or SVN repositories [36, 49], cryptographic keys, database backups, server status information, and `phpinfo` files [36].

One core problem with large-scale Web security notifications is to reliably reach the people responsible for fixing the issue. While previous work found that more individual communication channels such as telephone [49], physical mail [36, 38, 49], websites' contact forms, associated social media accounts [49], or manually identified contact email addresses [36, 38, 49] can lead to higher delivery rates, the involved overhead in terms of human resources and monetary cost makes these infeasible for notifying websites at scale. Hence, most Web security notification campaigns have used generic approaches to contact websites via email, either directly via generic email addresses [6] such as `info@DOMAIN` or `webmaster@DOMAIN` (RFC 2142 [9]), WHOIS contact information [6, 7, 12, 35, 49, 50, 61] or through trusted third parties including CERTs [35, 50], DNS nameserver operators [7], or hosting providers [6], depending on the investigated issue(s). Drawbacks of this approach include high bounce rates due to missing or outdated information in WHOIS records or non-use of RFC 2142 mailbox names [7, 50]. Use of intermediaries carries the risk of them not forwarding notifications [35].

Even if a notification email is correctly targeted, it still risks being considered spam or otherwise malicious by both mail servers and human recipients. Prior work has studied how to increase perceived message authenticity in notification campaigns by evaluating the effect of sender reputation [6, 38, 61], email format such as plaintext or HTML [49], text localization [35, 61], and use of S/MIME [49], but no clear "recipe" has emerged. Finally, findings also differ for the content of the notification message itself: While some studies did not identify a significant influence of message text on remediation rates [7, 61], others found that more detailed explanations had a significant positive influence [35, 59]. Message tone was found to not affect fix rates [49, 61]. Maass et al. compared existing work and outlined practical recommendations for future notification studies [37].

Privacy Notifications. Compared to abuse and security notifications, previous work that notified website owners about privacy issues at scale is scarce. Challenges lie in such a study requiring 1) determining if the examined website is subject to the privacy legislation of interest, 2) certainty that a given issue is regarded a violation of this privacy legislation, and 3) a high-accuracy detection mechanism to keep the number of false positives low and not cause unnecessary anxiety and costly investigations with people whose websites do not have a privacy problem. The third prerequisite is challenging, as privacy issues are hard to detect automatically [46] unless focused on specific services or vendors, where a common implementation can allow for simple yes/no checks of a value, parameter, or URL. Otherwise, complex heuristics are necessary that may require, for example, contextual analysis and natural language processing to determine if a privacy notice contains the required disclosures.

Consequently, prior privacy notification campaigns focused on specific vendors or consent frameworks. Maass et al. [38] notified the owners of 4,754 German websites about the lack of IP anonymization in their Google Analytics integration, which the German DPA deemed necessary for GDPR compliance [33] and can be remotely detected with certainty through URL parameters passed in the HTTP request to Google. While the study found that framing notifications as legal compliance issues led to increased fix rates, the analysis is limited to German sites and those with an imprint, thereby limiting insights to this selected group. Our work uses a much larger domain set without this bias. In May 2021, Austrian privacy NGO noyb notified more than 500 companies about consent notices on their websites that used techniques considered to be non-GDPR-compliant by various national DPAs [44]. While 42 % of individual violations were fixed within 30 days, 82 % of companies still exhibited some type of GDPR compliance issue. Since the goal of the notification campaign was to file complaints to national DPAs, the automated detection mechanism was tailored to a single CMP, OneTrust, and supported by manual review by legal experts, limiting this approach in coverage and scalability.

Automated Detection of Web Privacy Issues. Automatically detecting vendor-agnostic privacy problems on websites is challenging due to a lack of standardization and concrete guidelines by lawmakers, DPAs, and court rulings on how to implement key legal requirements on a technical level, including transparency mechanisms

mandated by privacy law such as privacy policies or consent notices. This is partly a deliberate decision to remain flexible towards future technological developments [46]. But even for concrete requirements, such as the wording of the "Do Not Sell My Personal Information" link mandated by the CCPA [48], actual implementations on websites widely vary [45, 58].

Still, there is previous work that worked towards automatic detection of the privacy issues at the heart of our study. A growing body of literature has tackled the problem to automatically find privacy policies on websites and download them for further analysis. Hosseini et al. [28] discuss and evaluate different approaches and identify best practices. Recent research has also shown growing interest in (cookie) consent notices. They originate in Article 15(3) of the EU's ePrivacy Directive, which requires website visitors' consent to store information on their devices unless the site cannot technically function otherwise, but are also used to obtain informed consent to data processing under Article 6(1)(a) GDPR. As with privacy policies, differences in implementation make automatic detection difficult [46]. Thus, past work has focused on specific consent frameworks or individual CMPs [4, 26, 39, 43, 44] or included manual analysis [10, 56]. Despite EU law requirements for free, prior and informed consent to data processing [3, 46], many consent notices were found not to offer sufficient choice [10, 56], use dark patterns to nudge people into giving consent [43, 56], or do not have a backend that ensures the visitors' selection is honored by the website [4, 39]. Bollinger et al. [4] trained a machine learning classifier on cookie–purpose mappings from CMP classifications and manual categorization by web developers. Examining the 30K websites from the Tranco 1M list that featured one of the investigated CMPs with cookie–purpose mappings, they found that 94.7 % exhibited at least one violation of a consent requirement. Like prior approaches to determine the purpose for which specific cookies are set [29], this approach suffers from a lack of reliable ground truth.

Our work builds upon some of these techniques to automatically detect privacy issues at scale and independent of software or vendor, focusing on keeping the number of false positives low to avoid unnecessary notification of compliant websites.

3 MEASUREMENT & NOTIFICATION SETUP

Our study setup first required us to identify the security and privacy issues we wanted to notify operators about and how to check their presence at scale. Further, we describe how we created a set of domains to monitor, notification messages, and report infrastructure. We also explain ethical aspects of our large-scale measurements and notifications and discuss limitations of this approach.

3.1 Investigated Issues & Implemented Checks

In Section 2 we already identified three core requirements for privacy issues to investigate in a large-scale notification study. Of particular importance is a low number of false positive cases to not erroneously alert recipients and potentially trigger costly and stressful investigations. Thus, for privacy notifications, we required a clear violation of unambiguous regulatory data protection requirements that can be automatically detected. This ruled out issues requiring human judgment, such as dark patterns. To estimate the prevalence of false positives for our checks, we manually verified, for each issue, whether it was indeed present on 250 websites randomly drawn from the set of all domains we found to have the respective problem. We ended up selecting four privacy issues which fit our requirements and implemented them as custom functions in an established measurement framework, OpenWPM [14]. To compare the effect of privacy notifications to those about a security vulnerability, we also selected one security issue already used in prior notification studies, publicly accessible Git repositories [36, 49]. For performance reasons, checks for the Git issue were not conducted with OpenWPM but with standard HTTP requests. All checks were performed once a day, launched shortly after midnight CET from CISPA servers on the premises of Saarland University in Germany. Performing the checks with an IP address in the EU is important because some websites, particularly with `.com` domains, show consent notices only to EU-based visitors [57].

3.1.1 No Privacy Policy. The GDPR requires that data subjects are informed about all processing of their personal data, which also comprises communications data such as users' IP addresses in web server logs [15], even if stored only temporarily. Thus, any website collecting such information must have a privacy policy that explains the use of visitors' personal data. To determine if a website had a privacy policy, we followed best practices identified by Hosseini et al. [28] and searched for privacy policy specific words in and around HTML link tags. For this, we extended a list of common words for privacy policy links, terms-of-service, and contact pages from a recent study [55] to cover all official EU languages. After a website had been fully loaded, we used the list of words identifying privacy policies to find links that likely lead to a privacy policy. If no such link was found, we also visited less privacy-specific subpages like terms-of-service and contact pages and searched them for words from the privacy list. If neither of these searches led to a match, the site was marked as violating the privacy policy requirement. For the manually checked sample of 250 sites drawn from those with this violation, 0.4 % were false positives.

*3.1.2 Use of Third-Party Cookies Without Consent Notice (*No Consent*) or Before Interaction With Consent Notice (*Before Consent*).* The setting of cookies is regulated by the EU's ePrivacy Directive (2009/136/EC) [54] and its implementations into national laws. Under its Article 5(3) the storing of information in a user's terminal equipment, including HTTP cookies that are not *strictly necessary* for the functioning of the website, is only allowed if the user has given prior, active consent. Mere continued use of the website does not constitute informed consent [16, 17]. For the *No Consent* and *Before Consent* issues we focused on cookies set by third-party providers for advertising, analytics, and social media, because EU DPAs have universally deemed these non-essential for the website's functioning [1, 32]. We used the WhoTracks.me database [23] to categorize a checked website's third-party requests by purpose and flagged those that included a `Set-Cookie` HTTP header and requested a third-party domain classified as "audio video player," "ad/pornvertising," "site analytics," or "social media". The presence of a cookie consent notice was determined based on two rule sets, a list of common HTML elements from the EasyList Cookie List [53] and the list of consent management providers vetted by IAB Europe's CMP Compliance Programme [30]. If one of the EasyList

rules matched or a script referring to one of the CMPs was found on the front page, we assumed that the site had a consent notice. If a website issued a third-party request that required prior consent but a consent notice was not detected, we considered the site a case of *No Consent*. If a consent notice was detected but the flagged request was issued despite our script not interacting with the website, the website was labeled as having the *Before Consent* issue. The prevalence of false positive cases for *No Consent* was 2 % in our manual check and 6.8 % for *Before Consent*. The latter involved a tradeoff between false negatives for the presence of a consent notice and false positives for a notice without a working consent mechanism.

*3.1.3 Input Fields for Personal Information Without HTTPS (*No HTTPS*).* The GDPR's requirements for "security of processing" (Article 32) and "data protection by design and by default" (Article 25) mandate the use of appropriate state-of-the-art technology for the collection and processing of personal data. One such mechanism is transport encryption of HTTP traffic via TLS, i. e., HTTPS. Better availability of certificates and browsers flagging HTTP-only connections as insecure have led to increased adoption, so that the majority of today's Web traffic uses HTTPS [10, 18]. Thus, it is considered a state-of-the-art technology to protect personal information [31, 52]. To detect if a website requested users' personal data without securing it with HTTPS, we created a list of terms for personal information (e. g., `firstname` or `password`) likely to be used as names for input fields that request the corresponding piece of information. In an iterative process we checked our list against the actual names of form fields used by websites, removed terms that led to many false positives, and added newly found, more specific terms (e. g., `login_email`). We ended up with a final list of 24 phrases (see Appendix A) that is not comprehensive but designed to reduce false positives. We flagged a site as violating the HTTPS requirement if one of the terms on the list was used in the `name` or `id` attribute of HTML input fields and the site did not use HTTPS. Manual validation yielded a prevalence of false positive cases of 3.6 % on the 250 sampled sites with this issue.

*3.1.4 Publicly Accessible Git Repository (*Git*).* If repositories for software version control systems such as Git or SVN are accidentally publicly accessible, they could potentially leak confidential information to outsiders, such as hardcoded encryption keys or credentials [49]. Considered a security vulnerability, this issue was already the subject of previous security notification campaigns [36, 49]. We selected it as a security issue to compare against our privacy notifications because it is still a common problem, can be accurately detected, and, like the privacy issues, concerns a specific *domain* rather than a specific server (that could host multiple domains). Most importantly, this issue can be tested in a non-intrusive manner, which is a core aspect of ethical security vulnerability checks [49, 50] and excludes any vulnerability for which even a proof-of-concept would require some server-side code execution, which could be considered illegal in some countries. To check websites for publicly accessible Git repositories, we used standard HTTP requests, as they were faster and less resource intensive than OpenWPM. We tried to access the file `domain.tld/.git/config`; if it contained the line `[core]`, we requested `.git/HEAD`. This either directly provided the hash of the currently checked out commit or pointed to a branch, so we could retrieve that branch's commit hash from `.git/refs/heads/<branch>`. If the commit hash could also be found on GitHub, we did not consider the domain problematic, assuming that a repository also published elsewhere does not increase the attack surface [49]. Our check did not further investigate if the repository indeed posed a security risk, because once the presence of a publicly accessible repository has been confirmed, it would be unethical to search its content for sensitive information.

3.2 Initial Domain Set

In order to obtain an a large and diverse initial set of websites to analyze, we leveraged a public domain list provided by the TheInternetBackup project [5], whose goal was to compile a list of every domain on the Internet. Our starting point was a domain list with 252 million domains from February 2020. To reduce the number of sites subject to resource-intensive checks, we defined additional criteria for a website to be a candidate for a detailed check:

- **EU-based**: To ensure that EU privacy laws applied to the monitored websites, we first resolved the domain names and checked that all requested IP addresses were within the EU, based on Maxmind's GeoIP database [40].
- **Not parked**: Next, we excluded domains for which the resolving nameserver was a known domain parking service. We identified these by manually extending the list by Vissers et al. [60]. These DNS-based checks reduced the number of candidate sites to 51 million.
- **Active web server**: We issued HTTP requests to all remaining domains to check whether they provide a website. If the HTTP response status was below 400, we kept the domain in our data set, leading to around 30 million candidate sites.
- **No previous opt-out:** We excluded 1,513 websites that had opted out of our previous notification studies [49, 50].
- **Public audience**: We excluded sites that only offered limited content or did not seem to be targeted at a public audience (e. g., "under construction" sites). As a metric we required at least five same-site links on the front page.

The check for internal links was part of a pre-study in which we visited all 30 million sites with our OpenWPM-based check infrastructure. It took three months to visit each domain once with our automated setup and check it for the presence of the four privacy issues. Overall we found 6,272,813 candidate sites (~21 %) with at least one privacy issue. Cases were not evenly distributed; most common was *No Privacy Policy* (17.44 %), followed by *Before Consent* (7.57 %) and *No Consent* (7.34 %). *No HTTPS* was rarest, with 2.85 % of sites. Checking all of these domains daily, let alone notifying all of them would have been infeasible given hardware restrictions, so we sampled 500,000 domains from the list of domains with at least one problem. Then, for each issue, we sampled up to 100,000 domains; since only about 1 in 10 problematic domains had *No HTTPS*, we only drew around 45,000 domains for this issue. This left us with 331,222 domains with privacy issues subject to further monitoring. Note that while this sampling was done in late September 2021, the set we sampled from contained all domains that had been identified as problematic once within the three preceding months. In addition, we found 58,715 domains with the *Git* issue, which we also added to the set of domains to check each day. In total, this yielded 388,825 domains for further consideration.

3.3 Notification Emails and Infrastructure

The notification process itself involved determining the email addresses to contact, the mail server setup, composition of the notification emails, and setting up a website that allowed participants to check the status of their website and learn about our study.

3.3.1 Contacted Email Addresses. To identify points of contact with the monitored websites, we investigated a potential alternative to manually identified or generic email addresses: automatically finding email addresses on websites. As part of our daily OpenWPM checks, we searched the websites for email addresses likely to belong to people involved in the website's technical or legal administration. For this, we identified links to privacy policy pages as described in Section 3.1. Using a regular expression, we searched the HTML code of these subpages for email addresses. If none were found on a privacy policy page, we extended the search to subpages expected to contain generic contact information, such as "About" or "Contact us" pages. To remove false positives (e. g., file names containing '@') and to prevent emailing someone unrelated to the domain, we used only addresses with matching domains. If this procedure yielded at least one email address for the inspected domain, we emailed up to three discovered email addresses in the order served by our database and flagged the domain as being notified through (a) *Parsed* address(es). If no email address meeting the above criteria could be found, the domain was flagged as *Generic* and we sent our notifications to three generic aliases (RFC 2142 [9]): `info@DOMAIN`, the most frequently found email address in the first step, plus `webmaster@DOMAIN`, and `abuse@DOMAIN`, the two most commonly used email generics according to Soussi et al. [47].

3.3.2 Mailserver Setup. To send notification emails, we used a designated server outside CISPA, i. e., hosted with an external server provider. This reduced the risk that our notifications negatively impacted our institution's normal email communication (e. g., in case we hit a spam trap). Both A and MX record of our subdomain `notify.cispa.de` point to this server. This subdomain was also used in the EHLO message. The server configuration followed best practices to increase the delivery rate, including SPF and DMARC records and DKIM signatures for outgoing emails. The policy in the DMARC record was set to 'none,' the percentage to 100 %, and the address for aggregated reports to `administration@notify.cispa.de`. We also configured the reverse DNS to point to `notify.cispa.de` to create another clear connection to our institution and S/MIME-signed all emails to enable validation by the receivers' email software. Finally, to reduce strain on receiving servers, we set the rate of delivery to our MX to at most one email every two seconds.

3.3.3 Notification Emails. In our emails we openly identified ourselves as researchers and their purpose as being a scientific study. The sender name was composed of the name and institution of the author responsible for the notification setup. Mails were sent from a designated email address, `notify@notify.cispa.de`. The emails' subject line was "[Security and] data protection issue[s] on your website [DOMAIN]" or "Security issue on your website [DOMAIN]", depending on the type of detected issues. Following prior findings that language did not significantly influence fix rates [61] and localization of notification messages may even increase the likelihood that recipients perceive them as malicious [35], all emails were sent in English. The message body introduced our project, the involved research groups and institutions, and the security and/or privacy issues identified on the respective website. We provided a description of the problem(s) and why they constituted a violation of privacy law or a potential security vulnerability. Appendix B contains an example notification email with all possible issues.

3.3.4 Report Website. To aid operators in fixing their websites, we followed prior work [36, 50] in providing a web interface that allowed them to track their website's status with regard to the investigated issues. Every email contained a link to a domain-specific online report, which again listed and described the issues found on the respective website, but was updated daily with the most recent check results. This allowed operators to learn if our checks still detected an issue or considered it fixed. To prevent incorrect feedback due to a flaky check, an issue's state was only reported as fixed if this was supported by the latest two checks. The online report also provided operators with the option to exclude their website from our checks and, for each detected issue, a form to report false positives. A screenshot of an example report is shown in Appendix D. The website serving the online reports was hosted at CISPA, from where the daily checks were conducted. It also contained an introduction to our research project (see Appendix C), an imprint, and a privacy policy explaining our data processing.

3.4 Research Ethics

To ensure our research followed ethical best practices, we requested approval from CISPA's IRB. We outlined that our measurement setup would not collect personal data beyond what was publicly available, i. e., email addresses found on websites or generic aliases. Beyond that we followed data minimization principles: Survey answers (see Section 5.1) were anonymized and did not contain any information that allowed us to identify the website or email address used for notification. The only information passed to the survey via URL parameters were the issues found on the website, notification round, email type, and study condition; see Sections 4.1 and 4.1.2. We received IRB approval without changes to the study protocol.

In addition, we followed best practices recommended by prior work for ethical network checks [13] and notification studies [37]. We communicated our identities and benign intentions at all points of contact with the monitored websites and their operators: In all notification emails and on the study website we identified ourselves as researchers and explained the purpose and scope of our checks and the whole research project. For the daily checks, we set the user agent of our OpenWPM crawler to "CISPA Web Analyzer (notify.cispa.de)" to point operators of the checked websites to our study website. Front office and IT staff at CISPA were briefed about the study, preparing them for operators potentially asking about the legitimacy of the study. Notified websites could use the opt-out functionality on the website with their report (see Appendix D) or send an email to be excluded from future checks. Web report accesses were collected in the report web application and disclosed in its privacy policy, drafted by our expert in data protection law. At the end of the study, we sent debriefing emails to still affected websites in the *Control* group (see Section 4.1), informing them about the detected issues and our study.

To ensure that the selected privacy issues were universally acknowledged violations of privacy law, the study was supervised by a legal expert with extensive knowledge of EU data protection law. In addition, we manually verified check results to ensure an as low as possible false positive (FP) rate. Still, sending email notifications for complex privacy issues at scale meant that we inevitably reached out to some domains that were FP cases for a privacy issue. When notification recipients emailed us about (presumed) FPs, we performed a timely manual inspection of their website and responded with the result to minimize recipients' time of uncertainty about the issue. We also routinely checked for FPs on domains whose operators had emailed us with an unrelated question. 75 out of 414 email conversations with recipients of privacy notifications concerned (presumed) false positives. On 33 of these 414 domains we manually found true FPs. The rest were presumed false positives, reasons for which we explore in Section 6.3.3. We assume that most recipients of a notification caused by an FP contacted us before investing a significant amount of time in the issue. Nearly 50 % of true FPs could also be quickly identified by non-experts, such as websites actually having a privacy policy or not targeting people in the EU. Thus, we believe that we did not cause undue burden on notification recipients and the benefit for the other notified operators outweighed the potential cost for the few true FPs.

3.5 Limitations

We defined technical constraints for privacy issues in collaboration with a legal expert, but since there are multiple steps involved that relied on external sources (e. g., for geolocation of IP addresses or classification of third-party requests), we can only aim to minimize, but not eliminate false positives. Our study design also did not focus on avoidance of false negatives, so we likely missed many websites that in fact did have privacy issues. For example, if a site provided a link to a privacy policy, we did not further investigate if that page actually contained the required disclosures. Due to use in a pre-study, we removed `.de` domains affected by the *Git* issue from the set of domains. However, we do not believe this had any effect on our findings. We openly communicated that the notification process was part of a scientific study, possibly prompting fixes that would not have taken place otherwise due to the observer effect [37].

4 MEASURED NOTIFICATION RESULTS

Evaluating the measurements first requires us to describe the final parameters we used when launching the notification campaign. After that, we present our results regarding website reachability, web interface usage, remediation rates, and the influence of warnings.

4.1 Final Study Parameters

4.1.1 Notified Domains. From the 388,825 domains initially considered, only 190,491 were still problematic when we started to send out notifications on October 20, 2021. The remainder was either fixed without our notification or could no longer be reached. To test our infrastructure, we sent out emails to 19,142 domains, which we removed from further consideration in our study. In this beta test, we noticed and fixed some minor issues. The full notification campaign started on November 3, 2021 and considered all 159,856 domains which were still flagged as problematic on November 1.

Table 1: Number of domains (*n*) and prevalence of issues within each study condition and email group. Websites could have multiple issues.

	n	No Priv. Policy	No Consent	Before Consent	No HTTPS	Git
Control	31,863	14,472	5,451	5,293	3,827	8,480
Warning	63,596	28,674	10,994	10,790	7,803	16,626
No Warning	63,576	28,750	10,832	10,700	7,654	16,816
Parsed	63,675	25,896	13,281	15,987	8,271	11,224
Generic	95,360	46,000	13,996	10,796	11,013	30,698
Total	159,035	71,896	27,277	26,783	19,284	41,922

4.1.2 Study Conditions. Beyond the main focus of this study we wanted to explore whether warning about potential consequences of persisting issues (e. g., fines under the GDPR or national laws implementing the ePrivacy Directive), so far only investigated in a context limited in scope [38], had any effect on notification success. For this, we divided the domain set into three groups, each of which we assigned one study condition: a *Control* group (20 %); a group which received a *Warning* (40 %) about potential fines; and one that only received information about the issues but *No Warning* (40 %). The concrete wording of the warnings can be found in the example notification in Appendix B. In our analyses we also differentiate domains by *email type*, i. e., whether we contacted them via a *Parsed* or a *Generic* email address. We do not consider these "true" experimental conditions as we did not have any influence on whether we were able to find an email address on a website.

4.1.3 Schedule. We monitored the websites in the above data set over the course of two months, November and December 2021. Between November 3 and 5, 2021, we sent *initial notifications* to 297,506 email addresses associated with the 127,172 domains in the *Warning* and *No Warning* groups that still had the originally detected security and/or privacy issues as of November 1. Between November 20 and 22, 2021 we sent *reminder* emails to websites on which we could still detect the initial issue(s) as of November 18. Excluding bounced emails, domains for which the report had been viewed, and opt-outs, reminders for 62,835 domains were sent to 98,079 addresses. We did not filter report visits for automated accesses, e. g., by URL scanners in email verification systems, but did not see spikes in access rates in early November that would have indicated many automated accesses. As recommended by prior work [37], we sent *debriefing* emails to 67,194 addresses for the 28,724 domains in the *Control* group on December 20, about seven weeks after the notification of experimental groups. The message text matched the initial notification for the *No Warning* group, including links to the info website, report, and survey. Overall, we contacted 47,574 domains in the *Parsed* group via one email address, 8,122 via two and 7,189 via three. All 93,832 domains in the *Generic* group were emailed via three generic addresses (see Section 3.3.1).

4.1.4 Opt-outs & Final Domain Set. Our initial set of notified domains comprised 159,856 domains. We received a total of 497 opt-out requests, 466 via the web interface and 31 by email. We also excluded 44 domains identified as false positives in email conversations (see Section 5). After removing another 280 domains that resolved to a domain parking service, we were left with a final

data set of 159,035 domains. Table 1 shows these domains by study condition, email type, and prevalence of identified issues.

4.2 Reachability

As expected with generic and automatically extracted email addresses, not all emails reached their recipients. Only for 70,542 (55.5 %) of initially notified domains at least one email was successfully delivered, which we assumed if our mail server was able to hand over the email to the recipient's mail server. While such a successful handover does not mean that an email will reach the recipient's mailbox, this metric still provides an upper bound for the notification delivery rate. The difference between *Parsed* and *Generic* email addresses was quite high: While 87.8 % of initial notifications to *Parsed* addresses were successfully delivered, this was only the case for 33.8 % of emails to *Generic* addresses.

4.3 Web Interface Usage Statistics

As described in Section 3.3.4, each notification email contained a link to our study website with more information about the project and a personalized report for each domain that also let participants report false positives or opt out of the study. Over the course of the study, 5,731 reports were viewed, 259 false positive checks were reported, and for 466 domains the web interface was used to opt out of daily checks. The number of opt-outs differed between issues: While 0.4 % (483) of the domains with a privacy issue opted out of our checks, only 0.1 % (58) of the domains with the *Git* issue requested their exclusion via the web interface. 75.2 % of report views occurred within 24 hours, indicating that the vast majority of recipients either reacted promptly or not at all.

4.4 Remediation Rates

To determine if our notifications had any measurable effect, we checked the monitored websites daily over the course of two months.

4.4.1 Sliding Window Approach. To avoid domains incorrectly flagged as fixed because of one-off checker timeouts or page maintenance, we implemented a sliding window approach to determine if an issue persisted: For a given day t_i, we considered a domain d to be `problematic` if at any point in time within a 7-day window (t_i, t_{i+6}) our checker had identified the reported issue to still be present on the site. This 7-day window allowed us to obtain at least one real measurement result (i. e., a true/false evaluation of the checked issue, not a returned error or a missing data point) for 98.5 % of evaluated windows, resulting in a robust check.

4.4.2 Evolution of Problematic Domains Over Time. Figure 1 (a)–(e) shows for each issue, experimental condition, and email group the percentage of domains considered `problematic` at a given point in time. We investigated the persistence of issues more closely at four distinct points in time: one and two weeks after both initial notifications and reminders. Table 2 in its % column lists `problematic` rates for November 10, 2021 and Appendix F for all four dates. Overall, we observe for a given privacy issue a similar downward trend in `problematic` domains across study conditions and email groups, with them mostly differing by only between 0–1 percentage points, though the *Control* group behaves as expected in yielding the highest rates of persisting issues. *Git* rates also follow this pattern but yield slightly higher differences to *Control*, in the dimension of 1–2 percentage points. This lack of a larger measurable effect is a direct result of our notifications' low delivery rates: With many domains in the treatment conditions never receiving a notification email, this subset is expected to behave similarly to the control group, exhibiting the same rates of natural decay of issues.

Table 2: Percentages of websites still considered `problematic` with regard to each specific issue according to our sliding window evaluation (see Section 4.4.1) on November 10, 2021, by study condition / email type. % indicates the percentage of still `problematic` domains, diff the difference in percentage points to the *Control* group, and p the p-values for Fisher's exact tests ($\alpha = 0.05$) compared to *Control*. Bold values indicate significance after Holm-Bonferroni correction.

	%	diff	p		%	diff	p
No Privacy Policy				*No HTTPS*			
Warning	98.19	-0.46	**0.0004**	Warning	97.42	-0.54	0.08151
No Warning	98.41	-0.24	0.0601	No Warning	97.75	-0.21	0.498
Parsed	98.22	-0.42	**0.0000**	Parsed	97.57	-0.39	0.08381
Generic	98.45	-0.19	0.0830	Generic	97.73	-0.23	0.4293
Control	98.65	–	–	Control	97.96	–	–
No Consent				*Git*			
Warning	96.15	-1.41	**0.0000**	Warning	91.18	-1.60	**0.0000**
No Warning	96.74	-0.82	**0.0039**	No Warning	91.08	-1.70	**0.0000**
Parsed	96.38	-1.18	**0.0000**	Parsed	94.04	1.26	0.1438
Generic	96.94	-0.62	**0.0138**	Generic	90.52	-2.26	**0.0000**
Control	97.56	–	–	Control	92.78	–	–
Before Consent							
Warning	97.59	-0.31	0.2194				
No Warning	97.83	-0.07	0.8166				
Parsed	97.70	-0.20	0.3509				
Generic	97.81	-0.09	0.6328				
Control	97.90	–	–				

For the two issues related to the use of third-party cookies, *No Consent* and *Before Consent*, the number of `problematic` domains steeply dropped between November 15 and 16, 2021. Inspecting the affected domains, we found the cause to be a Twitter cookie named `lang`, originating from `cdn.syndication.twimg.com`, that was no longer present from November 16. This coincides with major platform updates at Twitter, including migration to Twitter APIv2[1].

Figure 1 (f) compares the evolution of `problematic` domains by issue aggregated across all experimental conditions. Looking at that figure and the rates in Table 2 and Appendix F, we observe that if it were not for the Twitter drop, there would be a consistent difference of about 5 % between the plot for the *Git* security issue and those for the privacy issues. This suggests that either websites were more inclined to fix security vulnerabilities or privacy issues required more time to address. Given the two-month period of our experiments, we could not conclusively figure out the exact reason. With slightly higher effects (diff column in Table 2), *No Consent* appeared to be the privacy issue easiest to remediate.

4.4.3 Statistical Significance. We also investigated the significance of differences in remediation rates at the aforementioned four points in time. We conducted Fisher's exact tests [19, 20] on the null hypothesis that the number of `problematic` vs. no longer `problematic` websites in the experimental conditions (*Warning*

[1] See https://developer.twitter.com/en/updates/changelog for Nov 15, 2021.

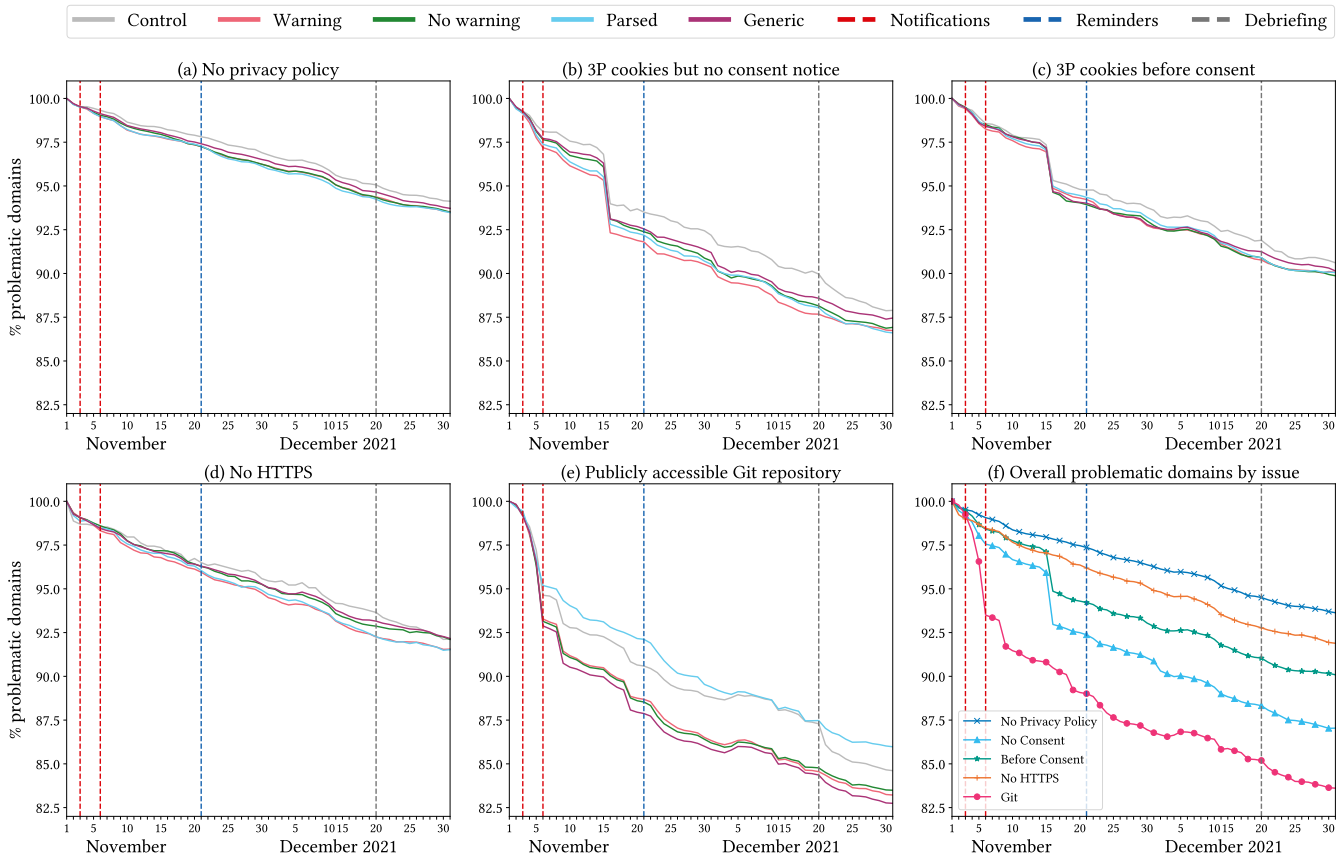

Figure 1: (a)–(e): Percentage of domains considered `problematic` according to our sliding window evaluation, by issue and experimental condition / email type. (f): Problematic domains by issue across all experimental conditions, including the control group. Vertical lines of the same color demarcate the periods or points in time when notification emails were sent. There are no data points for December 11–13 due to necessary hardware maintenance.

and *No Warning*) does not significantly differ from the *Control* group, and in an identical fashion for email types. Table 2 in its p column shows the results for November 10, 2021 and Appendix F for all dates. We used the Holm-Bonferroni method [27] to correct for multiple testing over time. While for November 10 the null hypothesis cannot be rejected for *Before Consent* and *No HTTPS* regardless of the presence of warnings, there is a significant difference in the distribution of `problematic` domains for *No Consent* and *Git* in both *Warning* and *No Warning* conditions.

As shown in Appendix F, these observations largely also hold true over time. The only cases where differences between *Control* and the treatment groups emerged later was *No Privacy Policy*, for which the *No Warning* condition did not lead to rejection of the null hypothesis on Nov 10 (while it did on any other date and across all dates for the *Warning* case), and *No HTTPS*, which only yielded significant differences to the control group for the *Warning* condition on November 28, a week after the reminder. This could confirm an earlier security notification study that found a limited effect of reminders [49], but differences between issues suggest that some privacy issues take a longer time to be addressed. For email type, we observe similar tendencies over time. Across issues, *Parsed* email addresses more frequently resulted in statistically significant differences in fix rates compared to the control group, except for the surprising case of *Git*.

4.4.4 Remediation by Website Popularity. Websites' willingness to remediate security and privacy shortcomings may depend on available human and monetary resources and how well a site is maintained in general. Hence, we investigated if issues are less likely to persist on more popular websites. To obtain popularity metrics for as many websites as possible, we queried the Google Chrome User Experience (CrUX) data set [24] via BigQuery as described by Durumeric [11] for the full domain rankings from November 2021, when we sent out notifications. For domains listed twice due to CrUX differentiating between HTTP and HTTPS, we used the higher rank. This yielded an overlap between the CrUX data (8,733,078 origins, 8,567,511 domains) and the sites we monitored (159,035) of 21,592 domains, 7 of which were ranked top 1,000 by CrUX, 43 top 10K, 432 top 100K, 3,800 top 1M and 17,721 top 10M. Due to the low number of domains per popularity bin we focused on comparing remediation rates between CrUX-ranked domains (n = 21,592) and unranked ones (n = 137,443), i. e., the long tail.

Table 3: Logistic regression models for *No Consent* and *Git*. Figures without brackets denote estimates and figures in brackets the standard error. $^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$.

	Nov 10	Nov 17	Nov 28	Dec 5
No Consent				
Intercept	-2.48 ***	-2.14 ***	-2.01 ***	-1.94 ***
	(0.05)	(0.04)	(0.04)	(0.04)
No warning	0.25 ***	0.18 **	0.25 ***	0.15 **
	(0.07)	(0.06)	(0.06)	(0.06)
Warning	0.31 ***	0.23 ***	0.27 ***	0.17 **
	(0.07)	(0.06)	(0.06)	(0.06)
Num. obs.	17,978	18,166	18,027	17,608
Git				
Intercept	-2.68 ***	-2.54 ***	-2.31 ***	-2.22 ***
	(0.04)	(0.04)	(0.04)	(0.04)
No warning	0.20 **	0.25 ***	0.32 ***	0.31 ***
	(0.06)	(0.06)	(0.05)	(0.05)
Warning	0.17 **	0.19 **	0.28 ***	0.28 ***
	(0.06)	(0.06)	(0.05)	(0.05)
Num. obs.	23,664	23,654	23,538	23,498

Contrary to our expectations, for most issues and points in time we found the rates of `problematic` domains for CrUX-ranked domains to exceed those for unranked ones by 0–1 percentage points. For Git, this was even more pronounced, with differences mostly between 2–3 percentage points, except for the *Git–Parsed* combination, which followed the overall 0–1 % pattern. We presume this difference to be mainly due to issues "fixing" themselves naturally, with unranked websites being less reliable to reach and more likely to be taken down permanently. Breaking this pattern, for *No Privacy Policy* notifications to *Parsed* email addresses, CrUX domains consistently exhibited lower rates of `problematic` checks, though the difference also mostly lay between 0–1 percentage points.

4.5 Influence of Warnings

For more insights into the effect of the presence of warnings on fix rates, we took a closer look at the set of domains for which at least one email could be successfully delivered according to our earlier definition, i. e., handed over to the next mail server. To determine the influence of warnings on fix rates, we computed logistic regression models [41, 42] for each issue and four different points in time: one and two weeks after the start of sending initial notifications and reminders, respectively. Table 3 shows the regression models for *Git* and *No Consent*, both relative to the *Control* group; Appendix G for all investigated issues. For *Git*, *No Privacy Policy*, and *No Consent*, we observe a statistically significant influence of the *Warning* and *No Warning* conditions on fix rates compared to the *Control* group, while the models do not show such influence for the *Before Consent* and *No HTTPS* issues. Still, even in the case where *Warning* and *No Warning* perform significantly better than the *Control* group, we cannot observe any difference between the estimates for these two conditions that does not fall within the standard error. This highlights that while receiving a notification significantly improved remediation, the presence or absence of a warning does not. This contrasts with prior work [38], which claimed warnings did play a role in remediation success (albeit limited to German websites).

5 GATHERING RECIPIENT FEEDBACK

In order to gain further insights into the measured changes (or lack thereof), we leveraged two channels of communication with recipients: an online survey and email conversations.

5.1 Survey

For consistency and to make sure participants received the survey invitation when the decision how to react to the notification was still fresh in their minds, we sent the survey link with all emails, i. e., initial notification and reminder for domains in experimental conditions (*Warning* and *No Warning*) and with the debriefing message for the *Control* group. The survey was implemented using a LimeSurvey instance hosted at Ruhr University Bochum. We first asked participants to assess the correctness of our checks (Q1), about prior awareness of the detected issue(s) (Q2), and plans to address them (Q3–4). Participants with privacy issues on their website were asked about the applicability of the GDPR (Q5–6), past changes to their website due to privacy legislation (Q7–8), and the influence of GDPR-mandated fines (Q9–10). Next, we asked all participants what type of support they would find useful to fix the issue(s) (Q11). We asked for participants' role(s) with regard to the website (Q12) to determine if we had reached a person with a suitable background to address the issue(s). The survey concluded with the opportunity to provide general feedback about our study (Q13). The full questionnaire can be found in Appendix E. One of the authors conducted a manual thematic analysis on open-ended survey answers to inductively identify common themes and sentiments. We categorized the answers via labels informed by the survey questions and additional themes found in the answers. Note that survey responses are subject to self-selection bias, which includes people with a strong (negative) experience with our notification process being more likely to provide feedback.

5.2 Email Communication

Sending automated emails at large scale inevitably results in large volumes of incoming mail, including automated responses from ticketing systems, delivery status notifications, and "out of office" messages. To support participants with fixing identified issues and obtain more information how our notifications were perceived, we focused on incoming responses composed by humans. We answered emails in German or English, depending on the language used by the sender, and if requested, we also sent a German translation of our notification message. When asked for advice, we only referred to third-party resources that either had been published by the hosting company of a website, by a data protection authority in the website operator's country, or by the vendors of third-party software already used by the website. Upon request we provided the names of third-party cookies or URLs to Git repositories that had triggered a `problematic` check result.

The researcher who had signed the notifications answered incoming emails from notification recipients according to these guidelines and categorized them on the conversation level by identifying recurring themes. After about 50 conversations we found the topics to have reached saturation. From the resulting list we removed very rare codes, refined the definitions of the emerged categories, and

Table 4: Survey participant sample (n = 212).

	n	%		n	%
Warnings?			**Issues**		
Warnings	95	44.8	No privacy policy	30	14.2
No warnings	117	55.2	No consent notice	22	10.4
Email type			Before consent	27	12.7
Parsed	105	49.5	No HTTPS	4	1.9
Generic	107	50.5	Git	140	66.0
Notif. Round			*Privacy only*	72	33.0
Initial	149	70.3	*Privacy & Git*	2	0.9
Reminder	40	18.9	*Git only*	138	65.1
Debriefing	23	10.9			

added (counter)examples. Our final codebook (see Appendix H) comprised 23 codes in the following six categories:

- **Sentiment** (3 codes): Expressions of gratitude and positive or negative sentiments towards our project.
- **More information** (8): Requests for more information, e. g., about our checks, the cookie or Git URL that had triggered them, about our research project, and if an issue was still detected after changes had been made to the website.
- **Performed actions** (3): Status reports from recipients, including already fixed issues, future plans to fix them, or forwarding the notification email to the responsible people.
- **Correctness** (4): (presumed) false positives, including the website being outside EU jurisdiction or (purportedly) not processing visitors' personal information.
- **Language** (2): Sentiments concerning language, such as translation requests or consternation about the emails being in English instead of the sender's / recipient's language.
- **Other** (3): Including sentiments that the notification could be spam, verification requests to other points of contact at our institution, and opt-out requests.

The remaining emails were single-coded by two coders, with uncertainties resolved via discussion. Each conversation could be assigned an arbitrary number of codes. As we had not passed unique IDs to the survey, we could not identify overlap between survey and email respondents, so some individually reported sentiments from these analyses may originate from the same person or domain.

6 PARTICIPANT FEEDBACK

6.1 Overview of Survey and Email Responses

6.1.1 Survey Participants. The survey link was clicked 1,890 times. 1,556 people only accessed the welcome page, 121 provided partial responses, and 213 completed the survey. We discarded all incomplete responses, plus one response for which the survey parameters had not been passed and questions based on specific detected issues had not been displayed. This left us with 212 complete responses. Table 4 shows the sample of participants who took the survey by experimental condition, email type, detected issues on the website, and notification round. While full responses were roughly equally distributed between the *Warning* and *No Warning* conditions, as well as email address type, most responses were collected through initial notifications (including control group debriefing), as opposed to the reminder. This corresponds with earlier findings that notification recipients either tend to act upon the first received email or not at all [35]. Two thirds of survey participants had been notified due to an open Git repository, while privacy issues less frequently motivated people to take the survey. This hints at security notifications being either taken more seriously or at least leading to recipients' higher willingness to interact with us. Appendix E provides an overview of all survey questions with response counts.

6.1.2 Email Communication. We received a total of 760 emails in 621 conversations with the operators of notified domains. 414 of these domains had been contacted because of a privacy issue and 167 because of a publicly accessible Git repository. 19 emails could not be assigned to a domain due to a lack of provided information. Most emails (662) had been sent to the two email addresses designated for this study. We also received 20 emails forwarded by the CISPA front office, and 85 had been sent to the institutional email address of the author who had signed the notification email. This was mainly done to verify if our notifications were legitimate. 7 emails had been sent to both the author and the project addresses. Appendix H shows how often each code (see Section 5.2) was assigned to email conversations with domains with security and privacy issues.

6.1.3 Who Did We Reach? If emails were successfully delivered, they had significant impact on fix rates, so we wanted to understand who the recipients were.

In the multiple-choice Q12 the majority of survey participants reported technical roles (developer and similar 57.1 %; administrator/operator 60.4 %), followed by roles related to the website's content (19.8 %) and product/project management (13.7 %). As for people in legal advisory roles, data protection officer ranked fifth (9.9 %), while the role as legal counsel was rarer (2.8 %). The involvement of these two roles was equally distributed between privacy and the Git issue(s). While survey-takers only represent a fraction of emailed websites, this provides evidence that the people who ultimately felt incentivized to react to the notification mainly hold responsibility for a website's technical administration or content.

In email conversations we looked for how often the recipient referred the handling of the issue to another person (code: *notified*). This was the case in 15.9 % of privacy conversations and in 10.2 % of those about the Git issue. Explicitly mentioned people or entities for both types of issues included the IT department or webmaster, in the Git case the security team, and for privacy notifications a lawyer, the cookie plugin provider, or the marketing department.

6.2 When Do Recipients (Plan to) Remediate?

Beyond the daily website checks, we wanted to gain additional insights into notification recipients' remediation behavior and future plans to address the issue(s) (or not). For this, we analyzed recipients' reported awareness of the issues and willingness to remediate them. In survey Q2 we found that while most participants (72.6 %) reported to have been unaware of the issue(s) prior to receiving the notification, this number was higher for security (81.4 %) than for privacy issues (56.8 %). Participants' reported plans to make subsequent changes to the website (Q3) also differ: While overall 81.6 % planned to make changes, this applies to 90.7 % of participants notified of Git issues but only 64.9 % of people with privacy problems. Pairing these results from Q2 and Q3, it seems that privacy issues are more often knowingly ignored.

Email analysis revealed similar differences in remediation intentions. Privacy notification recipients mostly told us that the issue(s) would be handled in the future (37.9 %, code: *will-handle*), while only 14.5 % stated that they had already been fixed. For the security notification, we saw the opposite: For *Git*, in 16.2 % of conversations the recipients stated that the issue would be handled in the future, while 44.9 % reported that the issue had already been fixed. As in the survey answers, this could either mean that privacy issues are more likely to deliberately be left unfixed – or that fixes simply take longer as they are more complex and may require the involvement of legal professionals, for example, to draft a privacy policy.

6.3 Roadblocks to Notification Success

6.3.1 Language Barrier. As we emailed domains from various countries in English, we may have contacted recipients in a language they do not understand, as first indicated by 3 survey participants in Q13 who found it hard to understand or assess the trustworthiness of an email written in English. More concrete evidence were the 17 translation requests we received via email. Most requests were for German, but some also for French and Czech.

6.3.2 Notifications Perceived as Spam or Otherwise Malicious. When asked for general feedback in survey Q13, 9 out of 76 participants (11.8 %) mentioned that they initially had been suspicious that the notification was spam or a scam attempt. To fight this impression, one participant suggested to point out the S/MIME signature in the email body, as "[s]pammers don't go out of their way to sign their emails from a public CA issued PEM certificate" (P1254).

Similarly, email analysis found in 7.2 % of conversations about Git-notified domains and in 12.1 % of privacy-related correspondence that the recipient was not sure if the notification email had been sent with benign intentions (code: *unsure-scam*). A special case were emails asking if the notification email had really been sent by our institution. 4 such verification requests were sent to dedicated project email addresses, 35 to the institutional address of the author who had signed the notification emails, and another 14 were sent to CISPA's front office. The language of the notifications may also have contributed to this. 6 email respondents wondered why we, as German research institutions, sent English notifications (code: *expected-german*); 3 of them stated they were not sure if our email was benign. This was also observed by Li et al. [35], who reported that emails the recipient expected to be in a different language (e. g., based on the sender's country of origin) were sometimes considered phishing or spam.

6.3.3 (Perceived) Incorrectness of Reports. In survey Q1 the majority of participants answered that the report was correct, with rates highest for *Git* (87.9 %) and lowest for *No Consent*[2] (45.5 %). Correspondingly, reported incorrectness rates were lowest and highest for these two cases, with 5.7 % and 22.7 % respectively. Uncertainty about the correctness of the report was highest for the *Before Consent* (29.6 %) and *No Consent* (18.2 %) cases. This hints at notification recipients often finding it difficult to determine which third-party cookies were present on their website and if they required user consent. While not a true false positive for the *Git* check, 3 survey participants notified about it replied in Q4 that they did not intend to make changes because the issue did not pose a security risk, as their Git repository did not contain any sensitive information, was not under their control or not accessible.

In emails we also received feedback about check results being false positives, for 18.1 % of conversations about a privacy issue but also for 6.0 % of those about Git. 16 emails stated that no sensitive data was stored inside the Git repository, though 3 reported that they had still made the repository inaccessible.

Manual checks revealed the majority of false-positive claims for the *Git* case to be due to failure to reproduce the issue. Many recipients of Git notifications tried to access `<domain>/.git/`, saw a "Forbidden" error page from their web server, and falsely assumed that this meant the Git repository was inaccessible, while in fact directory listing was forbidden. It is likely that they then stopped to further investigate the issue, leaving it unfixed.

6.3.4 Perceived Inapplicability of Privacy Legislation. One recurring theme in both the answers to multiple survey questions and email conversations about privacy notifications was recipients' perception that the privacy legislation in question did not apply to their website. In survey Q5 the 74 participants notified about a privacy issue were asked whether they thought the GDPR applied to their website. 63.5 % thought it did, while 20.3 % did not think so and 9.5 % were not sure. When asked in Q6 why they thought the GDPR did not apply, we received 14 replies. 6 answered that they were not located in the EU ("Because the UK is no longer in the EU" (P349)), illustrating unawareness of the GDPR's extraterritorial applicability to non-EU websites with EU visitors. Interestingly, several of these respondents were based in the UK, which still has a verbatim copy of the GDPR in its national legislation, so the same legal requirements apply. Regarding the material scope of the GDPR, 7 participants claimed to not process any personal data, unaware that even temporary storage of IP addresses is considered processing of personal data under EU law (see Section 3.1.1).

6.3.5 Privacy Indifference. Another demotivating factor that only emerged for privacy notifications was a general disdain for privacy legislation and its requirements. We found such sentiments in the answers to multiple open-ended survey questions. Asked in Q4 why they did not intend to add a privacy policy to their website, one participant replied "because these rules are plain stupid!" (P1131), and in Q8 another refused to make any changes at all due to the GDPR: "does not matter – GDPR is sucks [sic]" (P671).

6.4 Motivation for Remediation

The survey also investigated factors that motivated participants to (want to) take action, particularly awareness of fines mandated by the GDPR. In Q9 more than half of the 74 survey participants with privacy notifications (38, 51.4 %) had already been aware of these fines before the notification. Another 9 (12.2 %) had learned about them via the email, with 7 participants in the *Warning* and 2 in the *No Warning* condition. Still unaware of fines were 20 participants (27.0 %), 8 from the *Warning* and 12 from the *No Warning* group. Half of the 8 warned but unaware participants had only seen ePrivacy warnings, but the other 4 had (also) been warned about GDPR-mandated fines. This shows that notifications have limited educational impact about privacy legislation and potential fines.

[2]We do not consider *No HTTPS* here, as only 4 survey participants had this issue.

Q10 explored how knowledge of GDPR-mandated fines had influenced participants' decision to fix the detected issue. 13 participants (27.7 %) answered that they had not been influenced by the risks of fines, stating as their motivation that they "wanted to be responsible" (P45) or "believe[d] that GDPR and compliance with it [was] important" (P1505). Another 13 (27.7 %) explicitly acknowledged that fines were a motivating factor in their decision ("want to prevent paying fines" [P973]). 7 of them had received a warning notification and 6 an email without a warning. Though these answers may suffer from social desirability bias, this hints at fines for GDPR noncompliance being known and influencing fix rates, regardless of whether they were explicitly mentioned in the notification.

6.5 How Can We Help Websites Fix Issues?

Survey Q11 asked what additional information recipients would have wished for to better understand and fix the notification issue(s). We received 129 open-ended responses, many of which expressed generic sentiments: 37.7 % found the notification helpful and the information to be sufficient. 12.3 % would have appreciated more detailed guidelines or links to external resources on how to fix the issue(s). 16.5 % asked for additional documentation of our checks: 15 *Git*-notified participants suggested to add the URL for the repository in question, and 2 asked to include a check whether any sensitive information was present in the repository. While the first suggestion is an easy fix for future notifications, the latter would require extensive resources and would raise ethical concerns. Regarding the use of cookies without consent, we were repeatedly asked to add to the notification the names of the third-party cookie(s) that had triggered the `problematic` flag, which is also feasible.

Classification of email conversations confirmed this. A major category were requests for more information: 8.0 % of the 414 conversations regarding domains with privacy issues asked about privacy checks, especially the name of the problematic cookie (5.3 %); and questions concerning Git checks (10.8 % of the 167 Git conversations) most often were interested in the URL of the publicly accessible repository (5.4 %). We also received more generic requests, such as what to do in general, if certain changes would make the website compliant with privacy law, or about our research project.

6.6 How Were the Notifications Perceived?

It should be best research practice to notify affected parties about potential security or privacy issues on their systems, but there is a risk of backlash. Hence, in survey Q13 we asked for general feedback about our project and received 76 responses. Sentiments varied greatly between recipients of security and privacy notifications. 76.0 % of respondents notified about Git thanked us for the notification or voiced positive feedback ("This is an amazing project, please keep up the good work to make the internet a more secure place!" (P1675)), but only 50 % of respondents with privacy notifications did so ("Thank you for this hint! There are so much [sic] rules. For a little webmaster it's hard to know everything. It's really great to know there are some who help the little ones ;-)" (P840)). Negative sentiments were only expressed for privacy notifications ("the stated analysis is only 'may be' .... You have just wasted our time & energy" (P1685)). Repeated criticism included the privacy notifications being false positives, too threatening, unwanted, not sent in the participant's language, or sent with ill or monetary intentions. This confirms the sentiments reported in Section 6.3.

Email feedback contained similar differences in sentiment. Here people thanked us for the notification (code: *thanks*) in 74.9 % of conversations with recipients of security notifications, but only 56.0 % of privacy conversations. To distinguish between "thanks" and real enthusiasm for our project, we used the code *great-project*, assigned to 16.2 % of security and 2.9 % of privacy conversations. Correspondingly, the distribution of negative sentiments towards our notifications or project was reversed, assigned to 5.3 % of privacy- but only 1.8 % of security-related email conversations.

7 DISCUSSION

Our study identified recommendations both for future research in web privacy as well as for the public entities tasked with the application and enforcement of privacy legislation.

7.1 Privacy vs. Security Notifications

We compared the effects of notifications about privacy issues with those about a security problem. Challenges faced by both types include how to reach the responsible parties, language barriers, and lack of a trustworthy messaging channel. These have already been identified by previous work on security notifications and continue to pose significant problems for any campaign that aims to reach people via email at scale, as it is hard for both computer systems and humans to differentiate automated emails sent with beneficial intent from those sent for malicious purposes. Specific to privacy notifications is the obstacle that many recipients are not aware that certain legal requirements apply to their website, either because of misconceptions regarding the territorial scope of privacy laws or the website's data processing operations. Future research in this area is encouraged to educate notification recipients in this regard and provide concrete information why the respective law was deemed to apply to the recipient's website. This is especially important given our observation that the motivation to make changes due to privacy notifications appears to be extrinsic (awareness of fines) more often, while fixes of security issues tend to be more intrinsically motivated.

7.2 Message Tone and Content

Security and privacy notifications were met with very different sentiments, which could be rooted in message content (mentioning of privacy laws) or tone (presence of warnings). Security notifications evoked more positive sentiments and fewer threats of legal action. To relieve recipients' anxiety and anger, we recommend researchers in future privacy notification studies to *explicitly* explain that they will not pursue any legal action against the recipients. Recipient feedback also indicated widespread problems to identify the third-party plugin or subpage that had triggered the placement of a third-party cookie before or without the visitor's consent. We recommend that future notification studies provide the necessary details to help recipients pinpoint the problem, in our case the Git URL or concrete third-party service or cookie and subpage that triggered the detection mechanism. Links to concrete guidelines by DPAs or courts (e. g., that pre-ticked checkboxes do not constitute informed consent [16]), could aid notification recipients in understanding the issue and why it applies to their website.

7.3 Call for Guidelines and Standardization

Our work shows that it is generally feasible to identify privacy issues on websites independent of specific vendors or consent frameworks, with a prevalence of false positive cases in the low single digits in our set of manually verified websites. Still, these checks were designed to minimize the number of false positives for the purposes of this study and false negatives were not a concern. It may well be a requirement in other contexts, such as automated privacy audits designed to take the burden off DPAs in enforcing privacy legislation. Identifying vendor-agnostic privacy issues at scale with low numbers of both false positives and false negatives is still a significant challenge, given the differences in implementation of, for example, privacy policies or consent to the processing of personal information. Regulators could aid privacy researchers – and themselves in enforcing privacy laws – by issuing more concrete guidelines how to implement legal requirements, as in the example of the CCPA that requires a "Do Not Sell My Personal Information" link. The persistent challenge to enforce privacy laws online in the light of limited human and monetary resources requires long-term assistance via automated audits. The challenges with vendor-agnostic checks could be alleviated via standardization of how privacy-related information is presented on the Web. Hence, web researchers and standardization committees are encouraged to create new and build upon existing proposals how to unify the presentation and user control of a service's data processing practices. To prevent any such standard from suffering the same fate as past web privacy mechanisms relying on voluntary adoption, such as DNT or P3P, regulators should make adherence mandatory.

7.4 The Challenge of Reachability

Using email addresses found on websites had promising results in terms of reachability – 87.8 % of initial notifications for *Parsed* were successfully delivered, but only 33.8 % of emails to *Generic* addresses. While this does not guarantee that the correct person is reached and they act upon the notification, this approach can help overcome one of the obstacles in the way to reach the people responsible to fix security or privacy issues on websites. It needs to be noted that this approach can possibly introduce bias, as well-maintained websites are more likely to provide contact information, including an email address. The reachability problem also provides an opportunity for standardization: In the vein of `security.txt` [21, 22], a proposed standard to help web security researchers identify points of contact, a file `privacy.txt` could serve this information for privacy-related issues – and also communicate a website's data processing practices in a standardized format or at least contain a link to its privacy policy. Until then, future privacy notification studies could also leverage `security.txt` to contact websites about privacy issues.

7.5 The Future of Privacy Notifications

Our results show limited success of our notification campaign, especially when weighting this outcome against the resources required to detect and notify websites about privacy issues at scale. Similarly, the improvements proposed above will be futile if recipients of notification emails do not trust the sender do not consider privacy violations a problem whose remediation is an urgent matter. Still, we believe that large-scale privacy notifications can be a valuable tool in improving web privacy, but they need to be accompanied by other measures to overcome these obstacles.

To increase sender credibility and authority, researchers could cooperate with data protection authorities for future notification campaigns about issues related to data protection and privacy. The role of the DPA would be to provide sender credibility via their legitimization as a public authority and to accompany the campaign with information and enforcement capabilities to raise awareness for the issues. They could communicate to the general public the goal of the notifications, participating research institutions, investigated issues, guidelines on how to fix them, territorial and material applicability of the relevant laws, and possible consequences of non-compliance. A DPA informing about the latter is likely to be met with less adversity, as enforcing applicable privacy laws is their core task. Researchers, in turn, could supply the personal and technical resources that public authorities often lack and provide the notification infrastructure, expertise to more reliably detect compliance issues at scale, and (limited) support with fixing them.

Past campaigns show that this could be a promising approach. Some DPAs already have experience with privacy-related web measurements, such as the "Cookie Sweep" [2] carried out by multiple national DPAs in 2014 to inform EU institutions about websites' use of cookies and obtain first evidence for ePrivacy compliance. The notification campaign by privacy NGO noyb [44] (see Section 2) shows how external entities can support authorities in enforcing privacy legislation. Manual analyses limited the scope of these campaigns, but they both illustrate where DPAs and privacy researchers could benefit from each other to help enforce privacy legislation at scale. These multi-modal campaigns might not only have a broader effect, but also provide opportunities for further research, such as evaluating the usability of guidelines to fix privacy issues.

8 CONCLUSION

We conducted a large-scale email notification campaign to investigate if this approach can also help websites fix more complex privacy issues like missing privacy policies and incorrectly implemented consent notices and to determine how they compare to notifications about security vulnerabilities. Though overall fix rates are higher for security than privacy issues and the latter show tendencies to be addressed at a later point in time, we still find a statistically significant influence of our notifications on remediation rates. To overcome the problem of websites being hard to reach, we identify a promising approach in automatically extracting contact information from websites. Qualitative feedback from email conversations with recipients and survey responses shows that website owners are less open towards notifications about privacy issues than a security vulnerability. Reasons include limited willingness to make changes for privacy compliance, widespread misconceptions about privacy laws' applicability, and often greater necessary effort to identify and fix the problem. Even though warnings about potential fines do not increase remediation rates, they do at times incur anxiety and anger with recipients and corresponding backlash towards the senders. Future work is encouraged to explore if more specific information about the privacy problems and assurance of benign intent can yield more positive reactions and make email notifications a tool that can support large-scale privacy compliance.

ACKNOWLEDGMENTS

We thank Ahmed Ali for his help with implementing and manually validating the privacy checks and Simon Lenau and Charlotte Schwedes for assisting with the statistical analyses. This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

REFERENCES

[1] Article 29 Data Protection Working Party. 2012. Opinion 04/2012 on Cookie Consent Exemption. https://ec.europa.eu/justice/article-29/documentation/opinion-recommendation/files/2012/wp194_en.pdf. (cited on p. 3).

[2] Article 29 Data Protection Working Party. 2016. *Cookie Sweep Combined Analysis – Report.* Technical Report 14/EN WP 229. European Commission, Brussels, Belgium. https://ec.europa.eu/newsroom/article29/item-detail.cfm?item_id=640605 (cited on p. 13).

[3] Article 29 Data Protection Working Party. 2018. Guidelines on consent under Regulation 2016/679. https://ec.europa.eu/newsroom/article29/items/623051/en. (cited on p. 3).

[4] Dino Bollinger, Karel Kubicek, Carlos Cotrini, and David Basin. 2022. Automating Cookie Consent and GDPR Violation Detection. In *Proceedings of the 31st USENIX Security Symposium* (Boston, MA) *(USENIX Security '22)*. USENIX Association, Berkeley, CA, USA. https://www.usenix.org/conference/usenixsecurity22/presentation/bollinger (cited on p. 3).

[5] Matthew Bryant. 2020. TheInternetBackup – Crowdsourced Internet Domain Database. https://web.archive.org/web/20200110214540/https://theinternetbackup.com/#about. Original no longer online; archived link last accessed April 1, 2022. (cited on p. 4).

[6] Orçun Çetin, Carlos Gañán, Maciej Korczyński, and Michel van Eeten. 2016. Understanding the role of sender reputation in abuse reporting and cleanup. *Journal of Cybersecurity* 2, 1 (Dec. 2016), 83–98. https://doi.org/10.1093/cybsec/tyw005 (cited on p. 2).

[7] Orcun Cetin, Carlos Gañán, Maciej Korczyński, and Michel van Eeten. 2017. Make Notifications Great Again: Learning How to Notify in the Age of Large-Scale Vulnerability Scanning. In *Proceedings of the 16th Annual Workshop on the Economics of Information Security* (La Jolla, CA) *(WEIS 2017)*. https://weis2017.econinfosec.org/wp-content/uploads/sites/3/2017/05/WEIS_2017_paper_17.pdf (cited on p. 2).

[8] CMS Law.Tax. 2022. GDPR Enforcement Tracker. https://www.enforcementtracker.com/. (cited on p. 1).

[9] Dave Crocker. 1997. Mailbox Names for Common Services, Roles and Functions. https://www.ietf.org/rfc/rfc2142.txt. (cited on p. 2, 5).

[10] Martin Degeling, Christine Utz, Christopher Lentzsch, Henry Hosseini, Florian Schaub, and Thorsten Holz. 2019. We Value Your Privacy ... Now Take Some Cookies: Measuring the GDPR's Impact on Web Privacy. In *26th Annual Network and Distributed System Security Symposium* (San Diego, CA, USA) *(NDSS '19)*. Internet Society, Reston, VA, USA. https://doi.org/10.14722/ndss.2019.23378 (cited on p. 1, 3, 4).

[11] Zakir Durumeric. 2023. Cached Chrome Top Million Websites. Retrieved February 26, 2023 from https://github.com/zakird/crux-top-lists (cited on p. 8).

[12] Zakir Durumeric, Frank Li, James Kasten, Johanna Amann, Jethro Beekman, Mathias Payer, Nicolas Weaver, David Adrian, Vern Paxson, Michael Bailey, and J. Alex Halderman. 2014. The Matter of Heartbleed. In *Proceedings of the 2014 ACM Internet Measurement Conference* (Vancouver, BC, Canada) *(IMC'14)*. ACM, New York, NY, USA, 475–488. https://doi.org/10.1145/2663716.2663755 (cited on p. 1, 2).

[13] Zakir Durumeric, Eric Wustrow, and J. Alex Halderman. 2013. ZMap: Fast Internet-wide Scanning and Its Security Applications. In *Proceedings of the 22nd USENIX Security Symposium* (Washington, DC) *(USENIX '13)*. USENIX Association, Berkeley, CA, USA, 605–619. https://www.usenix.org/system/files/conference/usenixsecurity13/sec13-paper_durumeric.pdf (cited on p. 5).

[14] Steven Englehardt and Arvind Narayanan. 2016. Online Tracking: A 1-million-site Measurement and Analysis. In *Proceedings of the 26th ACM Conference on Computer and Communications Security* (Vienna, Austria) *(CCS '16)*. ACM, New York, NY, USA, 1388–1401. https://doi.org/10.1145/2976749.2978313 (cited on p. 3).

[15] European Commission. 2022. What is personal data? Retrieved February 17, 2022 from https://ec.europa.eu/info/law/law-topic/data-protection/reform/what-personal-data_en (cited on p. 3).

[16] European Court of Justice. 2019. Judgment of the Court of 1 October 2019 in Case C-673/17 – Planet49. https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:62017CJ0673. (cited on p. 3, 12).

[17] European Data Protection Board. 2018. Guidelines 05/2020 on consent under Regulation 2016/679. https://edpb.europa.eu/sites/default/files/files/file1/edpb_guidelines_202005_consent_en.pdf. (cited on p. 3).

[18] Adrienne Porter Felt, Richard Barnes, April King, Chris Palmer, Chris Bentzel, and Parisa Tabriz. 2017. Measuring HTTPS Adoption on the Web. In *Proceedings of the 26th USENIX Security Symposium* (Vancouver, BC, Canada) *(USENIX Security '17)*. USENIX Association, Berkeley, CA, USA, 1323–1338. https://www.usenix.org/conference/usenixsecurity17/technical-sessions/presentation/felt (cited on p. 4).

[19] Ronald A. Fisher. 1922. On the Interpretation of χ^2 from Contingency Tables, and the Calculation of P . *Journal of the Royal Statistical Society* 85, 1 (Jan. 1922), 87–94. https://doi.org/10.2307/2340521 (cited on p. 7).

[20] Ronald A. Fisher. 1935. The Logic of Inductive Interference. *Journal of the Royal Statistical Society* 98, 1 (1935), 39–82. https://doi.org/10.2307/2342435 (cited on p. 7).

[21] Edwin Foudil and Yakov Shafranovich. 2017. `security.txt` – *A proposed standard which allows websites to define security policies.* Retrieved March 29, 2022 from https://securitytxt.org/ (cited on p. 13).

[22] Edwin Foudil and Yakov Shafranovich. 2022. *RFC 9116: A File Format to Aid in Security Vulnerability Disclosure.* Retrieved May 1, 2022 from https://www.rfc-editor.org/rfc/rfc9116 (cited on p. 13).

[23] Ghostery GmbH. 2022. WhoTracks.me. https://github.com/whotracksme/whotracks.me. (cited on p. 3).

[24] Google, Inc. 2022. CrUX API. Retrieved February 25, 2023 from https://developer.chrome.com/docs/crux/api/ (cited on p. 8).

[25] Ilse Heine. 2021. *3 Years Later: An Analysis of GDPR Enforcement.* Retrieved June 24, 2022 from https://www.csis.org/blogs/strategic-technologies-blog/3-years-later-analysis-gdpr-enforcement (cited on p. 1).

[26] Maximilian Hils, Daniel W. Woods, and Rainer Böhme. 2020. Measuring the Emergence of Consent Management on the Web. In *Proceedings of the ACM Internet Measurement Conference* (Virtual Event, USA) *(IMC '20)*. ACM, New York, NY, USA, 317–332. https://doi.org/10.1145/3419394.3423647 (cited on p. 3).

[27] Sture Holm. 1979. A Simple Sequentially Rejective Multiple Test Procedure. *Scandinavian Journal of Statistics* 6, 2 (1979), 65–70. https://www.jstor.org/stable/4615733 (cited on p. 8).

[28] Henry Hosseini, Martin Degeling, Christine Utz, and Thomas Hupperich. 2021. Unifying Privacy Policy Detection. *Proceedings on Privacy Enhancing Technologies* 2021, 4 (July 2021), 480–499. https://doi.org/10.2478/popets-2021-0081 (cited on p. 3).

[29] Xuehui Hu, Nishanth Sastry, and Mainack Mondal. 2021. CCCC: Corralling Cookies into Categories with CookieMonster. In *Proceedings of the 13th ACM Web Science Conference* (Online) *(WebSci '21)*. ACM, New York, NY, USA. https://doi.org/10.1145/3447535.3462509 (cited on p. 3).

[30] IAB Europe. 2022. CMP List. https://iabeurope.eu/cmp-list/. (cited on p. 3).

[31] Information Commissioner's Office. 2014. Protecting personal data in online services: learning from the mistakes of others. https://ico.org.uk/media/for-organisations/documents/1042221/protecting-personal-data-in-online-services-learning-from-the-mistakes-of-others.pdf. (cited on p. 4).

[32] Irish Data Protection Commission. 2020. Guidance note: Cookies and other tracking technologies. https://www.dataprotection.ie/sites/default/files/uploads/2020-04/Guidance%20note%20on%20cookies%20and%20other%20tracking%20technologies.pdf. (cited on p. 3).

[33] Konferenz der unabhängigen Datenschutzaufsichtsbehörden des Bundes und der Länder. 2020. Hinweise zum Einsatz von Google Analytics im nicht-öffentlichen Bereich (Notes concerning the use of Google Analytics in the non-public sector). https://www.datenschutzkonferenz-online.de/media/dskb/20200526_beschluss_hinweise_zum_einsatz_von_google_analytics.pdf. (cited on p. 2).

[34] Marc Kührer, Thomas Hupperich, Christian Rossow, and Thorsten Holz. 2014. Exit from Hell? Reducing the Impact of Amplification DDoS Attacks. In *Proceedings of the 23rd USENIX Security Symposium* (San Diego, CA) *(USENIX '14)*. USENIX Association, Berkeley, CA, USA, 111–125. https://www.usenix.org/conference/usenixsecurity14/technical-sessions/presentation/kuhrer (cited on p. 1, 2).

[35] Frank Li, Zakir Durumeric, Jakub Czyz, Mohammad Karami, Michael Bailey, Damon McCoy, Stefan Savage, and Vern Paxson. 2016. You've Got Vulnerability: Exploring Effective Vulnerability Notifications. In *Proceedings of the 25th USENIX Security Symposium* (Austin, TX) *(USENIX '16)*. USENIX Association, Berkeley, CA, USA. https://www.usenix.org/conference/usenixsecurity16/technical-sessions/presentation/li (cited on p. 1, 2, 5, 10, 11).

[36] Max Maass, Marc-Pascal Clement, and Matthias Hollick. 2021. Snail Mail Beats Email Any Day: On Effective Operator Security Notifications in the Internet. In *Proceedings of the 16th International Conference on Availability, Reliability and Security (ARES 2021)*. ACM, New York, NY, USA. https://doi.org/10.1145/3465481.3465743 (cited on p. 1, 2, 3, 4, 5).

[37] 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 (ARES 2021)*. ACM, New York, NY, USA. https://doi.org/10.1145/3465481.3470081 (cited on p. 2, 5, 6).

[38] 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 *Proceedings of the 30th USENIX Security Symposium* (Virtual Event) *(USENIX '21)*. USENIX Association, Berkeley, CA, USA. https://www.usenix.org/conference/usenixsecurity21/presentation/maass (cited on p. 1, 2, 6, 9).

[39] Célestin Matte, Nataliia Bielova, and Cristiana Santos. 2020. Do Cookie Banners Respect my Choice? Measuring Legal Compliance of Banners from IAB Europe's Transparency and Consent Framework. In *Proceedings of the 2020 IEEE Symposium on Security and Privacy* (San Francisco, CA, USA) *(SP '20)*. IEEE, Piscataway, NJ, USA, 791–809. https://doi.org/10.1109/SP40000.2020.00076 (cited on p. 3).

[40] MaxMind, Inc. 2022. GeoLite2 Free Geolocation Data. https://dev.maxmind.com/geoip/geolite2-free-geolocation-data. (cited on p. 4).

[41] Peter McCullagh and John A. Nelder. 1989. *Generalized Linear Models.* Chapman & Hall, London, United Kingdom. https://www.routledge.com/Generalized-Linear-Models/McCullagh-Nelder/p/book/9780412317606 (cited on p. 9).

[42] John A. Nelder and Robert W. M. Wedderburn. 1972. Generalized Linear Models. *Journal of the Royal Statistical Society, Series A* 135, 3 (1972), 370–384. https://doi.org/10.2307/2344614 (cited on p. 9).

[43] Midas Nouwens, Ilaria Liccardi, Michael Veale, David Karger, and Lalana Kagal. 2020. Dark Patterns after the GDPR: Scraping Consent Pop-ups and Demonstrating their Influence. In *Proceedings of the ACM CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI '20)*. ACM, New York, NY, USA, 1–13. https://doi.org/10.1145/3313831.3376321 (cited on p. 3).

[44] noyb. 2021. noyb files 422 formal GDPR complaints on nerve-wrecking "Cookie Banners". https://noyb.eu/en/noyb-files-422-formal-gdpr-complaints-nerve-wrecking-cookie-banners. (cited on p. 1, 2, 3, 13).

[45] Sean O'Connor, Ryan Nurwono, and Eleanor Birrell. 2021. (Un)clear and (In)conspicuous: The right to opt-out of sale under CCPA. , 15 pages. arXiv:2009.07884 https://arxiv.org/abs/2009.07884 (cited on p. 3).

[46] Cristiana Santos, Nataliia Bielova, and Célestin Matte. 2020. Are cookie banners indeed compliant with the law? Deciphering EU legal requirements on consent and technical means to verify compliance of cookie banners. *Technology and Regulation* 2020 (2020), 91–135. https://doi.org/10.26116/techreg.2020.009 (cited on p. 2, 3).

[47] Wissem Soussi, Maciej Korczyński, Sourena Maroofi, and Andrzej Duda. 2020. Feasibility of Large-Scale Vulnerability Notifications after GDPR. In *2020 IEEE European Symposium on Security and Privacy Workshops* (Online) *(EUROS&PW 2020)*. IEEE, 531–536. https://doi.org/10.1109/EuroS&PW51379.2020.00077 (cited on p. 5).

[48] State of California Legislative Counsel. 2018. *Assembly Bill No. 375 – Chapter 55.* https://leginfo.legislature.ca.gov/faces/billTextClient.xhtml?bill_id=201720180AB375 (cited on p. 3).

[49] Ben Stock, Giancarlo Pellegrino, Frank Li, Michael Backes, and Christian Rossow. 2018. Didn't You Hear Me? – Towards More Successful Web Vulnerability Notifications. In *25th Annual Network and Distributed System Security Symposium* (San Diego, California) *(NDSS '18)*. Internet Society, Reston, VA, USA. https://www.ndss-symposium.org/wp-content/uploads/2018/02/ndss2018_01B-1_Stock_paper.pdf (cited on p. 1, 2, 3, 4, 8).

[50] 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 *Proceedings of the 25th USENIX Security Symposium* (Austin, TX) *(USENIX '16)*. USENIX Association, Berkeley, CA, USA. https://www.usenix.org/conference/usenixsecurity16/technical-sessions/presentation/stock (cited on p. 1, 2, 4, 5).

[51] Samuel Stolton. 2020. *GDPR enforcement held back by lack of resources, report says.* Retrieved May 5, 2022 from https://www.euractiv.com/section/data-protection/news/gdpr-enforcement-held-back-by-lack-of-resources-report-says/ (cited on p. 1).

[52] TeleTrusT and ENISA. 2021. IT Security Act (Germany) and EU General Data Protection Regulation: Guideline "State of the art" – Technical and organisational measures. https://www.teletrust.de/fileadmin/user_upload/2021-09_TeleTrusT_Guideline_State_of_the_art_in_IT_security_EN.pdf. (cited on p. 4).

[53] The EasyList authors. 2022. EasyList Cookie List. Retrieved June 11, 2022 from https://github.com/easylist/easylist/tree/master/easylist_cookie (cited on p. 3).

[54] The European Parliament and the Council of the European Union. 2002. Directive 2002/58/EC of the European Parliament and of the Council of 12 July 2002 concerning the processing of personal data and the protection of privacy in the electronic communications sector. Official Journal of the European Communities. (cited on p. 3).

[55] Christine Utz, Sabrina Amft, Martin Degeling, Thorsten Holz, Sascha Fahl, and Florian Schaub. 2023. Privacy Rarely Considered: Exploring Considerations in the Adoption of Third-Party Services by Websites. *Proceedings on Privacy Enhancing Technologies* 2023, 1 (Jan. 2023), 5–28. https://doi.org/10.56553/popets-2023-0002 (cited on p. 3).

[56] 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, NY, USA, 973–990. https://doi.org/10.1145/3319535.3354212 (cited on p. 3).

[57] Rob van Eijk, Hadi Asghari, Philipp Winter, and Arvind Narayanan. 2019. The Impact of User Location on Cookie Notices (Inside and Outside of the European Union). In *Workshop on Technology and Consumer Protection* (San Francisco, California) *(ConPro '19)*. IEEE. https://www.ieee-security.org/TC/SPW2019/ConPro/papers/vaneijk-conpro19.pdf (cited on p. 3).

[58] Maggie Van Nortwick and Christo Wilson. 2022. Setting the Bar Low: Are Websites Complying With the Minimum Requirements of the CCPA? *Proceedings on Privacy Enhancing Technologies* 2022, 1 (Jan. 2022), 608–628. https://doi.org/10.2478/popets-2022-0030 (cited on p. 1, 3).

[59] Marie Vasek and Tyler Moore. 2012. Do Malware Reports Expedite Cleanup? An Experimental Study. In *5th Workshop on Cyber Security Experimentation and Test* (Bellevue, WA) *(CSET '12)*. USENIX Association, Berkeley, CA, USA. https://www.usenix.org/conference/cset12/workshop-program/presentation/Vasek https://www.usenix.org/conference/cset12/workshop-program/presentation/Vasek. (cited on p. 2).

[60] Thomas Vissers, Wouter Joosen, and Nick Nikiforakis. 2015. Parking Sensors: Analyzing and Detecting Parked Domains. In *Proceedings of the 2015 Network and Distributed System Security Symposium* (San Diego, CA, USA) *(NDSS 2015)*. Internet Society, Reston, VA, USA. https://doi.org/10.14722/ndss.2015.23053 (cited on p. 4).

[61] Eric Zeng, Frank Li, Emily Stark, Adrienne Porter Felt, and Parisa Tabriz. 2019. Fixing HTTPS Misconfigurations at Scale: An Experiment with Security Notifications. In *2019 Workshop on the Economics of Information Security* (Boston, MA) *(WEIS '19)*. https://weis2019.econinfosec.org/wp-content/uploads/sites/6/2019/05/WEIS_2019_paper_16.pdf (cited on p. 1, 2, 5).

A WORDS TO DETECT PII INPUT FIELDS

Names for input fields that likely contain personal data:

"newsletter", "login", "email", "username", "e-mail", "name", "firstname", "lastname", "gender", "birthdate", "bday", "dob", "dateofbirth", "sso", "signin", "signin__email", "login_email", "loginmodel-username", "connection_mail", "email_address", "login-user", "email_1_db", "login_pwd_db", "user_login"

B NOTIFICATION EMAILS

Hello,

We are a group of security and privacy researchers from the CISPA Helmholtz Center for Information Security and Ruhr University Bochum in Germany. As part of our current research project, we analysed potential security and data protection issues on websites.

We would like to raise your attention to the following security and data protection issue(s) on your website [DOMAIN]. Please note that we do not offer a conclusive legal assessment or consultancy on an individual website's legal compliance.

No privacy policy. For public websites that use European domains, are hosted in the EU, or may be used by European users, any collection of users' personal data is governed by the EU General Data Protection Regulation (GDPR). If a website meets these conditions, the operator is legally required by Article 13 of the GDPR to have a privacy policy explaining the use of their visitors' personal data. Personal data also encompasses the processing of communications data such as IP addresses of users even if no additional information is collected. The privacy policy has to inform users about the use of their personal data in a concise, transparent, intelligible, and easily accessible form.

Our automated analysis of your website did not detect a privacy policy, which may indicate noncompliance with the GDPR's information requirements.

Input fields for personal information without HTTPS. Article 32 of the GDPR requires data controllers such as website owners to implement appropriate technical and organisational measures

to ensure a level of security appropriate to the risk, taking into account the state of the art. Protection of users‘ communication and interactions with your website via HTTPS is considered state of the art in data security.

Our automated analysis detected input fields on your website that allow users to enter personal data without using HTTPS secure communication to prevent eavesdropping and phishing. This may indicate noncompliance with the GDPR's data security requirements.

Use of third-party cookies without consent notice. Under Article 5 Paragraph 3 of the EU ePrivacy Directive (Directive 2009/136/EC) and respective implementations of the Directive into national law of the EU member states, the setting of individual cookies on the user's terminal equipment that are not strictly necessary for the functioning of the website is only allowed if the user has given his or her prior consent.

Our automated analysis did not detect such a consent form for the third-party cookies on your website. This may indicate noncompliance with EU ePrivacy requirements.

Third-party cookies set before interaction with consent notice. Under Article 5 Paragraph 3 of the EU ePrivacy Directive (Directive 2009/136/EC) and respective implementations of the Directive into national law of the EU member states, the setting of individual cookies on the user's terminal equipment that are not strictly necessary for the functioning of the website is only allowed if the user has given his or her prior consent. Such consent has to be given in advance via a meaningful interaction by the user. According to our automated analysis, your website does provide users with a cookie notice or consent form, but the cookies are set before any meaningful interaction of a user with the consent form takes place. This lack of explicit consent may indicate noncompliance with EU ePrivacy requirements.

Publicly accessible Git repository. If the configuration folder for Git (`.git`) is reachable through HTTP, an attacker may copy the content of this repository. This allows an attacker to access the source code versioned in this repository, including any credentials or other sensitive data possibly stored there. Our automated analysis detected a publicly accessible Git repository on your website. Note that we only check for the existence of a repository and do not attempt to download any actual content. Hence, we cannot state if it contains any sensitive information.

If in Warning condition:

- *If Git:* Please note: In the worst case, access to configuration files with credentials could lead to an attacker taking over your entire website.
- *If No Privacy Policy or No HTTPS:* Noncompliance with GDPR requirements could lead to fines of up to 10 million euros or up to 2 percent of the global turnover of the preceding fiscal year according to Article 83 Paragraph 4 GDPR.[3]
- *If No Consent or Before Consent:* Fines for noncompliance with ePrivacy requirements may vary depending on national laws.

You can review more detailed information about the security and data protection issues and their remediation status on your website by visiting our web interface at https://notify.cispa.de/reports/[DOMAIN]/report-[UNIQUE_ID].

Since this notification is part of an ongoing research project, we will re-check your website to verify if the issues have been fixed. If you wish us to stop this check, please visit our web interface at https://notify.cispa.de/reports/[DOMAIN]/report-[UNIQUE_ID] to opt out or contact us at `info@notify.cispa.de`.

Help us improve our notification process with anonymous feedback at: https://notify.cispa.de/reports/[DOMAIN]/report-[UNIQUE_ID]/notification-survey.[4]

Should you need further information or have any other questions, please do not hesitate to contact us using the same email address.

Best regards,
Matthias Michels
Security Researcher
CISPA Helmholtz Center for Information Security
Stuhlsatzenhaus 5
66123 Saarbrücken
Germany

C INFO WEBSITE

We are security and privacy researchers from the Secure Web Applications Group (https://swag.cispa.saarland/) at the CISPA Helmholtz Center for Information Security and the Systems Security group (https://informatik.rub.de/syssec/) at Ruhr-Universität Bochum, both in Germany. We are currently conducting a research project on large-scale security and data protection notifications. With our notifications we would like to help website owners identify and fix security and data protection issues on their websites.

Our analysis tool checks websites for the presence of a privacy policy and a cookie consent notice, whether third-party cookies are being set before consent, potentially unprotected personal information in input fields, and publicly accessible code versioning repositories. If our tool detects an issue, we notify the website owner about it via e-mail. The checks are performed in a non-intrusive way. Our tool will never try to exploit a vulnerability on your server or interfere with your services.

In case you would like to contact us about this research, you can send an email to `info@notify.cispa.de`. If you want your websites to be excluded from our analysis, you can email us the domains, IP addresses, or IP ranges which should be excluded. Alternatively, if you have received an individual report for your website from us, you can use the opt-out buttons in that report.

[3]Due to oversight we did not differentiate in the warning text between the two tiers of fines in Article 83 GDPR: While not having a privacy policy (violates Article 13) is subject to the fines in Article 83(5) (20M/ 4% of annual turnover), non-use of HTTPS (violates Article 32) falls under Article 83(4) (10M / 2 % of turnover). We do not expect this difference in maximum fines to have any significant impact on notification recipients' remediation behavior.

[4]Clicking this link triggered a redirect to the survey. The `UNIQUE_ID` was only used to look up the notification issues, study conditions, and email group associated with the website, which were then translated into URL parameters for the survey link. No unique identifier was passed to the survey.

D SAMPLE REPORT

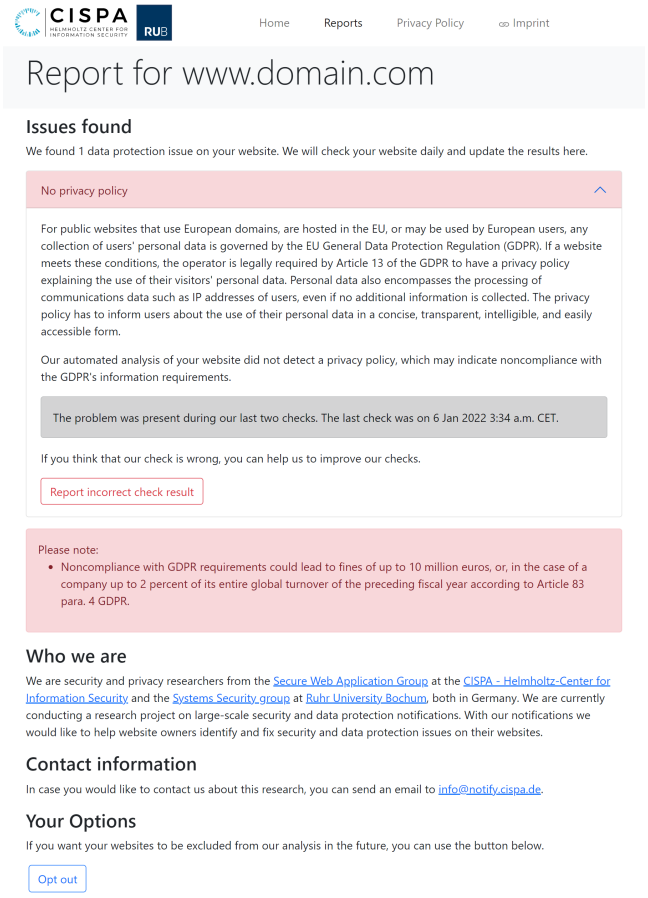

Figure 2: An example report for a website in the *No Warning* condition with a missing privacy policy. Email recipients could access it via the link "web interface at https://notify.cispa.de/reports/www.domain.com/report-[UNIQUE_ID]" in the email. The color of the accordion menu for each detected issue changed to display its current status: red for not detected as fixed, yellow for detected as fixed once within the last two days, and green for detected as fixed for at least two days. The red "Please note" box was only shown for websites in the *Warning* condition and displayed the corresponding warning message(s).

E SURVEY

Survey Title

Survey on Security and Data Protection Notifications

Intro Text

We are security and privacy researchers from the CISPA Helmholtz Center for Information Security and Ruhr University Bochum in Germany. In our current research we are trying to better understand how to notify websites about security and data protection issues. We recently emailed you a security and data protection notification from notify@notify.cispa.de.

You can help us improve our notification process through completing this survey. The survey is short and anonymous, and all questions are optional, so please answer the ones that you feel comfortable with. Your feedback is very valuable to us and we really appreciate your time.

Privacy Policy & Consent

We take great care in protecting our survey participants' privacy in accordance with the provisions of the General Data Protection Regulation (GDPR). Your answers to this survey will be stored securely on a server hosted by Ruhr University Bochum, Germany. Any of the survey data will only be accessible by the researchers involved in this project and will not be correlated with other data or otherwise used to identify individual participants. If we make data from this research available to the research community or the interested public, we will only publish it in an aggregated form that does not allow anyone to identify you or the website for which we sent you a notification email. You can find the contact information of the responsible data protection officers at https://notify.cispa.de/privacy_en.html.

Participation. Your participation in this research is completely voluntary. Once you have started the survey, you may cancel at any time by clicking the "Exit and clear survey" button in the upper right part of the screen, and your answers will be discarded.

Contact Information. If you have any questions, comments, or concerns about the study either before, during, or after participation, please contact us at info@notify.cispa.de.

Survey Questions & Responses

First we would like to ask you about the security and data protection issues we found on your website.
Here is a list of the issues we found: **[of the following, only the detected issues were shown]**

- No privacy policy
- Use of third-party cookies without consent notice
- Third-party cookies set before interaction with consent notice
- Input fields for personal information without HTTPS
- Publicly accessible Git repository

Table 5: Survey Questions & Responses. "Displayed" indicates how many participants had seen each question; "N/A" indicates how many of them did not provide an answer.

Q1: Do you think our report is correct regarding each of the detected issues? [list of detected issues as shown above; single choice for each]

	No Privacy Policy		No Consent		Before Consent		No HTTPS		Git	
	n	%	n	%	n	%	n	%	n	%
Yes	21	70.0	10	45.5	17	63.0	1	25.0	123	87.9
No	5	16.7	5	22.7	2	7.4	1	25.0	8	5.7
Uncertain	2	6.7	4	18.2	8	29.6	2	50.0	8	5.7
N/A	2	6.7	3	13.6	0	0.0	0	0.0	1	0.7
# Displayed	30		22		27		4		140	

Q2: Were you aware of this/these issue(s) before we contacted you? [single choice]

	Security		Privacy		All	
	n	%	n	%	n	%
Yes	18	12.9	21	28.4	39	18.4
No	114	81.4	42	56.8	154	72.6
Don't know	7	5.0	6	8.1	13	6.1
N/A	1	0.7	5	6.8	6	2.8
# Displayed	140		74		212	

Q3: Are you planning to make any changes to the website after receiving our message? [single choice]

	Security		Privacy		All	
	n	%	n	%	n	%
Yes	127	90.7	48	64.9	173	81.6
No	4	2.9	17	23.0	21	9.9
Don't know	7	5.0	3	4.1	10	4.7
N/A	2	1.4	6	8.1	8	3.8
# Displayed	140		74		212	100.0

Q4 **(If "No")**: Why are you not planning to make any changes? [free text, multiple codes per answer possible]

	n	%
Non-applicability of privacy law (generic)	1	4.8
Non-EU	4	19.0
No third-party cookies used	3	14.3
No personal data collected	7	33.3
Other	5	23.8
N/A	1	4.8
# Displayed	21	

Q5 **(If any privacy issue was detected)**: Do you think the European Union's General Data Protection Regulation (GDPR) applies to your website? [single choice]

	n	%
Yes	47	63.5
No	15	20.3
Don't know	7	9.5
N/A	5	6.8
# Displayed	74	

Q6 **(If "No")**: Why do you think the GDPR does not apply to your website? [free text]

	n	%
Non-EU	6	40.0
No personal data collected	7	46.7
Other	1	6.7
N/A	1	6.7
# Displayed	15	

Q7 **(If any privacy issue was detected)**: In the past, did you already make changes to the website because of the GDPR or other privacy legislation? [single choice]

	n	%
Yes	32	43.2
No	36	48.6
Don't know	1	1.4
N/A	5	6.8
# Displayed	74	

Q8 **(If "Yes")**: What changes did you make because of this privacy legislation? [free text; multiple codes per answer possible]

	n	%
Made changes to privacy policy	5	15.6
Installed cookie plugin or banner	13	40.6
Removed third-party service/cookies	6	18.8
Enforced HTTPS	2	6.3
Other	9	28.1
N/A	6	18.8
# Displayed	32	

Q9 **(If any privacy issue was detected)**: Were you aware of potential fines mandated by GDPR before you received our message? [single choice]

	n	%
Yes, since you emailed me	9	12.2
Yes, even before you emailed me	38	51.4
No, I'm not aware of them	20	27.0
N/A	7	9.5
# Displayed	74	

Q10 **(If either "Yes" option was selected)**: In which way did this knowledge of fines influence your decision to fix the issue(s)? [free text; single code per answer]

	n	%
Reported influence of fines	13	27.7
No reported influence of fines	13	27.7
Unrelated answer	6	12.8
N/A	15	31.9
# Displayed	47	

Q11: What type of support would you find helpful to fix the issue(s) we found on your website? [free text; multiple codes per answer possible]

	n	%
Info in notification was sufficient	80	37.7
Better documentation of checks	35	16.5
More information about fixes	26	12.3
Other	6	2.8
N/A	83	39.2
# Displayed	212	

Q12: What is / are your role(s) with regard to the website we notified you about? [multiple choice]

	n	%
Product or project manager	29	13.7
Content creator or contributor	42	19.8
Social media manager	8	3.8
Marketing	11	5.2
Sales	7	3.3
Quality assurance	12	5.7
User experience	11	5.2
(Web) developer, programmer, or software engineer	121	57.1
Administrator or (web) operator	128	60.4
Legal counsel	6	2.8
Data protection officer	21	9.9
Customer service / customer support / c. relations	11	5.2
Other: [free text]	21	9.9
N/A	11	5.2
# Displayed	212	

Q13: Is there anything you want to tell us about our checks, notifications, or any other issue related to this research or to security and data protection notifications in general? [free text; multiple codes per answer possible]

	n	%
Positive sentiment / thanks	51	24.1
Negative sentiment	4	1.9
Email first seemed suspicious	9	4.2
More information required	4	1.9
N/A	136	64.2
# Displayed	212	

F RATES OF PROBLEMATIC DOMAINS AND FISHER'S EXACT TEST RESULTS OVER TIME

Table 6: Percentages of websites still considered `problematic` with regard to each specific issue according to our sliding window evaluation (see Section 4.4.1) on the respective date (in 2021), by study condition / email type. % indicates the percentage of still `problematic` domains, diff the difference in percentage points to the *Control* group, and p the p-values for Fisher's exact tests ($\alpha = 0.05$) compared to *Control*. Bold values indicate those still significant after Holm-Bonferroni correction, while *italics* indicate no longer significant values.

	Nov 10			Nov 17			Nov 28			Dec 5		
	%	diff	p	%	diff	p	%	diff	p	%	diff	p
No Privacy Policy												
Warning	98.19	-0.46	**0.0004**	97.61	-0.55	**0.0002**	96.46	-0.66	**0.0003**	95.85	-0.62	**0.0018**
No Warning	98.41	-0.24	0.0601	97.69	-0.47	**0.0015**	96.46	-0.66	**0.0003**	95.85	-0.62	**0.0018**
Parsed	98.22	-0.42	**0.0000**	97.63	-0.54	**0.0000**	96.37	-0.75	**0.0000**	95.70	-0.76	**0.0000**
Generic	98.45	-0.19	0.0830	97.83	-0.33	**0.0057**	96.72	-0.40	**0.0064**	96.12	-0.34	*0.0381*
Control	98.65	–	–	98.16	–	–	97.12	–	–	96.46	–	–
No Consent												
Warning	96.15	-1.41	**0.0000**	92.23	-1.64	**0.0001**	90.75	-1.80	**0.0001**	89.45	-2.09	**0.0000**
No Warning	96.74	-0.82	**0.0039**	92.98	-0.89	**0.0337**	91.28	-1.28	**0.0056**	89.85	-1.69	**0.0005**
Parsed	96.38	-1.18	**0.0000**	92.68	-1.19	**0.0003**	90.99	-1.56	**0.0001**	89.90	-1.64	**0.0001**
Generic	96.94	-0.62	**0.0138**	93.03	-0.85	**0.0166**	91.63	-0.93	**0.0201**	90.15	-1.40	**0.0009**
Control	97.56	–	–	93.87	–	–	92.55	–	–	91.54	–	–
Before Consent												
Warning	97.59	-0.31	0.2194	94.71	-0.49	0.1959	93.20	-0.81	0.0529	92.49	-0.71	0.1064
No Warning	97.83	-0.07	0.8166	94.49	-0.72	0.0602	93.32	-0.69	0.0939	92.48	-0.72	0.0993
Parsed	97.70	-0.20	0.3509	94.81	-0.39	0.1504	93.53	-0.48	0.1035	92.64	-0.56	0.0847
Generic	97.81	-0.09	0.6328	94.58	-0.62	0.0658	93.22	-0.79	*0.0406*	92.61	-0.59	0.1381
Control	97.90	–	–	95.20	–	–	94.01	–	–	93.20	–	–
No HTTPS												
Warning	97.42	-0.54	0.0815	96.54	-0.59	0.0962	95.05	-1.00	**0.0164**	94.13	-1.09	*0.0155*
No Warning	97.75	-0.21	0.4980	97.05	-0.08	0.8602	95.45	-0.60	0.1452	94.68	-0.54	0.2263
Parsed	97.57	-0.39	0.0838	96.75	-0.38	0.2229	95.14	-0.91	**0.0081**	94.37	-0.85	**0.0143**
Generic	97.73	-0.23	0.4293	96.94	-0.19	0.4649	95.61	-0.44	0.1681	94.72	-0.50	0.1904
Control	97.96	–	–	97.13	–	–	96.05	–	–	95.22	–	–
Git												
Warning	91.18	-1.60	**0.0000**	89.91	-1.91	**0.0000**	86.85	-2.35	**0.0000**	86.34	-2.61	**0.0000**
No Warning	91.08	-1.70	**0.0000**	89.80	-2.01	**0.0000**	86.71	-2.49	**0.0000**	86.25	-2.71	**0.0000**
Parsed	94.04	1.26	0.1438	92.64	0.82	0.8899	90.01	0.81	0.7519	89.12	0.17	0.0961
Generic	90.52	-2.26	**0.0000**	89.38	-2.44	**0.0000**	86.27	-2.93	**0.0000**	85.99	-2.96	**0.0000**
Control	92.78	–	–	91.82	–	–	89.20	–	–	88.95	–	–

G LOGISTIC REGRESSION MODELS

Table 7: Logistic regression models for all issues. Figures without brackets denote the estimates and figures in brackets the standard error. *** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$.

	No Privacy Policy				**No Consent**				**Before Consent**			
	Nov 10	Nov 17	Nov 28	Dec 5	Nov 10	Nov 17	Nov 28	Dec 5	Nov 10	Nov 17	Nov 28	Dec 5
Intercept	-5.81 ***	-5.44 ***	-5.12 ***	-4.89 ***	-2.48 ***	-2.14 ***	-2.01 ***	-1.94 ***	-2.51 ***	-2.23 ***	-2.17 ***	-2.09 ***
	(0.15)	(0.13)	(0.11)	(0.10)	(0.05)	(0.04)	(0.04)	(0.04)	(0.05)	(0.05)	(0.05)	(0.05)
No warning	1.02 ***	0.93 ***	1.03 ***	0.90 ***	0.25 ***	0.18 **	0.25 ***	0.15 **	0.03	0.08	0.16 **	0.05
	(0.18)	(0.15)	(0.13)	(0.12)	(0.07)	(0.06)	(0.06)	(0.06)	(0.07)	(0.06)	(0.06)	(0.06)
Warning	1.24 ***	1.12 ***	1.23 ***	1.10 ***	0.31 ***	0.23 ***	0.27 ***	0.17 **	0.04	0.08	0.12 *	0.07
	(0.17)	(0.15)	(0.12)	(0.11)	(0.07)	(0.06)	(0.06)	(0.06)	(0.07)	(0.06)	(0.06)	(0.06)
Num. obs.	43,512	43,809	43,453	42,450	17,978	18,166	18,027	17,608	19,490	19,739	19,625	19,077
	No HTTPS				**Git**							
	Nov 10	Nov 17	Nov 28	Dec 5	Nov 10	Nov 17	Nov 28	Dec 5				
Intercept	-2.53 ***	-2.40 ***	-2.34 ***	-2.31 ***	-2.68 ***	-2.54 ***	-2.31 ***	-2.22 ***				
	(0.06)	(0.06)	(0.06)	(0.06)	(0.04)	(0.04)	(0.04)	(0.04)				
No warning	0.11	-0.04	0.07	-0.06	0.20 **	0.25 ***	0.32 ***	0.31 ***				
	(0.08)	(0.08)	(0.08)	(0.08)	(0.06)	(0.06)	(0.05)	(0.05)				
Warning	0.17 *	-0.03	0.12	0.12	0.17 **	0.19 **	0.28 ***	0.28 ***				
	(0.08)	(0.08)	(0.08)	(0.08)	(0.06)	(0.06)	(0.05)	(0.05)				
Num. obs.	12,226	12,304	12,174	11,915	23,664	23,654	23,538	23,498				

H CODEBOOK FOR EMAIL CLASSIFICATION

Table 8: Codebook for email classification. Numbers in the % columns are relative to the total number of email conversations about domains with security issues (n = 167) and privacy issues (n = 414), respectively.

Code	Description	Examples	Counter-examples	Requires	# of Conversations			
					Security		Privacy	
					n	%	***n***	%
Sentiment								
thanks	The recipient thanks us for the notification.	"Thank you for your notification"			125	74.9	232	56.0
great-project	The recipient expressed that they liked our project.	"Thank you for your work", "We need more projects and people like you!", "good luck with the project", "In case you find any other vulnerabilities I'd be extremely grateful if you would let me know"	"Thank you for your notification", "Many thanks for your two messages, including the valuable advice"		27	16.2	12	2.9
negative	The recipient did not like our project or our notification.	"PISS OFF!!!", "telling a UK business what to do is completely unacceptable", "stop sending threatening emails, it's stupid"	"I would like to be excluded from your project"		3	1.8	22	5.3
More information								
more-info	The recipient asks for more information about, e. g., our project, our checks, about the GDPR, about Git.	"Do you think the GDPR applies to us?", "What needs to be changed?"	"How are you?", "Can you fix this for us?", "Can you exclude us?"		42	25.1	131	31.6
privacy-check	The email contains a question about our privacy checks.	"How does your check recognize a privacy policy?"	"What must be included in a privacy policy?"	more-info	1	0.6	33	8.0
cookie-name	The email contains a question for a cookie name.	"Could you be so kind to specify name of the cookie you are referring to?"		privacy-check	1	0.6	22	5.3
git-check	The email contains a question about our Git check.	"Could you please provide more details about your findings and the actions performed by your automated [Git] analysis?"		more-info	18	10.8	0	0.0
git-url	The email asks for the URL of the Git repository or a file (e. g., config) inside the repository.	"At which URL have you been able to access the repository?"		git-check	9	5.4	0	0.0
project-info	The email contains a question about our research project in general.	"How have you selected our website?", "Is it also possible to trigger this check one way or another?"		more-info	8	4.8	20	4.8
state	The email asks if the issue is still present on the website.	"Could you check again?"		more-info	9	5.4	28	6.8
fix-this-plz	The recipient asks us to fix the issue on their behalf.	"How do you fix this? Can you do that?"		more-info	2	1.2	0	0.0
Performed actions								
fixed	The email states that the issue has (presumably) been fixed.	"This has been resolved.", "I've updated my nginx configuration to deny all access to '.' directories", "the security issue should be fixed now"			75	44.9	60	14.5
will-handle	The email states that the recipient will look into the issue or fix the issue in the future.	"I will arrange according to your advice", "You can assume that the website's communication will be encrypted within the next hours", "We will fix it asap"			27	16.2	157	37.9
notified	The recipient notified someone else in order to fix or look into the issue.	"I will get in touch immediately with the person that created our website", "I've forwarded your message to domain owner"	"We will handle this."		17	10.2	66	15.9

Continued on next page

Table 8 – *Continued from previous page*

Code	Description	Examples	Counter-examples	Requires	# of Conversations			
					Security		Privacy	
					n	%	***n***	%
Correctness								
false-positive	The recipient thinks that the current state of their website is secure / compliant.	"I thought GDPR does not apply to our website", "We do not gather any third party cookie data from visitors"			10	6.0	75	18.1
git-no-sensitive	The recipient thinks that the Git repository does not contain any sensitive data.	"the git repo doesn't contain any confidential information", "the repository is also published at URL"			14	8.4	0	0.0
laws-not-apply	The recipient thinks that the privacy laws do not apply to them.	"I thought GDPR does not apply to our website", "Our web page is not public", "We do not process personal data", "We do not have cookies for visitors acceptance and only visitors that subscribe newsletter provide their email"		false-positive	0	0.0	40	9.7
laws-not-in-uk	The recipient thinks that EU privacy laws do not apply to them because they are in the UK.	"Why do you contact an [sic!] UK business?"		laws-not-apply	0	0.0	5	1.2
Language								
expected-german	The recipient asks why we sent emails in English and not German, the language of our institutions' country.	"Why do you send an English email to a German as a German research institute?	"Feel free to also contact me in German "		0	0.0	6	1.4
translate	The recipient asks for a translation into another language (most frequently German).	"If you want to communicate with me, then please write in German!", "In German, please", "is there possibly a 'German version' of this email?"	"Feel free to also contact me in German"		1	0.6	12	2.9
Other								
unsure-scam	The recipient is unsure if the mail is spam / a scam.	"is this a real email or a phishing attempt", "This looks extremely suspicious to me in its content, tone, and method of delivery", "this looks like spam", "Is this a legitimate email?"	"somehow sounds legitimate"		12	7.2	50	12.1
really-cispa	The recipient is unsure if the email is really from CISPA.			unsure-scam	10	6.0	35	8.5
exclude	The recipient wants to be excluded from our study. Includes conditional exclusion requests.	"Either you call us or I have to ask you to exclude our website"			7	4.2	34	8.2