Computers & Security 160 (2026) 104682

Contents lists available at ScienceDirect

Computers & Security

journal homepage: www.elsevier.com/locate/cose

Full length article

"I believe it's incredibly difficult to fight against this flood of spam": Towards enhancing strategies for creating effective vulnerability notifications

Anne Hennig [a,*], Maxime Veit [a], Leoni Schmidt-Enke [a], Fabian Neusser [b], Dominik Herrmann [b], Peter Mayer [c]

[a] Karlsruhe Institute of Technology, Karlsruhe, Germany
[b] University of Bamberg, Bamberg, Germany
[c] South Denmark University, Odense, Denmark

ARTICLE INFO

Keywords:
Website hacking
Website compromise
Notification experiment
Website vulnerabilities
Web security
Redirect hack
Awareness
Vulnerability notification
Interview study

ABSTRACT

Identifying the most effective and scalable methods for notifying website owners about compromises or vulnerabilities remains an enduring challenge. Although some success factors have been identified, results regarding effective senders and notification framing are often inconsistent, and the understanding of how recipients perceive vulnerability notifications is still limited. Heading towards a better understanding, we conducted a 3×3 randomized controlled notification experiment, examining the impact of three distinct senders and three variations of notification framings for $n = 581$ compromised German websites. Our findings revealed a promising trend: receiving any notification significantly increased remediation compared to the absence of one. Remarkably, the choice of sender and framing played only a minor role in our notification experiment, which underscores the importance of notifying compromised websites and should motivate those who find vulnerabilities to take action. Yet, despite these encouraging results, a staggering 58% of the notified websites failed to remediate. To delve deeper into this phenomenon, we conducted follow-up interviews with 42 website owners who did not remediate their websites. The insights were revealing: while our notifications were delivered, many interviewees admitted they either overlooked or dismissed them as spam. This pattern persisted across different senders and framings, highlighting a critical challenge for future notification campaigns. Moving forward, future research should focus on finding ways to cut through the overwhelming amount of daily "spam" and explore strategies for how notifications can effectively convey their importance in recipients' inboxes. Exploring strategies to raise the general awareness for cybersecurity, encouraging website owners to provide a security.txt, or providing additional assistance in the form of a self-service tool, are some proposals to increase remediation rates. We further recommend that future work should consider theories from communication science or psychology, e.g., Protection Motivation Theory (PMT) or the Elaboration-Likelihood Model, when designing notification campaigns.

1. Introduction

It was reported that in 2022, each website had to endure, on average, 172 attacks per day (SiteLock, 2022). While some types of website attacks, e.g., defacement or Denial of Service (DoS) attacks, are immediately noticeable by website owners, other types, like cloaking[1] or unauthorized third-party redirect hacks,[2] are designed to remain hidden (Samarasinghe and Mannan, 2021). Attackers may have compromised a system long before the attack is finally detected. Even experienced developers might struggle to identify and remediate such unauthorized third-party redirect hacks, as described in BitofWP (2019). In a preliminary investigation, a wide variety of websites were found to be affected by these redirect hacks (mindUp Web & Intelligence GmbH, 2025a), with some cases remaining undetected for several months. This

* Corresponding author.
E-mail addresses: anne.hennig@kit.edu (A. Hennig), maxime.veit@kit.edu (M. Veit), leoni.schmidt-enke@kit.edu (L. Schmidt-Enke), fabian.neusser@stud.uni-bamberg.de (F. Neusser), dominik.herrmann@uni-bamberg.de (D. Herrmann), mayer@imada.sdu.dk (P. Mayer).

[1] Cloaking describes the technique of returning different versions of a website for search engine crawlers and users.

[2] Unauthorized third-party redirect hacks, as described in Section 2.1, are redirects placed on benign websites without authorization of the website owner to redirect users to malicious websites.

https://doi.org/10.1016/j.cose.2025.104682
Received 28 May 2025; Received in revised form 5 September 2025; Accepted 21 September 2025
Available online 21 October 2025

highlights the necessity to notify the website owners about the compromises. The primary goal of our study was to notify affected website owners and make them aware of their websites being compromised.

While searching for the most effective, yet feasible-to-scale method of notifying website owners about vulnerabilities, we found that previous research sometimes provides conflicting results, e.g., regarding the effect of the sender (Çetin et al., 2016; Maaß et al., 2021c). Furthermore, most studies have focused on specific settings, such as using only Computer Emergency Response Teams (CERTs) (Lone et al., 2022) or hosting providers (Çetin et al., 2019b) as senders. The influence of incentives is also still unclear (see Section 2).

To address these limitations, our study is the first to combine factors from related work that positively affect remediation into a single study design. We are the first study that partnered with two German hosting providers and the CERT of the German Federal Office for Information Security (Federal CERT) to sent out notifications directly to website owners. We conducted a mixed-methods investigation, combining a notification experiment to examine the effects of sender and framing of a notification, with qualitative interviews to explore reasons for non-remediation. Based on the results of our experiment and the feedback from affected website owners, we provide design recommendations for practitioners and researchers involved in notification campaigns. Additionally, we outline a comprehensive path for future work in the area of vulnerability notifications.

Our work was guided by the four research questions described below. We motivate these in detail from the context of the related work in Section 2.

RQ 1 [Framing] *Which framing has what impact on the remediation of compromised websites?*

We analyzed the effects of three incentives that previous work either proposed (reputational incentives (Vasek and Moore, 2012; Çetin et al., 2017)) or investigated (technical incentives (Çetin et al., 2019a,b), no incentives). Our goal, thereby, was to identify whether providing consequences, other than legal incentives which are only applicable in certain circumstances (Maaß et al., 2021c), has a significant effect on remediation. To the best of our knowledge, reputational incentives as well as a direct comparison of different incentives within a single study have not been investigated before. We did not find any significant differences in the remediation rates between the three framings, indicating that for vulnerability notifications, senders can choose any of the framings in our study, whichever they feel most comfortable with (Section 4.1).

RQ 2 [Sender] *Which sender has what impact on the remediation of compromised websites?*

We investigated the effect of three senders (university, hosting provider, Computer Emergency Response Team (CERT) of the German Federal Office for Information Security (BSI)) in comparison to one another, which have only been investigated separately in distinct studies in previous work. Further, we directly partnered with contact persons in the respective entities to make sure that notifications are sent. Our goal, thereby, was to provide clear evidence on the effects of different senders. We found that all notified websites exhibited significantly higher remediation rates and fewer days to remediation than the unnotified control group. However, we did not find a significant difference between the three senders university, hosting provider, Federal CERT (Section 4.1), indicating that the sender appears to have no impact on remediation.

RQ 3 [Framing × Sender] *Can we identify an interaction effect between sender and framing with respect to the remediation of compromised websites?*

We aimed to identify whether we can observe interaction effects between any of the sender and framing conditions in our study, e.g., whether a technical framing was more effective when sent out by the hosting provider. Previous work indicated such interaction effects, albeit for a combination of legal sender and legal framing (Maaß et al., 2021c; Utz et al., 2023). All senders sent out notifications for all framings. However, we did not find significant differences between the different framings when correlated with a corresponding sender, further supporting the notion that sending notifications is the crucial part, while framing and sender matter less (Section 4.1).

RQ 4 [Reasons] *What are reasons for website owners to not remediate their websites even after being notified twice?*

As related work has revealed surprisingly low remediation rates, we additionally conducted qualitative interviews with website owners, rather than using static surveys, as in Durumeric et al. (2014), Li et al. (2016a), Maaß et al. (2021c), Stock et al. (2018), StopBadware and Commtouch (2012), Zeng et al. (2019), Çetin et al. (2017) and Lone et al. (2022) to gain in-depth insights regarding *reasons for non-remediation behavior* (in contrast to *remediation strategies* as analyzed in Rodríguez et al. (2021)). We identified a variety of reasons why website owners did not remediate the unauthorized third-party redirect hacks. Most prominent are the visibility of vulnerability notifications among other emails and the website owners' perceived relevancy of the compromise for their website (Section 4.2).

2. Related work

A wide range of work has been conducted in the area of communicating security risks, particularly notifying users about potential security issues (e.g., password breaches (Mayer et al., 2021), password reuse (Golla et al., 2018; Albayram and Walker, 2024), data breaches (Zou et al., 2019; Huang et al., 2022), identity theft (Muniz et al., 2024), or IoT hacking (Rostami et al., 2022), etc.). While these findings are relevant to our work, we will focus on notification experiments that explicitly notified *website owners* about potential attacks on their websites in the following. In particular, we focused on unauthorized third-party redirect hacks, as this compromise has not been extensively investigated, and we are unaware of any studies that have raised awareness for this compromise among website owners. In our experimental design, we focused on the effects of *sender* and *framing*, as well as *interaction effects* between the two. This approach was motivated by contradictory results or research gaps identified in the existing literature. Additionally, we collaborated with specific contact persons in the Federal CERT and the hosting companies to address the limitations of previous work (Lone et al., 2022; Nosyk et al., 2023; Kührer et al., 2014; Li et al., 2016b). We agreed on a common cover letter that each sender used, and ensured that all framings were used by each sender, and all notifications were sent.

2.1. Threat model

Unauthorized third-party redirect hacks, as investigated in this paper, involve hijacking legitimate websites to lead users to malicious sites, such as fake online pharmacies or fraudulent casinos (mindUp Web & Intelligence GmbH, 2025b). These unauthorized third-party redirect hacks are stealth attacks, meaning they are designed to hide inside an otherwise benign system (Cazorla et al., 2018). The redirects are only active when users access the website through search engine results that advertise the malicious websites. The original website can still be accessed normally, i.e., no redirect takes place. At first glance, the website may appear uncompromised, but when searched using the "site:" - operator, unusual entries are visible in the search results, redirecting the visitor to an illegitimate online shop (see Fig. 1 for an example). This behavior is often referred to as (malicious) search

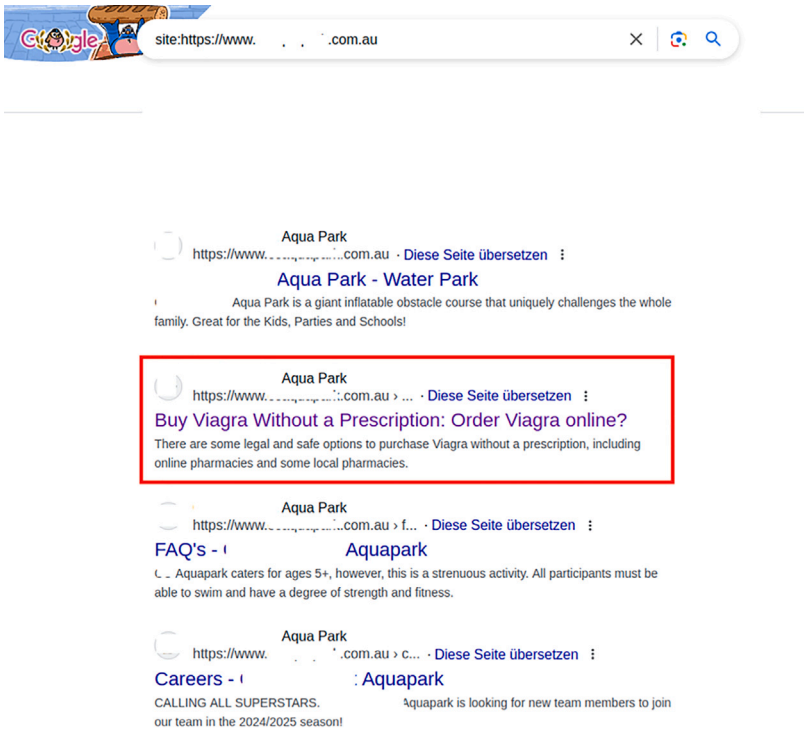

Fig. 1. Search results of a compromised website. While there are mostly legitimate results, the highlighted link would redirect the user to a fake pharmacy selling Viagra.

engine Spam or SEO Spam (Martori, 2020), website hijacking (Li et al., 2016b), Pharma Hack (BitofWP, 2019), Japanese Hack (Goodchild, 2024), or WordPress Hack (Sundaram, 2022). All these attacks describe different variations of the same underlying problem: the ranking of a legitimate website is maliciously used to redirect to other websites. However, in this paper, we specifically investigated a variation in which the legitimate website is compromised. To avoid using a term with an unclear or misleading definition, we refer to the compromise as "unauthorized third-party redirect hacks".

While seemingly "harmless", these redirects can be seen as indicators of compromise, as mentioned in Kumar et al. (2019) and described in Halder (2025), since an attacker needs write access to the website data to place the unauthorized third-party redirect hack. Soska and Christin (2014) describe risk factors for websites *before* they become malicious (e.g., for phishing or distributing malware), and based on our findings, we would add irregular search engine entries as such risk factors. As this specific unauthorized third-party redirect hack has not been analyzed on a large scale, there is only case-by-case evidence of how attackers infiltrated the systems and the extent of control they have (e.g., Halder, 2025; Martori, 2020; BitofWP, 2019; Sundaram, 2022; Goodchild, 2024).

2.2. Motivating research questions

Notification content and framing. There are several design factors, such as general notification design (Stock et al., 2018), translation of notification texts into native languages (Zeng et al., 2019; Li et al., 2016b), or variations in subject lines (Zeng et al., 2019) that were studied in previous work, but showed no significant effects on remediation rates. On the other hand, prior research has shown that providing detailed information in the notifications (Li et al., 2016a,b; Stock et al., 2018; Vasek and Moore, 2012; Zeng et al., 2019; Çetin et al., 2016) as well as describing the security issue and highlighting its importance (Zeng et al., 2019) increases the effectiveness of a notification.

Regarding the framing of the notification, Zeng et al. (2019) found that pointing out consequences, i.e., framing the attack as more or less severe, has no impact on remediation rates. However, other authors suggested the opposite. Çetin et al. (2019a,b) could show that technical incentives, such as quarantining infected websites, can increase remediation rates. Maaß et al. (2021c) found that providing legal incentives, in the form of fines, can significantly increase remediation rates. Reputational incentives, such as declining search rankings (Vasek and Moore, 2012) or publicly naming compromised websites (Çetin et al., 2017), have been proposed to increase remediation rates; however, their effectiveness has not been researched yet.

Due to conflicting results or missing research about incentives and message framing in the literature, we wanted to focus on this aspect in more detail. Our goal is to compare the effects of technical, reputational, and no incentives for the same problem within a single study, providing further insight into how to frame effective notifications. Thus arose our research question **RQ 1**: "*Which framing has what impact on the remediation of compromised websites?*"

Sender of the notification. Çetin et al. (2016) notified website owners via email from three different senders with varying reputations (an individual researcher, a university, and an anti-malware organization). Although the authors found a significant difference between the control group and the treatment groups, they were unable to find a significant difference between the three senders. The authors concluded that the remediation rate could neither be improved by choosing a certain sender, nor was the willingness to remediate affected by the sender's reputation (Çetin et al., 2016). Later studies could also not identify statistically significant differences between the respective sender groups (e.g., individual researcher vs. research group (Stock et al., 2018), researcher vs. Google Search Console (Zeng et al., 2019), (inter)national CERTs (Kührer et al., 2014; Gasser et al., 2017; Lone et al., 2022), or hosting provider (Zhang et al., 2017; Çetin et al., 2019b, 2018; Rodríguez et al., 2021; Çetin et al., 2019a; Bouwmeester et al., 2021; Kührer et al., 2014)).

In contrast, Maaß et al. (2021c) showed that notifications with a legal framing sent from the university's law group led to significantly higher remediation rates, compared to a university computer science

group, indicating that either sender or framing (or a combination of both) might make a difference. Since a vulnerability notification always requires a sender, who potentially influences the perception of the notification, our goal is to complement existing results by comparing different senders with high reputation that were used in previous studies but not yet evaluated in comparison to each other. This motivated our research question **RQ 2**: "*Which sender has what impact on the remediation of compromised websites?*"

Maaß et al. (2021c) and Utz et al. (2023) found that a framing referencing fines, ideally in combination with a sender that has the authority to impose such fines (Maaß et al., 2021c) leads to higher remediation. Thus, we also wanted to investigate whether we can transfer their results to other framing–sender combinations, identifying whether a specific sender is more successful with a specific framing, which could then be a further design factor for future notifications. Thus arose our research question **RQ 3**: "*Can we identify an interaction effect between sender and framing with respect to the remediation of compromised websites?*"

Remediation rates. Previous research has shown that notifying website owners increases remediation rates compared to non-notification scenarios (Durumeric et al., 2014; Kührer et al., 2014; Li et al., 2016a,b, 2019; Maaß et al., 2021a,c; Stock et al., 2018, 2016; Vasek and Moore, 2012; Zeng et al., 2019; Çetin et al., 2017, 2019b, 2016). However, remediation rates, in general, turned out to be low. For website-related vulnerability notifications, Stock et al. (2018) achieved a remediation rate of 24% when notifying website owners about publicly accessible Git repositories, and an even lower remediation rate of 17% when notifying about cross-site scripting vulnerabilities in WordPress. Maass et al. achieved a medium remediation rate of 56.6% ranging from 33.9% to 76.3% (Maaß et al., 2021c) depending on the condition[3] when informing website owners about missing IP anonymization while using Google Analytics. Zeng et al. reported a remediation rate between 7% and 34% when notifying website owners about different HTTPS misconfigurations.

Additionally, Zeng et al. (2019) also observed and notified website owners about soon-to-be distrusted Symantec certificates. While in this case, notifications had no statistically significant effect compared to not notifying, remediation reached an outstanding rate of 90% after 40 days across all groups. Durumeric et al. (2014) also reached a relatively high remediation rate of around 57% when notifying website owners about the infamous Heartbleed vulnerability. For both experiments, the authors acknowledged that the prominence of the respective issues had a huge impact on the high remediation rate (Durumeric et al., 2014; Zeng et al., 2019).

Stöver et al. (2023) identified further reasons why website owners would not remediate the misconfiguration by analyzing the email and survey responses that Maaß et al. (2021c) got during their notification experiment. Among the most prominent reasons were lack of awareness for the problem, lack of technical knowledge, or lack of resources, such as time. Other reasons included deliberate lack of maintenance, ambiguous responsibilities, and complex organizational structures, which slow down or hinder remediation processes. As a result, we recognized the need to not only measure remediation rates as an indicator for the success of our notification campaign, but also to identify reasons why website owners notified about unauthorized third-party redirect hacks remain inactive. In doing so, we aim to clarify the considerable variation in the remediation outcomes reported in the literature. This motivated our research question **RQ 4**: "*What are reasons for website owners to not remediate their websites even after being notified twice?*"

[3] The authors compared twelve different conditions: university law group, university computer science group, and an individual researcher as senders; privacy, GDPR, and GDPR with fine as framings; as well as email and letter as notification channels.

2.3. Related work informing design decisions

Notification channel. While most studies used email notifications (Durumeric et al., 2014; Hennig et al., 2022a; Kührer et al., 2014; Li et al., 2016a,b; Maaß et al., 2021a,c; Stock et al., 2018, 2016; Vasek and Moore, 2012; Zeng et al., 2019; Çetin et al., 2017, 2019b, 2016), some also tried other channels, such as Google Search Console messages, which were not more effective than email (Zeng et al., 2019). In addition, letters and phone calls, as well as social media, performed better than email (Stock et al., 2018). Letters have also proven to be more effective than emails in Maaß et al. (2021c). This results in the alternative channels having a slightly higher remediation rate, at the cost of manual effort to retrieve the addresses (Stock et al., 2018), print and envelope the letters, as well as the postage fees (Maaß et al., 2021c). Although sending letters proved to be less efficient, manually retrieving email addresses seems to pay off. Previous studies have shown that email bounces can be decreased (Stock et al., 2018) and deliveries can be improved by manually retrieving (Hennig et al., 2022a; Maaß et al., 2021c; Stock et al., 2018) or automatically crawling (Utz et al., 2023) for email addresses compared to using WHOIS (Stock et al., 2018, 2016; Zeng et al., 2019) or generic emails (Stock et al., 2016; Çetin et al., 2016; Utz et al., 2023). *Thus, we decided for email as the notification channel and manually collected email addresses from the websites' imprints.*

Notification content and study design. We also derived specific design decisions for our notification text and study design from interviews with $n = 25$ website owners affected by unauthorized third-party redirect hacks, that were conducted prior to this study. The methodology is not described in this paper, as further information can be found in the original publication (Hennig et al., 2022b). As both studies are based on the same type of unauthorized third-party redirect hacks, we found it necessary to draw our design decisions from our pre-study specifically.

During the interviews, most participants expressed a general distrust in vulnerability notifications. Asked to identify a suitable sender, only a few interviewees could intuitively name a sender they considered appropriate. Among specific senders that were regarded suitable, the police was named 12 times, hosting provider was named nine times, research facilities were named six times, and the Federal Office for Information Security (BSI) was named four times (multiple answers permitted). A summary of the results is provided in Appendix A.2, Table 3.

Regarding suitable notification channels, 17 interviewees identified email as the most appropriate. Email notifications are valued for speed, ease of use, and practicality. Two interviewees also stated that email would be the most logical notification channel for a digital problem. On the other hand, five interviewees explained they would be somewhat suspicious of notifications via email, but no one explicitly refused to be notified via email. Phone calls were deemed most suitable by 11 interviewees, and two indicated that they would prefer either a letter or a dedicated web portal.

Concerning the content of a notification, the results supported previous research (Li et al., 2016a,b; Maaß et al., 2021c; Stock et al., 2018; Vasek and Moore, 2012; Zeng et al., 2019; Çetin et al., 2016). In our pre-study, interviewees especially emphasized that a clear description of the attack, a clear motivation for the notification, and, if applicable, instructions on how to solve the malicious redirect should be included in a vulnerability notification. Furthermore, providing contact information – such as a phone number or email address, a signature, a letterhead, or an imprint – helps the recipients verify the sender. Also, four interviewees said that they would appreciate a personalized salutation. Two interviewees stated that they pay attention to correct orthography and spelling, while two others requested a meaningful subject. All these factors would make a vulnerability notification credible and comprehensible, thereby raising awareness for the legitimacy of the described attack. But awareness does not automatically lead to remediation: Two interviewees stated that they deemed the described compromise negligible, even if the notification was credible.

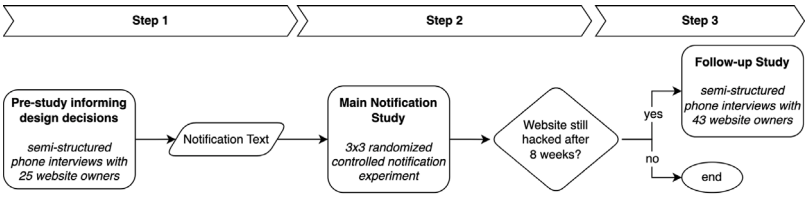

Fig. 2. General structure with all steps of our study design. Italic indicates the method used in each step.

Implications for our study design. We considered all the results from our pre-study and related work, as described above, in designing our main study (see Section 3.2 for details). Firstly, we defined suitable senders and notification channels for our notification experiment based on our pre-study (Hennig et al., 2022b). Since the police was named as a suitable sender most frequently, we pursued avenues to cooperate with cybercrime divisions of the German police as senders for our notification experiment. However, this ultimately proved impossible (see Section 3.3) and we had to settle for the next three most frequently named options: university/research institution, hosting provider, and BSI. Secondly, we chose email as notification channel. Most of the interviewees either named email as a suitable notification channel or were not strictly opposed to email. We found this to be the most cost-effective solution and, therefore, the most likely to be relevant in practice. We also decided to retrieve email addresses from the imprint of the websites manually. Thirdly, we made clear note of the individual features that the interviewees emphasized as necessary for the notification text, e.g., a clear description of the attack, and used them in designing the notification text. Finally, as it was found that website owners did not take the unauthorized third-party redirect hacks seriously, we deemed it necessary to point out consequences.

3. Methodology

To run our experiment, we utilized data from a web crawling service that identified websites affected by unauthorized third-party redirect hacks (Section 3.1). We then implemented a technique to continuously monitor the remediation status of the identified websites independently of the crawling service to find the compromised domains (Section 3.1). We used the results of our pre-study (Hennig et al., 2022b), which complemented the best practices for vulnerability notifications from the literature, to make informed design decisions for our study design (Section 3.2).

Based on the findings of previous work, we designed a quantitative 3×3 randomized controlled notification experiment to investigate **RQ 1**–**RQ 3**, analyzing the effects of three different senders (university, hosting provider, Federal CERT) and three different framings (neutral, technical, reputation) (Section 3.2.1). To answer **RQ 4**, we designed qualitative follow-up interviews with website owners who had not remediated the unauthorized third-party redirect hack within 56 days after our initial notification (Section 3.2.2). The overall structure of our study design is shown in Fig. 2.

We end this section by discussing our ethical considerations (Section 3.3), the limitations of our work (Section 3.4), and our data analysis methods (Section 3.5). As will be described in Section 3.3, all parts of our experiment, including notification texts, interview guideline, and notification process, have been approved by the ethics board of our university.

3.1. Background

Sampling compromised websites. We obtained a dataset of websites affected by unauthorized third-party redirect hacks through a service specializing in large-scale web crawling (mindUp Web & Intelligence GmbH, 2025a). No specific datasets (e.g., Tranco lists) were used to identify affected websites. Instead, the crawling identified affected websites based on public search engine results using keywords for the search. The service employed a multi-step process. First, they conducted a web search using a broad range of keywords known to be used to attract victims to malicious target websites. The keywords range from the shopping context, e.g., well-known brands, to topics such as Bitcoin or casino-related terms, and were defined by the crawling service (see Appendix A.1, Table 2 for a translated list of keywords of the websites in our sample). The URLs that were found, using the keywords-based search, were analyzed automatically to check known indicators of compromise, e.g., variations in website topics (like pharmaceutical content on a car seller's website), unusual phrases (like brand names of luxury watches on a school's website), page behavior with and without search engine identification, and other specific behaviors which the crawling service found typical for unauthorized third-party redirect hacks. Based on these indicators, the service compiled a list of potentially compromised websites every month throughout the notification experiment (i.e., 20 months). We received these monthly lists, which contained a minimum of 47 and a maximum of 926 worldwide domain names, along with the respective Google search results URLs. We used our monitoring system to check the compromised websites. We received no further demographic data for the websites (e.g., popularity, sector, or business size). We did not consider it useful to classify the websites ourselves (e.g., by industry branch or sales revenue), as these information were not available for all websites in our sample, and would, thus, mainly be drawn from the authors' assumptions based on publicly available information (e.g., the website or, if applicable, trade registers). Overall, the websites in the monthly lists were very diverse, comprising small and medium-sized enterprises, entrepreneurs, associations, projects, schools, or universities. For each month and each website, the presence of unauthorized third-party redirects was then verified by our monitoring system before the website was added to our sample.

Monitoring websites during experiment. We used a monitoring system to regularly check if the unauthorized third-party redirects remained active during our 56-day observation period for each website.[4] The monitoring system was implemented in Python, checking the presence of the redirects every six hours. It accessed each website by simulating a click from a Mozilla Firefox browser on Google search results, which were also provided by the crawling service, using the HTTP headers *Referer* and *User-Agent* respectively. We implemented several measures to counteract cloaking: a user agent string from a standard browser (Mozilla Firefox 97 on Windows 10, the most common variant found at the time of our study), an IP from a university client pool instead of a datacenter server IP, HTTP referrer header from search engines, and manual checks. We manually checked that our detection of the redirect worked as intended for each of the websites at the beginning of each website's 56-day observation interval. Our monitoring system then recorded whether a website's existing redirect behavior changed (e.g., a website's redirect disappeared due to the unauthorized third-party redirect hack being fixed).

[4] Note that since we obtained a new list every month, the start time for the observation period varied for all websites. However, each website was monitored for at least 56 days.

If the website could not be accessed or access took more than 30 s, the access was retried four more times in quick succession before it was assumed that the website was taken down. The following rules were applied:

(1) It is assumed that the website is *still compromised*

- if it contains the respective keyword, or
- if there is a redirect to another website using HTTP redirect, JavaScript window location, or meta refresh.

(2) It is assumed that the website is *no longer compromised*

- if an HTTP client or server error response indicates that the page has been removed, or
- if the website has no page content, indicating that the malicious content has been removed, or
- if the domain can no longer be resolved through the Domain Name System (DNS), indicating that the website has been taken down.

(3) If none of the above applies, no automatic decision can be made, and the monitoring system flags the website status as *unknown*, which triggers manual review.

3.2. Study design main study

Notification channel. As described in Section 2.3, we decided to use email as the channel for our notification experiment. We purposely excluded other notification channels from our experiment, despite possible alternatives such as postal letters, social media, or phone calls (Maaß et al., 2021a; Stock et al., 2018) potentially being more effective. None of these were practically possible, since none of the external senders identified as relevant in our pre-study (hosting provider and the Federal CERT) were able to integrate them into their processes for large-scale notification campaigns. Furthermore, as discussed by Stock et al. (2018), the success rate of alternative communication channels was low compared to the effort and costs. It cannot, therefore, be considered cost-efficient, especially when notifying an even larger number of websites – Stock et al. (2018) notified a total of 264 websites, 173 of them via alternative channels, including postal letters (67 websites), web forms (69 websites), social media (91 websites), and phone (46 websites). The results of our pre-study also supported this design decision. Thus, using email was the best choice, considering both the application of our research results at scale in practice and the notification channels accepted by website owners. To ensure that our email notifications are delivered, we followed recommendations from previous work (Hennig et al., 2022a; Maaß et al., 2021b; Stock et al., 2018) and manually collected contact information from the imprint of each website. Specifically, we visited each website in our sample manually, without the aid of any automated script, and collected email addresses and the names of the persons responsible for the website from the imprint. Furthermore, we ensured that our email servers are configured according to state-of-the-art best practices, with valid SPF, DKIM, and DMARC records set to minimize the risk of our notifications being marked as spam. We purposively did not include any tracking mechanism (as Stock et al., 2018 did to measure reachability) to decrease the likelihood that our emails would be filtered out.

Notification text. We developed the text for our vulnerability notification based on the results of our pre-study and related work (see Section 2.3). We started with the most personal salutation possible, i.e. "Dear Mr./ Ms. [Lastname]". Next, we provided information on our motivation and on how to verify the unauthorized third-party redirect hack via the "site:"-operator in the search engine. We also explained that this operator will list all search engine entries of their domain, revealing the malicious redirects. Then, we either closed with a request to remediate the unauthorized third-party redirect hack and provided a link to our project website for further information, followed by our name, and a footer (neutral framing). Alternatively, we asked website owners to remediate the unauthorized third-party redirect hack, provided a link to our project website for further information, and subsequently included a technical or reputational incentive (technical or reputational framing). We also closed with our name and a footer that provided the sender's contact information. Since Maaß et al. (2021c) proposed reminder notifications to increase awareness, we decided to send one reminder notification to those website owners who did not respond or remediate within two weeks after our initial notification. The reminder email had the same content as the original email, but the subject line was changed to "Reminder: Important information on your website *'domain'*". All notifications were sent only in German and addressed to German domains (see Section 3.2.1). Both the text and the framings were reiterated with the hosting providers and the Federal CERT to ensure that we use a wording that all senders can identify with. We provide the translated text and the framings in Appendix B.1 and B.2.

Framings. Regarding consequences, previous studies suggested to either provide technical incentives (like quarantining a vulnerable website at least temporarily (Çetin et al., 2018, 2019b)) or pointing out negative consequences (like legal prosecution (Maaß et al., 2021c) or reputational damage (Vasek and Moore, 2012; Çetin et al., 2017)) to raise awareness. However, the effectiveness of reputational framing, as well as the different framings in comparison to each other, has not been analyzed yet. To test whether – and if so which – incentives can increase remediation rates, we compared three different framings that could either be used by all senders or were phrased to fit a particular sender (see Table 1 for an overview of the different groups, and Appendix B.2 for the wording of the framings): (1) a neutral framing with no incentives (`neutral`); (2) a technical framing stating that the website can be blocked by the hosting provider (`technical-generic`), the website will be suspended after a specific date (`technical-hoster1`), or further actions will be taken by the hosting provider (`technical-hoster2`); (3) a reputational framing stating that the website suffers reputational damage (`reputation-generic`), or that web reputation services might suspend the website based on negative reputation scores (`reputation-CERT`). While we did not explicitly consider behavioral theories when designing the framings, the wording is guided by Protection Motivation Theory (PMT). All framings are inspired by related work, although none of them is exactly like previous ones, as the content needed to be adapted to our specific type of compromise. Furthermore, the technical and reputational consequences had to be iterated with the hosting providers and the Federal CERT in order to find a phrasing that all senders could approve of on behalf of their institution. In particular, the decision to announce the take-down of websites was suggested by the respective hosting provider, as this is part of their regular notification process. Note that all senders sent out notifications using each of the framings, contrary to Lone et al. (2022). Based on the remediation for each of these framings, we can answer **RQ 1**.

Sender. To answer **RQ 2**, we partnered with dedicated contact persons from IONOS and GoDaddy, two major hosting providers in Germany, as well as the BSI (German Federal CERT), serving as senders. Additionally, we sent notifications ourselves for the university condition. In total, we had four senders, i.e., one person in each of the four entities. The same person always sent the notifications to achieve homogeneity within the groups. All senders can also be related to one of the framings, similar to Maaß et al. (2021c), where the legal framing was related to the law group. Hosting providers can be associated with technical framing due to their oversight of the infrastructure on which the compromised websites are hosted, and their resulting technical authority. The Federal CERT can be related to a reputational framing, as it maintains public lists of malicious websites and could potentially add affected websites to these lists. The university relates to the neutral

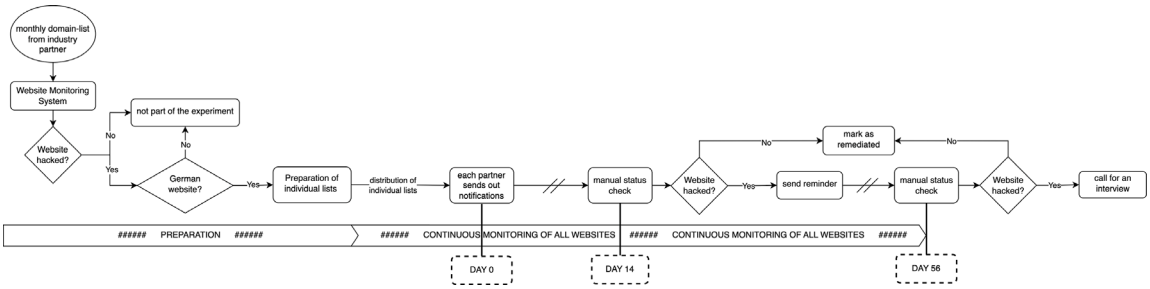

Fig. 3. Procedure RQ 1 - RQ 3.

Table 1
Final number of notifications sent per sender and framing incl. control group.

Sender	Framing	# notifications
Federal CERT	Neutral	73
	Technical-generic	70
	Reputation-CERT	69
	Sum	**212**
Hosting provider	Neutral	54
	Technical-hoster 1 & 2	54
	Reputation-generic	55
	Sum	**163**
University	Neutral	68
	Technical-generic	68
	Reputation-generic	69
	Sum	**206**
Control	–	**205**

framing, with neither apparent technical nor reputational authority. Each of the three senders sent out email notifications using each of the three framings (see Table 1 for an overview of the different groups), allowing us to answer **RQ 3**. Some websites were not notified as a control condition to allow a baseline comparison of the notifications (see Section 3.2.1).

3.2.1. *Procedure RQ 1–RQ 3: Framing & Sender*

The notifications were sent out monthly over 20 months between April 2022 and November 2023. Each month, we processed a completely new list obtained from the crawling service, totaling 20 lists with between 47 and a maximum of 926 domains from around the world per month. Websites that were already detected in any of the previous months or that had already been called (see Section 3.2.2) were excluded from the sample. Websites that were already remediated, non-German, or lacked valid contact information were also excluded from the sample. The procedure each month was as follows (see also Fig. 3).

(1) We obtained a domain list with vulnerable websites detected in the previous month from the crawling service. This domain list was then uploaded to the monitoring system to test all websites for their most recent status (see Section 3.1).
(2) All websites showing indicators of unauthorized third-party redirect hacks were manually checked for language and headquarters by the authors. For German websites, email addresses and recipients' names were then manually collected from the imprint. Websites from other countries were not included in the sample.
(3) Websites with valid contact information were then allocated semi-purposively to individual lists for each of our partners. The two hosting providers were assigned their own customers based on their Autonomous System Number (ASN). The remaining websites were then assigned to the Federal CERT, the university, or the control group. Note that websites, where it seemed critical to inform them about the vulnerability (i.e., kindergartens, schools, hospitals, doctors, lawyers) were never allocated to the control group.
(4) Finally, each website from the treatment groups was randomly assigned to one of the three framings. The framings were distributed evenly to all senders throughout the study. Meaning that if we could not evenly distribute the three framings to the number of domains we found for each partner in one month (i.e., when we had a number of domains that could not be divided by three), we then balanced it out in the following month.
(5) The individual lists, as well as further instructions, were sent to our contact persons at the hosting companies and the Federal CERT, respectively. Email notifications were sent by all senders every Wednesday in the third week of each month (Day 0). Two weeks later (Day 14), a reminder email was sent to websites that were still compromised.

3.2.2. *Procedure RQ 4: Reasons*

To answer **RQ 4**, we, as researchers from the university, called website owners who were still compromised eight weeks after the first notification email was sent out (Day 56). In the initial phone call, we invited them to a follow-up interview to get more in-depth information on why they had not remediated. Again, we processed each month's list individually to ensure that we do not contact website owners from any previous list. The phone number was retrieved from the imprint or the contact information given on the website. Website owners were contacted a total of no more than three times by one of the authors if we were unable to reach them on the first call. In the call, we specifically asked for the person who received and processed our notification. If the person did not agree to an interview, we briefly noted down what they told us regarding our notification (see Appendix C.3 for the categories we used to code the phone calls).

If the person agreed to participate in an interview, a return call was scheduled. We then conducted semi-structured telephone interviews. Note that all interviews were conducted in German as it was the first language for both participants and the interviewer. The ethics committee of our institution approved the interview guideline. Informed consent was obtained before the interviews. The participants were not compensated, and the interviews took 15 min on average, depending on the amount of information the interviewees recalled (see Appendix C.1 for the translated interview guideline, including informed consent). All interviews were recorded and manually transcribed verbatim without using any transcription AI, while anonymizing personal data in the process. Only the anonymized transcripts were analyzed. For coding the interviews, we applied open coding with three coders. The lead researcher developed a first draft of the codebook based on the interview guideline. After that, three of the authors independently coded three interviews and then met to discuss new codes and disagreements within existing codes. Thus, while the development of the initial codebook followed a deductive approach, further codes were added in an inductive approach. This process was repeated twice. After 23% of

the interviews had been independently coded by all three coders, and an IRR of $\kappa = 0.82$ was reached in the last iteration, the remaining interviews were coded independently by only two coders. However, all three coders met regularly to discuss any questions that arose. In the end, all codings were checked by the lead author.

3.3. Ethical considerations

All parts of our study were approved by the ethics committee of our institution. There were no ethical concerns about the permissibility of our research. In designing our study, we carefully considered any critical factors reported in related work. Firstly, neither we nor the crawling service we used exposed the websites or web servers to any risk, e.g., through performing any form of attack (in contrast to, e.g., Wu and Lu, 2021). We only sampled websites that showed indicators of being compromised as explained in Section 3.1.

Secondly, we carefully discussed the necessity and permissibility of a legal framing, i.e., a reference to legal obligations to remediate the hacking. As mentioned in the context of the Princeton study (Princeton University, 2021; Princeton, 2021), legal framings can cause fear and anxiety among recipients, and may lead to anger about the notification and mistrust in the senders. Since legal framings, i.e., referring to the General Data Protection Regulation (GDPR) and potential fines, had proven most effective in other notification studies (Maaß et al., 2021c; Utz et al., 2023), we discussed with several legal entities and law enforcement agencies (e.g., several cybercrime divisions of the polices in different states, a university law group, and contact persons at the Federal CERT) the feasibility of a legal framing similar to the one used by Maaß et al. (2021c). However, due to German legislation, this was impossible. There were only vague grounds for holding affected websites liable in case of unauthorized third-party redirect hacks, and prosecution varied significantly across the federal states. Therefore, we deliberately abandoned this condition – even if it limits the generality of our results – so as not to put legal pressure on the recipients or cause them unnecessary stress.

Another potential issue was the processing of personal data, including names, phone numbers, and email addresses. It was not possible to obtain consent for processing this personal data in advance without contacting participants at least once. Doing so would have compromised the integrity of our experiment. For contacting the affected websites, we used only publicly available data. In our email notification, we linked to our project website to connect the notification with the research project. A proper debriefing took place when we were able to reach the website owners via phone. Informed consent was obtained from all interviewees (see Appendix C.1). During the interviews, the participants were not exposed to any physical or mental risk at any time. While all personal data were deleted at the end of the experiment at the latest (e.g., personal data in the transcripts was removed during transcription), the domain names of the websites are retained for ten years as part of our sample documentation, in accordance with guidelines for proper scientific practice.

3.4. Limitations

Internal validity. The internal validity of our results may be compromised by sampling bias. As described in Section 3.1, compromised websites were identified by searching for specific keywords. This approach disregards websites that might be compromised by unauthorized third-party redirect hacks, but cannot be identified via the selected keywords. To mitigate this bias, the crawling service we used updated its crawling model to include new keywords over time. We must acknowledge that we did not have control over the keywords used for the crawling, which means that there are likely affected websites that were not found. However, this would only affect the potential sample size and would not impact our results. As described in Section 3.5, we required a sample size of 400 websites for our statistical analysis, which we successfully achieved.

We took countermeasures against cloaking before we added the websites to our sample (see Section 3.1). Nevertheless, there is a potential risk that attackers could have enhanced their cloaking technique during the 56-day observation period, enabling them to evade our countermeasures. This would result in false positives, where "compromised" websites are erroneously assumed to be "not compromised" by our monitoring system. However, random manual checks indicated that such behavior is unlikely to have occurred; therefore, we are confident that this has not affected our results.

Another limitation concerns our follow-up interviews. Our sample only consists of people who answered our calls *and* agreed to conduct an interview. Website owners who chose to participate in our interviews might systematically differ from those who did not, potentially leading to self-selection bias. Thus, our interview data likely represent a small and non-representative sample of website owners. Furthermore, our data might be affected by recall biases. However, we reached code saturation and gained a wide range of answers, allowing us to identify trends. Thus, we are confident in the value of our data. Furthermore, we only contacted website owners when the unauthorized third-party redirect hack was *not* fixed. In turn, we did not interview website owners who fixed the unauthorized third-party redirect hack. While interviewing website owners who had fixed the unauthorized third-party redirect hack could have provided interesting results, it was out of scope for this experiment, and we leave it to future work to address this.

By ensuring that our email servers had valid SPF, DKIM, and DMARC records, we addressed potential issues related to reachability. With our sample being diverse in terms of receiving mail servers we are confident that we did not encounter the problems with one service, like Google, filtering most of our notifications, as described in Stock et al. (2018). However, there is still the general risk that recipients email spam filters rejected our notifications. Reputation monitoring systems or spam blocking lists are helpful to check the reputation of the sender's email address (Maaß et al., 2021b). As we wanted the experiment to be integrated into the regular notification process of the Federal CERT and the hosting providers, we had no control over their mail server configurations and could not subscribe to any spam reporting service. We acknowledge that potential differences could have an impact on deliveries itself as well as delivery rates and time (also affecting the appearance of our notifications in recipients' inboxes), which might have influenced whether the notification was read.

However, we can be sure that these institutions regularly check the reputation of their email addresses. Our notifications were treated the same as any other of their usual communications. We recorded if one of our emails bounced and also asked our partners to report emails that bounced. We are aware of a total of two emails that could not be delivered, indicating a negligible number of emails experiencing technical delivery issues. We only checked whether our university address and the CERT's address are listed on Spamcop.net and Spamhaus.org after completing our experiment, but no problems were found with either address. There is still a high probability that our notifications did not pass the "human" spam filter, as our results in Section 4.2 show. Unfortunately, we cannot quantify how many notifications might have been discarded as spam by the recipients.

External validity. The non-representativeness of our sample limits the external validity of our findings. We could only include websites that were discovered in the web crawling. Furthermore, our notification experiment is exclusively focused on German websites, constraining its external validity. The findings might not be generalizable to other geographical locations, as websites may operate under different legal, cultural, or technological environments. However, we wanted to focus on only one country first to avoid introducing additional variables to our experiment (e.g., legal requirements in different countries, policies of CERTs or hosting providers, or translation of the notification). Furthermore, as Maaß et al. (2021c) noted, restricting notification campaigns to a specific country has the advantage that notifications are better understood and the names of senders are somewhat familiar,

which increases trust in the organization sending the notification. Nevertheless, replicating our study in other countries and for other regulatory and cultural contexts represents an important direction for future work to confirm that our results are valid for a global audience as well.

Lastly, regarding the interaction effect between sender and framing, we cannot be certain that our participants share our view of relating the Federal CERT to the reputational framing, hosting providers to the technical framing, and universities to the neutral framing, which might have affected the effectiveness of our framings.

3.5. Data analysis

For our notification experiment, we measured and controlled two independent variables (framing and sender) with three nominal characteristics each (neutral, technical, reputation, and university, hosting provider, Federal CERT). The existence of third-party redirects was continuously measured by our monitoring system four times a day throughout the observation period, as defined in Section 3.1 (see also Fig. 3). Thus, we measured remediation as our dependent variable as days until remediation (continuous).

We analyzed the differences between the framings (**RQ 1**), and the differences between the senders, and between the senders and the control group (**RQ 2**), using a single-factor ANOVA. To identify a possible correlation between sender and framing, we used a two-factor ANOVA (**RQ 3**). Additionally, we used survival analysis to compare the time until remediation between our treatment groups. For feasibility reasons, we only interpreted the results of the automatically determined status once a day. Therefore, we decided to use right-censoring and Kaplan–Meier estimator, as in previous studies where survival analysis was used (Maaß et al., 2021c; Lone et al., 2022; Çetin et al., 2016, 2018) and Zeng et al. (2019), instead of, e.g., interval censoring.

We used an alpha level of .05 for all statistical tests, and applied post-hoc Holm-Bonferroni correction to counter alpha error cumulation. We also used a priori power analysis to determine the estimated sample size for our statistical analysis. We found that by assuming a medium effect size, we would need 280 websites to run a single-factor ANOVA with four groups (university, hosting provider, Federal CERT, control), and 400 websites to run a two-factor ANOVA with six groups (three senders, three framings). In the absence of similar rules for survival analysis, we aimed for a minimum sample size of 280 websites and attempted to collect up to 400 websites. With a final sample size of $n = 467$ websites, we surpassed that threshold.

4. Results

Between April 2022 and November 2023, we notified 581 website owners about unauthorized third-party redirect hacks. We did not notify 205 website owners who were in our control group. In total, two emails bounced and were not delivered. These were added to the control group. We are not aware of any other delivery failures and, therefore, assume that all our notifications reached their intended destinations. In total, our sample included 786 websites. Table 1 provides the number of notifications sent out in each group. The overall remediation rate for our notification experiment was 42.0%, which means that 58.0% of websites were still compromised eight weeks after two notifications.

Our sample also included websites with an unknown status, for which our monitoring system could not determine a status (see Section 3.1). If the monitoring system returned "unknown", we manually checked the websites periodically, but not daily. Thus, for these websites, we only have a few manually verified data points within the 56-day observation period. To measure the remediation in days as continuous variable we excluded the websites for which we could not clearly define the status or recorded a change in status over the 56-day period, and created a sub-sample for our longitudinal analysis containing $n = 467$ websites.

4.1. RQ 1 –RQ 3: Effect of framing and sender

To determine the effects of **sender** on the time to remediation as a continuous variable, we first used one-way ANOVA with the sub-sample ($n = 467$). University, hosting provider, and Federal CERT did not show any outliers in the box-plot-diagram, but the control group had 16 extreme outliers, which seemed reasonable to us: Contrary to the otherwise long time to remediation – or non-remediation until the end of the observation period – these cases were indeed remediated and sometimes even relatively quickly. Data were not normally distributed (Shapiro–Wilk test, $p < .001$), but ANOVA has proven robust against violations of the assumption of normality (Blanca et al., 2017; Lix et al., 1996; Schmider et al., 2010; Harwell et al., 1992; Glass et al., 1972). There was no homogeneity of variance (Levene's test, $p < .001$), and we, thus, interpreted the results of Welch's test. For **sender** we had defined four categories: control group, university, hosting provider, and Federal CERT. The mean time to remediation was 51.61 days (95%-CI[49.30, 53.91]) for the control group, 39.83 days (95%-CI[35.41, 44.26]) for university, 39.61 days (95%-CI[34.77, 44.45]) for the hosting provider, and 41.93 days (95%-CI[37.75, 46.11]) for the Federal CERT. We observed that the time to remediation was significantly different for at least one group, but only with a small effect, Welch's $F(3, 230.87) = 16.68$, $p < .001$, $\eta^2 = .054$. There were no statistically significant differences between the three treatment groups ($p > .05$). Still, each sender was significantly different from the control group (Games-Howell post-hoc analysis, $p < .001$: university - Control: −11.77, 95%-CI[−18.31, −5.24], Federal CERT - Control: −9.68, 95%-CI[−15.93, −3.42], provider - Control: −12.0, 95%-CI[−19.02, −4.98]). Applying the Holm-Bonferroni correction confirmed these findings (see Appendix D, Table 4 for all results).

Second, we performed a two-way ANOVA to assess the effects of **framing**, **sender**, and the **interaction effect** between framing and sender on the time to remediation as continuous variable. For framing, there were no outliers, as indicated by the box-plot diagram. In none of the categories, the data were normally distributed (Shapiro–Wilk test, $p < .001$), but again, we can assume robustness against violations of the assumption of normality, since our sample size was larger than 15 domains for each group. We also determined homogeneity of variances using Levene's test, which showed that equal variance could be assumed ($p = .054$). We excluded the control group as sender from our data, as we only wanted to determine effects if a notification was sent ($n = 334$). For **framing** we had defined three categories: neutral, technical, and reputation. The mean time to remediation was 42.59 days (95%-CI[38.24, 46.95]) for the neutral framing, 40.05 days (95%-CI[35.74, 44.36]) for the technical framing, and 39.02 days (95%-CI[34.32, 43.72]) for the reputational framing. Interestingly, as shown in Fig. 4(a), the reputational framing resulted in the highest remediation rates for Federal CERT as sender, and the neutral framing resulted in the highest remediation rates for university as sender. However, our assumption that there is a general interaction effect between the framing and sender was not valid for the hosting provider as sender, where both the reputational and the neutral framing resulted in higher remediation rates than the technical framing. Furthermore, our main model was not significant, $F(8, 325) = 0.53$, $p = .834$, and neither of our framings (neutral, technical, reputation), $F(2, 325) = 0.09$, $p = .916$, nor any of our senders (university, hosting provider, Federal CERT), $F(2, 325) = 0.41$, $p = .661$, or the interaction between framing and sender, $F(4, 325) = 0.75$, $p = .559$, were significant.

Third, we conducted survival analysis using the sub-sample ($n = 467$). The websites that were notified with a reputational **framing** were estimated to remediate with a mean of 39.02 days (95%-CI[34.39, 43.65]), requiring the least time until remediation compared to a technical framing (M = 40.05, 95%-CI[35.81, 44.3]), or a neutral framing (M = 42.59, 95%-CI[38.25, 46.93]). However, the results of the log-rank test showed that survival distributions do not differ significantly, $\chi^2(2) = 0.299$, $p = .861$ (see Fig. 4(b)).

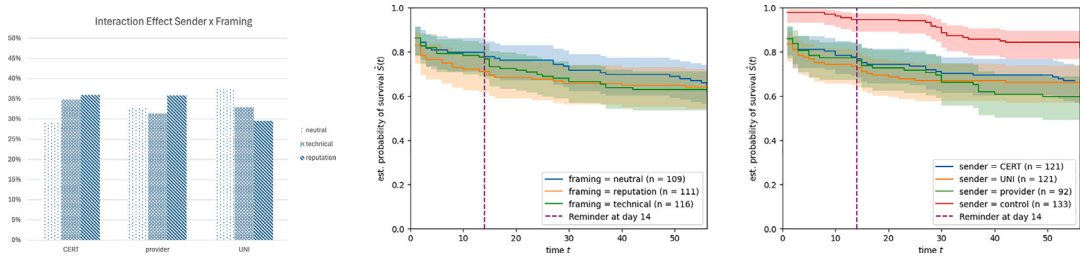

(a) Percentage of remediated websites. (b) Kaplan Meier curve for framing. (c) Kaplan Meier curve for sender.

Fig. 4. RQ1–RQ3: Effect of framing and sender.

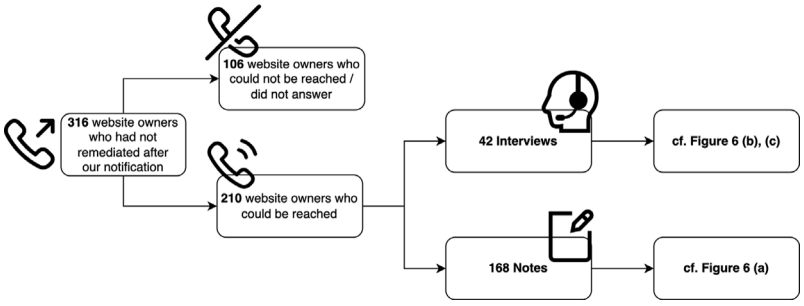

Fig. 5. Illustration of the process of the follow-up interviews.

We also analyzed the effect of **sender** within our survival analysis. The websites that the hosting providers notified were estimated to remediate with a mean of 39.61 days (95%-CI[34.86, 44.36]), requiring the least time until remediation, whereas websites that were notified by the university (M = 39.84 days, 95%-CI[35.48, 44.2]), or the Federal CERT (M = 41.93 days, 95%-CI[37.76, 46.12]) took longer. All treatment groups required significantly fewer time until remediation than the control group (M = 51.61 days, 95%-CI[49.33, 53.88]). The log-rank test confirms that significant differences exist between at least two of the four groups, $\chi^2(3) = 17.49$, p <.001. Pairwise post-hoc log-rank tests showed statistically significant differences in the survival distributions of the control group and the Federal CERT, $\chi^2(1) = 10.63$, p = .001, the control group and the provider, $\chi^2(1) = 15.79$, p <.001, as well as the control group and the university, $\chi^2(1) = 13.05$, p <.001. However, there were no significant differences in the survival distribution of the treatment groups (see Fig. 4(c)).

4.2. *RQ 4: Reasons for non-remediation*

Approximately eight weeks after our initial notifications, we contacted the website owners who had not remediated their websites. Between October 2022 and December 2023, we called 316 website owners via phone. We managed to reach 210, of whom 42 agreed to an interview. While not all agreed to an interview, most website owners provided some information in the phone calls, e.g., whether they received our notifications or already knew about the attack, which we also noted down as codes (see Appendix C.3 for an overview of the codes we used). Thus, we were also able to analyze the responses we received during the phone calls. See Fig. 5 for an illustration of the process and Figure 6 for the results.

Of those who had not agreed to an interview ($n = 168$), 75 website owners said they were not aware of our email notification. In contrast, 26 website owners said they got some information, and eight already knew about the attack before our notification. We were unable to collect further information from 51 website owners, who either would not speak with us (14 website owners) or did not respond as promised during the initial call (37 website owners). Instead, in our interviews ($n = 42$), we mainly talked to website owners who had received our notification: 28 interviewees had received and read our notification, two website owners said they had not received our notification, and 12 did not remember if they received it.

All 14 interviewees who did not – knowingly – receive and, thus, read our notification stated that they overlooked our email, probably because it got buried among the daily spam. Or, as [P29] summarizes: *"Well, I have to be honest and say that I think a message like this [...] comes so unprepared and you receive such an abundance of emails every day suggesting that something needs to be done in some way, um, that I think it's incredibly difficult to fight against [...] this flood of spam and it's, therefore, incredibly difficult to actually convey the seriousness that it needs to be received and read [...] at all"*. P[36] also explained the email address we wrote to was too general, and it probably got lost due to vacation: *"You sent it to a rather general address. [...] Maybe someone was on vacation, I don't know, but in any case it went through"*. P[29] further added that the email address we used belongs to an old website and is only infrequently checked: *"[I]t is our former company account, [which] no longer exists, and [...] is only checked very rarely"*. Some participants also mentioned that the subject line we used did not help the email to stand out. P[7] explained: *"'Important information' sounds a bit more like... I get a lot of e-mails like that, where it says something like 'very important, very important'. That actually tends to get thrown out, because most of the time [...] it's just the opposite, when it somehow says 'super important' or something like that and then I click on it and then I see, ok, that's just something that's not important at all"*.

Even if they had knowingly received and read the notification, 11 interviewees said they regarded the notification as spam and, thus, did not immediately react to it. P[4] said the sender was not familiar (*"I'd say that the average user, who sells his cake on the internet, won't know the CERT anyway"*). P[37] mistrusted Federal CERT as sender, questioning their motivation, and found this kind of notification unfamiliar: *"So, [...] cert@BSI didn't mean anything to me"*. P[22] also said that a federal office scanning websites at no cost and informing users about vulnerabilities seemed unrealistic to them: *"Or CERT [...] this seems implausible itself. Because, I think, why should some irrelevant website be of interest [...]? You know[...] in those offices, where they usually have other*

things to do, [...] and then they additionally write emails to people [...] So why should they do that? [...] [A]nd then I thought, no, not really, that can only be spam".

In addition, nine interviewees did not consider the unauthorized third-party redirect hack relevant enough to take immediate action. The main reason, given by three interviewees, was that the website has no priority for them. P[4] said they did not have time to fix the unauthorized third-party redirect hack immediately and then forgot about it: *"To be honest, I haven't had time to take care of it. It sounds a bit weird that we as a web-agency don't take care of our website, but we're so busy with work that we never really actively advertise on the website. So it wasn't a big priority for me at first".* P[7] felt that troubleshooting is more expensive than just launching a new website, so they did not react: *"[...] if I hire a web designer or a programmer who spends, I don't know, 3 days searching for something, it will cost me more than building a new website".* Two interviewees said they had not understood the problem, and four could not reproduce it. P[20] could not reproduce the problem because they did not read the advice on how to verify the unauthorized third-party redirect hack carefully enough. Two interviewees stated that they had checked their files but could not identify anything suspicious. Two interviewees admitted that they had understood the problem and found it relevant enough to act on, but did not know how to remediate, and 11 interviewees thought the attack was already remediated. We provide relevant quotes from the interviews in Appendix E, Table 5.

5. Discussion and future work

As discussed in Stock et al. (2018) three key parameters contribute to the success of notification campaigns: (1) successful delivery of the notifications; (2) trust in the notification process; and (3) enhancing recipients' ability to remediate.

We must assume that our manual efforts in retrieving email addresses from the imprints of the websites were worthwhile, as the majority of notifications seemingly reached recipients' inboxes. This is supported by the fact that only two notifications bounced.[5] However, we observed that the second and third parameter, "reaching out" and "breaking through", were significant challenges in our notification experiment.

5.1. Reaching out: Notification channel and sender

We learned that a significant number of website owners, who did not remediate the unauthorized third-party redirect hack within 56 days, could not recall receiving our notification. Some interviewees mentioned that the contact information we used was too unspecific or outdated, which diminished the reliability of our email notifications. One alternative could be to reach out to those affected via alternative notification channels, such as phone calls — the second most preferred notification channel identified in our pre-study (Hennig et al., 2022b). However, as highlighted in previous research (Maaß et al., 2021c; Stock et al., 2018), and confirmed during our follow-up interviews, calling website owners is time-consuming and costly, making it unsuitable for large-scale notification campaigns.

Another reason why some email notifications went unnoticed or were regarded as spam was the lack of familiarity with the sender, distrust regarding the sender's intention, or simply the fact that they were notified at all. When we asked for recommendations for suitable senders, we found that opinions varied greatly among the interviewees. Website owners indicated a preference for notifications from a trustworthy and reputable sender but could not provide specific examples. We assumed that the Federal CERT, as also investigated in Lone et al. (2022) and Nosyk et al. (2023), or the respective hosting providers are perceived as reputable senders. However, even when websites were notified of these trusted entities, our experiment demonstrated that remediation is not significantly higher than if the notifications are sent by the university. This, thus, confirms that the perception of a sender's trustworthiness depends highly on the individual recipient, as already stated in Hennig et al. (2023). However, since all our treatment groups performed significantly better than the control group, we can support previous research (Durumeric et al., 2014; Kührer et al., 2014; Li et al., 2016a,b, 2019; Maaß et al., 2021a,c; Stock et al., 2018, 2016; Vasek and Moore, 2012; Zeng et al., 2019; Çetin et al., 2017, 2019b, 2016) in that sending out notifications does indeed encourage remediation.

We encourage future research to investigate these findings in the light of the Elaboration Likelihood Model (ELM) of persuasion, a theoretical model from communication science in the context of media effects research. We can consider our notification as a stimulus in the form of a piece of persuasive information that should motivate the recipients to change their attitude or behavior, i.e., remediate their vulnerable websites. The ELM proposes two "routes" by which recipients interact with the information (e.g., in a notification) based on the level of elaboration that is stimulated. If deemed convincing, the notification leads to a change in attitude or behavior, i.e., results in the remediation of the described issue. First, when elaboration is high, e.g., if the recipient has a high need for cognition and is motivated to interact with the notification, the information is processed on the *central route*. Here, the content and its arguments are closely scrutinized, and the information is carefully examined. However, based on what we learned in our follow-up communication with the recipients, we have to assume that vulnerability notifications via email are processed on the *peripheral route*, where elaboration is low and information is mainly processed based on heuristic principles, e.g., credibility heuristics where the sender's perceived expertise is assessed. Peripheral cues, i.e., external factors such as the perceived credibility of the sender, have a particularly strong effect on the information processing via the *peripheral route*, as they serve to assess whether a change in attitude or behavior is necessary.

In the context of vulnerability notifications, this leaves researchers and practitioners with two challenges. First, external circumstances should be improved so that the recipient is empowered and motivated to process a notification directly on the *central route*. Taking sender and notification channel as an example, elaboration motivation could be stimulated by personal relevance (e.g., choosing a sender that has a personal relationship to the participant), and elaboration ability could be stimulated by sufficient prior knowledge about the topic or a non-distracting environment (e.g., choosing a notification channel that is exclusively used to address tech-savvy contacts directly, as opposed to a universal notification channel that receives many irrelevant messages and is mainly managed by non-tech-savvy contacts who do not expect to handle critical vulnerability notifications). Future work should use the ELM to systematically identify external factors in the context of vulnerability notifications that increase elaboration motivation and elaboration ability.

RFC 9116 (Shafranovich and Foudil, 2022) proposes a file format (security.txt) in which dedicated contact information for technical contacts are provided so that vulnerability notifications can be sent in a "non distracting environment". Unfortunately, related work has shown that the adoption of security.txt is alarmingly low (Poteat and Li, 2021; Findlay and Abdou, 2022). Yet, we still lack a fundamental understanding of the reasons behind this hesitance among website owners to provide such contact information. To address this gap, we recommend future work to directly engage with website owners, explore potential misconceptions or obstacles they face, and identify resources website owners need. Furthermore, it might be beneficial to raise awareness for the importance of providing such contact information while fostering a broader understanding of website security within a comprehensive awareness campaign. Future work might evaluate the impact of concepts like a "Web Security Awareness Month", which could serve as a rallying call for better practices in the community. Organizations such

[5] As described in Section 3.2, we purposively did not include further tracking mechanisms, so we can only assess delivery success based on the number of bounces.

as chambers of commerce, industry associations, national CERTs, and external service providers – including hosting and content management platforms like WordPress – are in a unique position to lead these efforts. By coming together to offer comprehensive information, workshops, or webinars focused on website security, they can create a powerful network of support and knowledge-sharing, inspiring website owners to embrace best practices and enhance their security posture. Based on the recommendations of Gerber et al. (2025), website owners may find the problem and its remediation less daunting if they feel supported and have the opportunity to connect with others, which leads to an increase in website owners' self-efficacy and consequently, increases their ability to remediate.

Future work should also investigate how theoretical foundations from the context of organizational information security awareness campaigns (see, e.g., Bada et al. (2015) for an overview), could shape the organization of and the communication around such events, as has been explored in related but distinct contexts, such as Social Network Analysis to select security champions in companies (Dang-Pham et al., 2017), or Transactive Memory System Theory (TMS) to facilitate sharing IT security knowledge within a group (Alahmari et al., 2023). Further factors that influence information security behavior of a website owner, such as job characteristics, personality traits, or real-life exposure (Furnell and Rajendran, 2012), should also be investigated to identify different types – or personas – of website owners. This is helpful information for ultimately tailoring such events to different target groups and making them as effective as possible for participants.

Second, peripheral cues have to be researched in more detail. This is especially important as it has to be considered that, depending on the context, every cue – or factor – has different effects. A sender's reputation, for example, might affect the degree of elaboration, i.e., whether the information is processed on the central or the peripheral route. Or it may affect decision-making within each of the routes, i.e., serves as a peripheral cue within the peripheral route, to motivate the recipient to take the notification as credible based on their credibility heuristic, or influences the valence of elaboration within the central route by increasing personal relevance. In general, every factor that accompanies the communication can be a peripheral cue. For communication in (mass) media, factors such as characteristics of the speaker (e.g., voice, presentation style, appearance, prominence, etc.), characteristics of the situation (e.g., environment in which a message is presented), characteristics of the message itself (e.g., length, repetition), or personal characteristics of the recipient (e.g., mood, distraction, attitude towards a message or the sender, attitude of other persons towards the message, etc.) are described as influencing factors (Burkart, 2021; O'Keefe, 2015; Wirth and Kühne, 2013). It would be interesting to know which of these factors are especially effective within risk communication, i.e., vulnerability notifications, in contrast to other persuasive communication, such as advertisements.

Our results regarding suitable channels and senders are twofold. Firstly, since none of the senders in our experiment proved to be most effective, we recommend that practitioners designate a single authority, like the Federal CERT, to handle security or privacy-related vulnerability notifications. This entity should be promoted nationally, possibly through partnerships with business associations to amplify its reach. Secondly, for future research on notification campaigns, we recommend exploring additional strategies to raise a general awareness among website owners for cybersecurity, and encourage them to provide accurate security contacts (e.g., security.txt). Although email scales well for notification campaigns, it should be accompanied by efforts to increase general awareness of security notifications to ensure they stand out amidst other emails. Within the context of the ELM, enhancing background knowledge and the relevance of a topic positively impacts elaboration ability and motivation, thereby increasing the likelihood that a notification is processed via the *central route*.

5.2. Breaking through: Framing and notification content

We discovered that website owners who did not remediate often did not consider the notification relevant enough to take (immediate) action against the unauthorized third-party redirect hacks. Interestingly, in contrast to previous research (e.g., Maaß et al., 2021c), we see that trust in the notification *content* was less of an issue, a finding supported by our remediation rate. Once the notification was considered relevant, a significant percentage (42%) of website owners remediated the unauthorized third-party redirect hacks. However, as described above, it appears the notifications did not convey sufficient relevance to motivate elaboration and, thus, to surpass the daily "flood" of unsolicited messages. This was also echoed in our interviews: Website security was not a main priority for the website owners we interviewed. Some were unable to gauge the extent of the attacks and were often reluctant to invest more resources than absolutely necessary in website maintenance, especially considering that it is typically not their bread and butter.

Providing incentives to raise awareness for the problem's severity did not significantly enhance the effectiveness of our notifications. Neither of the two framings with incentives (reputation, technical) was significantly more effective. We observed a slight advantage of the reputational framing in terms of time to remediation compared to the technical or the neutral framing, but the differences were not significant. We also examined whether a framing sent from an authoritative source – one that could impose fines for non-compliance, as suggested by Maaß et al. (2021c) – might enhance the effectiveness of our vulnerability notifications. In the absence of a legal framing in our experiment, we investigated the effect of other framing–sender combinations (i.e., technical framing from a hosting provider, and reputational framing from the Federal CERT). Our study did not uncover any correlation between sender and framing, making it impossible to generalize the results from Maaß et al. (2021c) and Utz et al. (2023) to other framing–sender combinations.

While we could not generalize results from Çetin et al. (2019a,b), who found that technical incentives in the form of quarantining compromised domain owners is effective, or Maaß et al. (2021c) and Utz et al. (2023) who proposed that certain framings are only effective in combination with a corresponding sender, we were able to confirm results from Zeng et al. (2019) in that framings have no (additional) effect on remediation. As both of our framings with incentives were only slightly more effective than the framing without any incentive, we do not observe a significant effect of a certain type of notification being discarded as spam more often (e.g., notifications with a neutral framing being discarded more frequently than those with a reputational framing). Perhaps the wording of the technical and reputational framings were not explicit enough to motivate the recipients to remediate and protect themselves more than was the case with the neutral framing.

Possible explanations can be derived from the Protection Motivation Theory (PMT) (Rogers, 1975). According to PMT, reactions towards warnings are based on two appraisal processes: the *coping appraisal* and the *threat appraisal*. In the process of *threat appraisal* individuals assess their perception of a threat, based on three factors: perceived severity, which refers to the expected degree of harm caused by the threat, perceived vulnerability, which indicates the likelihood of experiencing that harm, and maladaptive rewards, meaning the aspects that reinforce insecure behaviors. For individuals to feel motivated to react to a warning, the perceived severity and perceived vulnerability must outweigh any maladaptive rewards. In our notification experiment, we aimed to increase the expected degree of harm and enhance participants' perceived severity by using technical and reputational incentives. However, it is possible that the incentives we provided were not strong enough to effectively influence participants' perceptions.

Some researchers propose that increasing the perceived risk associated with vulnerabilities, or threatening website owners with public disclosure if they do not remediate, could provide stronger stimuli to

evoke action and improve remediation rates significantly (Li et al., 2016a; Stock et al., 2018). Based on our results, we would argue against employing threatening incentives. One interviewee proposed that notifications should be addressed to browser vendors, allowing them to directly block malicious websites in the respective browsers and, thereby, alert website owners about the severity of the unauthorized third-party redirect hack. While such warnings can effectively attract the website owner's attention and significantly increase remediation, they must be accompanied by additional information about the compromise and clear remediation instructions, as noted by Li et al. (2016b). Furthermore, stimulating threat perception must not lead to fear arousal (Mayer et al., 2017; Zou et al., 2024). Research has shown that heightened pressure and concern may lead to refusal and mental overload, as observed in contexts such as passwords (Alkaldi and Renaud, 2018; Dupuis et al., 2021) or cybersecurity incidents (Dupuis and Renaud, 2021; Renaud et al., 2021a). Related work has also shown that security and privacy topics in general are perceived as complex and frightening by users, leading to people feeling overwhelmed and frustrated (Renaud et al., 2021b; Gerber et al., 2025; Da Silva and Jensen, 2022). Instead of stimulating negative emotions that might inhibit recipients to take action, recent work recommends to frame security and privacy topics as "more engaging and enjoyable" (Gerber et al., 2025). In the context of vulnerability notifications, prior results (Lone et al., 2022; Hennig et al., 2023) also propose that positive reinforcement – such as emphasizing that remediation reduces the threat to others – may be more effective in motivating website owners to remediate compromises.

In the context of the PMT, positive reinforcement can also take place in the process of *coping appraisal*, where individuals assess their ability to cope with a threat based on three factors: self-efficacy, response efficacy, and response costs. To motivate protective actions, e.g., remediation in our context, Mayer et al. (2017) describe that self-efficacy and response efficacy must outweigh response costs. In the context of vulnerability notifications this means that the resources the recipients need to spent for remediation (e.g., time, knowledge, or costs for staff or external support) must be outweighed by the possibilities they have to remediate (e.g., access to the vulnerable system, technical knowledge), and a perceived high probability that the vulnerability will be removed after remediation. Thus, another reason our framings were not effective is that our notification did not stimulate the recipients' ability to cope with the threat sufficiently to offset the expected costs.

Reasons mentioned by website owners for not remediating included the belief that they had already remediated, or that they were unsure about the appropriate solution. This points to transparency issues, where individuals are not able to perceive the status of their system. Some interviewees expressed a desire for more information such as a PDF attachment or links to further resources. Previous studies also indicated that providing a self-service tool or detailed reports alongside a notification can be valuable (Stock et al., 2018, 2016; Çetin et al., 2016, 2017; Maaß et al., 2021c). Furthermore, some of our interviewees suggested that notifications should include offers of support, such as specific assistance for a fee or help from the Federal CERT. This indicates a distinct need to understand the systems individuals use and regain control over it. Future work in this regard should investigate creating additional materials based on cognitive and behavioral theories, such as Protection Motivation Theory (PMT), which was, for example, used to design interventions that should encourage users to change their passwords after a data breach (Zou et al., 2024), or Theory of Planned Behavior (TBP) (Ajzen, 2002, 1991), as was, for example, used by Bulgurcu et al. (2010) who investigated the intention of employees to comply with security policies in their company. Specifically, factors like "attitude" (TBP), which was found to have a reliable medium effect on increasing secure information security behavior (Mayer et al., 2017) and a large effect on security policy compliance (Cram et al., 2019), "self-efficacy" (PMT, TBP), which is a factor in both theories and has proven to have a reliable weak (Mayer et al., 2017) to medium positive effect (Cram et al., 2019) especially in combination with "controllability" (Mayer et al., 2017), as well as "subjective norms" (TBP) and "response efficacy" (PMT) (reliable weak positive effect (Mayer et al., 2017), medium to large effect (Cram et al., 2019)), should be investigated to address the issues and needs identified in our interviews.

Another possible reason why our incentives had no (additional) effect might be habituation effects. As several interviewees mentioned, they often receive emails that urge them to take action. Therefore, any incentive that demands action might be dismissed as yet another spam message. However, this underlines the importance of sending notifications without any additional call-to-action, and helping recipients to distinguish legitimate vulnerability notifications from those with marketing interests.

It might be worth discussing whether a larger sample size could reduce the likelihood of underestimating potential differences between our treatment groups (type II error) in our statistical analysis. However, our findings revealed statistically significant differences between the control group and the treatment groups, and we met the required sample size determined by a priori power analysis. Therefore, simply increasing our sample size would not alter our findings in meaningful ways. Nonetheless, as discussed in Section 3.4, replicating our study with a more diverse sample that includes websites from different geographical regions and legal contexts would enhance the robustness of our results.

Rather than using threats to compel website owners to remediate compromises, we recommend supporting website owners in their coping appraisal by encouraging them, offering support, and including references to further information in a broader awareness campaign. Additionally, we suggest sharing information materials with third parties, such as CERTs, hosting providers, or browsers (e.g., Google Safe Browsing), which can provide further support or distribute warnings. Future research should develop materials based on theoretical foundations such as Protection Motivation Theory or Theory of Planned Behavior and explore whether supplementary materials enhance remediation or foster mistrust. We also recommend investigating encouraging framings for future work. Our notification will be freely available to facilitate its use in practice,[6] but we advice practitioners to use it without any specific framing.

6. Conclusion

In our study, we focused on unauthorized third-party redirect hacks, which are not easily recognizable by non-experts. While the websites display normal content when accessed directly, infected websites will list unusual and eventually malicious URLs in search engine results. This is not only harmful to a website's reputation, but it also indicates that an attacker gained write access to the website. The goal of our research was to combine results of previous work and, based on qualitative interviews with 25 website owners conducted in a pre-study (Hennig et al., 2022b), develop an effective notification process. We then conducted a quantitative 3 × 3 randomized controlled notification experiment, measuring differences between senders that were deemed suitable by the interviewees (i.e., university, hosting provider, and Federal CERT), and different framings that should incentivize website owners to remediate. Between April 2022 and November 2023 we notified 581 website owners via email, observing an additional 205 that were in our control group. We found that only 42.0% of websites remediated within 56 days. We can confirm previous research,

[6] Before publication, we will revise our notification template based on the feedback from the interviews.

which also found that we could not identify a sender or framing that was more effective than others (Stock et al., 2018; Zeng et al., 2019; Çetin et al., 2016). It may be that, for example, a legal framing or a sender with legal authority might be more effective; however, the use of these conditions should be carefully considered in light of potential ethical issues. Nevertheless, remediation rate was significantly better in the treatment groups compared to the control group, indicating that notification campaigns are effective in increasing remediation.

In addition to notifying website owners via email, we called them when the unauthorized third-party redirect hack was not remediated within 56 days. We conducted 42 qualitative follow-up interviews to find reasons for non-remediation. The main reason for non-remediation was that website owners could not recall receiving our notification or regarded it as spam. Thus, we found increasing the perceptibility of notifications to be a major issue. We recognize a general need to raise awareness about security notifications and the provision of proper security contacts by website owners, especially for those whose websites are not their primary business. We also recommend that future work investigates how email notifications can stand out from the mass of daily emails in recipients' inboxes, e.g., by embedding notifications in a broader awareness campaign, by analyzing framings that encourage recipients to open an email, or by following up with website owners who have remediated the problem to investigate reasons for remediation. We also found that providing additional support, e.g., in the form of a self-service tool or a PDF attachment, might increase recipients self-efficacy and encourage website owners to take action.

We hope that our results and the recommendations derived from them can guide researchers in finding more effective ways to notify website owners of compromised websites. We make our revised notification texts freely available online to help practitioners when notifying the victims of compromises.

CRediT authorship contribution statement

Anne Hennig: Writing – review & editing, Writing – original draft, Visualization, Validation, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Maxime Veit:** Writing – original draft, Software. **Leoni Schmidt-Enke:** Writing – original draft, Data curation. **Fabian Neusser:** Writing – original draft, Data curation. **Dominik Herrmann:** Writing – original draft. **Peter Mayer:** Writing – review & editing, Supervision, Conceptualization.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Anne Hennig reports financial support was provided by Federal Ministry of Education and Research Berlin Office. Leoni Schmidt-Enke reports financial support was provided by Federal Ministry of Education and Research Berlin Office. If there are other authors, they declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This research is supported by the German Federal Ministry of Education and Research as part of the INSPECTION project (Zuwendungsnummer 16KIS1113), and by funding from the topic Engineering Secure Systems, topic 46.23.01 Methods for Engineering Secure Systems, of the Helmholtz Association (HGF), Germany and by KASTEL Security Research Labs. Special thanks to Alexandra Pawelek, Elly Reich, Lauritz Kanyi, and Miriam Mutter who helped at different stages of this research as part of their jobs as student assistants.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.cose.2025.104682.

Data availability

Data will be made available on request.

References

Ajzen, Icek, 1991. The theory of planned behavior. Organ. Behav. Hum. Decis. Process. 2, 179–211. http://dx.doi.org/10.1016/0749-5978(91)90020-T, https://www.sciencedirect.com/science/article/pii/074959789190020T.

Ajzen, Icek, 2002. Perceived behavioral control, self-efficacy, locus of control, and the theory of planned behavior. J. Appl. Soc. Psychol. 4, 665–683. http://dx.doi.org/10.1111/j.1559-1816.2002.tb00236.x.

Alahmari, Saad, Renaud, Karen, Omoronyia, Inah, 2023. Moving beyond cyber security awareness and training to engendering security knowledge sharing. Inf. Syst. E Bus. Manag. 21, 123–158. http://dx.doi.org/10.1007/s10257-022-00575-2.

Albayram, Yusuf, Walker, Jaden, 2024. Investigating effectiveness of informing users about breach status of their email addresses during website registration. Int. J. Hum. Comput. Interact. 1–20. http://dx.doi.org/10.1080/10447318.2024.2404721.

Alkaldi, Nora, Renaud, Karen, 2018. Encouraging password manager adoption by meeting adopter self-determination needs (Extended version). SSRN Electron. J. http://dx.doi.org/10.2139/ssrn.3259563.

Bada, Maria, Sasse, Angela, Nurse, Jason, 2015. Cyber security awareness campaigns: Why do they fail to change behaviour? pp. 118–131.

BitofWP, 2019. WordPress infected with the pharma hack? How to detect, clean and secure your site from it - DEV community. (last Accessed 19 April 2024). https://dev.to/bitofwp/wordpress-infected-with-the-pharma-hack-how-to-detect-clean-and-secure-your-site-from-it-4fja.

Blanca, M.J., Alarcón, R., Arnau, R., Bendayan, R., 2017. Non-normal data: Is ANOVA still a valid option? Psicothema 29 (4), 552–557. http://dx.doi.org/10.7334/psicothema2016.383.

Bouwmeester, Brennen, Turcios Rodriguez, E.R., Gañán, Carlos, van Eeten, Michel, Parkin, Simon, 2021. "The thing doesn't have a name": Learning from emergent real-world interventions in smart home security. In: Proceedings of the 17th Symposium on Usable Privacy and Security. SOUPS 2021, In: Proceedings of the 17th Symposium on Usable Privacy and Security, SOUPS 2021, USENIX Association, pp. 493–512.

Bulgurcu, Burcu, Cavusoglu, Hasan, Benbasat, Izak, 2010. Information security policy compliance: An empirical study of rationality-based beliefs and information security awareness. MIS Q. 34 (3), 523–548, http://www.jstor.org/stable/25750690.

Burkart, Roland, 2021. Kommunikationswissenschaft. Boehlau Verlag.

Cazorla, Lorena, Alcaraz, Cristina, Lopez, Javier, 2018. Cyber stealth attacks in critical information infrastructures. IEEE Syst. J. 12 (2), 1778–1792. http://dx.doi.org/10.1109/JSYST.2015.2487684.

Çetin, Orçun, Altena, Lisette, Gañán, Carlos, Eeten, Michel van, 2018. Let me out! evaluating the effectiveness of quarantining compromised users in walled gardens. In: Fourteenth Symposium on Usable Privacy and Security. SOUPS 2018, USENIX Association, Baltimore, MD, pp. 251–263, https://www.usenix.org/conference/soups2018/presentation/cetin.

Çetin, Orçun, Gañán, Carlos Hernandez, Altena, Lisette, Kasama, Takahiro, Inoue, Daisuke, Tamiya, Kazuki, Tie, Ying, Yoshioka, Katsunari, van Eeten, Michel, 2019a. Cleaning up the internet of evil things: Real-world evidence on ISP and consumer efforts to remove mirai. In: Proceedings 2019 Network and Distributed System Security Symposium. https://www.ndss-symposium.org/ndss-paper/cleaning-up-the-internet-of-evil-things-real-world-evidence-on-isp-and-consumer-efforts-to-remove-mirai/.

Çetin, Orçun, Gañán, Carlos, Altena, Lisette, Tajalizadehkhoob, Samaneh, Eeten, Michel van, 2019b. Tell me you fixed it: Evaluating vulnerability notifications via quarantine network. In: 2019 IEEE European Symposium on Security and Privacy. EuroS&P, EuroS&P, Vol. 00, 326–339. http://dx.doi.org/10.1109/eurosp.2019.00032.https://ieeexplore.ieee.org/document/8806733,

Çetin, F.O., Ganan, C. Hernandez, Korczynski, M.T., van Eeten, M.J.G., 2017. Make notifications great again: learning how to notify in the age of large-scale vulnerability scanning. WEIS 2017, In: 16th Workshop on the Economics of Information Security, San Diego, pp. 1–23, http://resolver.tudelft.nl/uuid:621f4a4f-e5d9-4f04-abc4-46252f9db3db.

Çetin, Orçun, Jhaveri, Mohammad Hanif, Gañán, Carlos, Eeten, Michel van, Moore, Tyler, 2016. Understanding the role of sender reputation in abuse reporting and cleanup. J. Cybersecur. 2 (1), 83–98. http://dx.doi.org/10.1093/cybsec/tyw005, https://academic.oup.com/cybersecurity/article/2/1/83/2629556.

Cram, W. Alec, D'Arcy, John, Proudfoot, Jeffrey G., 2019. Seeing the forest and the trees: a meta-analysis of the antecedents to information security policy compliance. MIS Q. 43 (2), 525–554. http://dx.doi.org/10.25300/MISQ/2019/15117.

Da Silva, Joseph, Jensen, Rikke Bjerg, 2022. "Cyber security is a dark art"': The CISO as soothsayer. Proc. ACM Hum. Comput. Interact. 6 (CSCW2), 31. http://dx.doi.org/10.1145/3555090, Article 365.

Dang-Pham, Duy, Pittayachawan, Siddhi, Bruno, Vince, 2017. Applications of social network analysis in behavioural information security research: Concepts and empirical analysis. Comput. Secur. 68, 1–15. http://dx.doi.org/10.1016/j.cose.2017.03.010, https://www.sciencedirect.com/science/article/pii/S0167404817300639.

Dupuis, Marc, Jennings, Anna, Renaud, Karen, 2021. Scaring people is not enough: An examination of fear appeals within the context of promoting good password hygiene. In: Proceedings of the 22nd Annual Conference on Information Technology Education. SIGITE '21, Association for Computing Machinery, New York, NY, USA, pp. 35–40. http://dx.doi.org/10.1145/3450329.3476862.

Dupuis, Marc, Renaud, Karen, 2021. Scoping the ethical principles of cybersecurity fear appeals. Ethics Inf. Technol. 23 (3), 265–284. http://dx.doi.org/10.1007/s10676-020-09560-0.

Durumeric, Zakir, Li, Frank, Kasten, James, Amann, Johanna, Beekman, Jethro, Payer, Mathias, Weaver, Nicolas, Adrian, David, Paxson, Vern, Bailey, Michael, Halderman, J Alex, 2014. The matter of heartbleed. In: Proceedings of the 2014 Conference on Internet Measurement Conference, IMC '14. IMC '14, Association for Computing Machinery, Vancouver, BC, Canada, pp. 475–488. http://dx.doi.org/10.1145/2663716.2663755.

Findlay, William, Abdou, AbdelRhaman, 2022. Characterizing the adoption of security.txt files and their applications to vulnerability notification. In: Proceedings of the 2022 Workshop on Measurements, Attacks, and Defenses for the Web. http://dx.doi.org/10.14722/madweb.2022.23014, https://people.scs.carleton.ca/abdou/findlay2022_madweb_authors_copy.pdf.

Furnell, Steven, Rajendran, Anish, 2012. Understanding the influences on information security behaviour. Comput. Fraud Secur. 2012, 12–15. http://dx.doi.org/10.1016/S1361-3723(12)70053-2, https://www.sciencedirect.com/science/article/pii/S1361372312700532.

Gasser, Oliver, Scheitle, Quirin, Denis, Carl, Schricker, Nadja, Carle, Georg, 2017. Security implications of publicly reachable building automation systems. In: 2017 IEEE Security and Privacy Workshops. SPW, pp. 199–204. http://dx.doi.org/10.1109/SPW.2017.13.

Gerber, Nina, Zimmermann, Verena, von Preuschen, Alexandra, Renaud, Karen, 2025. Unpacking the social and emotional dimensions of security and privacy user engagement. In: Proceedings of the 21th Symposium on Usable Privacy and Security. SOUPS 2025, In: Proceedings of the 21th Symposium on Usable Privacy and Security, SOUPS 2025, USENIX Association.

Glass, Gene V., Peckham, Percy D., Sanders, James R., 1972. Consequences of failure to meet assumptions underlying the fixed effects analyses of variance and covariance. Rev. Educ. Res. 42 (3), 237–288. http://dx.doi.org/10.3102/00346543042003237.

Golla, Maximilian, Wei, Miranda, Hainline, Juliette, Filipe, Lydia, Dürmuth, Markus, Redmiles, Elissa, Ur, Blase, 2018. "What was that site doing with my facebook password?": Designing password-reuse notifications. In: Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security. CCS '18, Association for Computing Machinery, New York, NY, USA, pp. 1549–1566. http://dx.doi.org/10.1145/3243734.3243767.

Goodchild, Paul, 2024. Rectifying google rankings: A primer on Japanese keyword hack recovery. (last Accessed 05 September 2024). https://getshieldsecurity.com/blog/japanese-keyword-hack/.

Halder, Stephan, 2025. Website redirects im umfeld von fake webshops und SEO fraud. (last Accessed 14 February 2025). https://www.bdo.de/de-de/insights/aktuelles/assurance/website-redirects-im-umfeld-von-fake-webshops-und-seo-fraud.

Harwell, Michael R., Rubinstein, Elaine N., Hayes, William S., Olds, Corley C., 1992. Summarizing Monte Carlo results in methodological research: The one- and two-factor fixed effects ANOVA cases. J. Educ. Stat. 17 (4), 315–339. http://dx.doi.org/10.3102/10769986017004315.

Hennig, Anne, Dietmann, Heike, Lehr, Franz, Mutter, Miriam, Volkamer, Melanie, Mayer, Peter, 2022a. "Your cookie disclaimer is not in line with the ideas of the gdpr. Why?" In: Human Aspects of Information Security and Assurance. HAISA 2022, In: IFIP Advances in Information and Communication Technology, vol. 658, Springer, Cham, pp. 218–227. http://dx.doi.org/10.1007/978-3-031-12172-2_17.

Hennig, Anne, Neusser, Fabian, Pawelek, Aleksandra Alicja, Herrmann, Dominik, Mayer, Peter, 2022b. Standing out among the daily spam: How to catch website owners' attention by means of vulnerability notifications. In: Extended Abstracts of the 2022 CHI Conference on Human Factors in Computing Systems (New Orleans, LA, USA). In: CHI EA '22, Association for Computing Machinery, New York, NY, USA, p. 8. http://dx.doi.org/10.1145/3491101.3519847, Article 317.

Hennig, Anne, Vuong, Nhu Thi Thanh, Mayer, Peter, 2023. Vision: What the hack is going on? A first look at how website owners became aware that their website was hacked. In: Proceedings of the 2023 European Symposium on Usable Security (Copenhagen, Denmark). EuroUSEC '23, Association for Computing Machinery, New York, NY, USA, pp. 312–317. http://dx.doi.org/10.1145/3617072.3617101.

Huang, Yue, Obada-Obieh, Borke, Beznosov, Konstantin, 2022. Users' perceptions of chrome compromised credential notification. In: Eighteenth Symposium on Usable Privacy and Security. SOUPS 2022, USENIX Association, Boston, MA, pp. 155–174, https://www.usenix.org/conference/soups2022/presentation/huang.

Kührer, Marc, Hupperich, Thomas, Rossow, Christian, Holz, Thorsten, 2014. Exit from hell? Reducing the impact of amplification DDoS attacks. In: 23rd USENIX Security Symposium. USENIX Security 14, USENIX Association, San Diego, CA, pp. 111–125, https://www.usenix.org/conference/usenixsecurity14/technical-sessions/presentation/kuhrer.

Kumar, Patchmuthu Ravi, Raj, Perianayagam Herbert, Jelciana, Perianayagam, 2019. A framework to detect compromised websites using link structure anomalies. Adv. Intell. Syst. Comput. 72–84. http://dx.doi.org/10.1007/978-3-030-03302-6_7.

Li, Frank, Durumeric, Zakir, Czyz, Jakub, Karami, Mohammad, Bailey, Michael, McCoy, Damon, Savage, Stefan, Paxson, Vern, 2016a. You've got vulnerability: Exploring effective vulnerability notifications. In: 25th USENIX Security Symposium. USENIX Security 16, pp. 1033–1050, https://www.usenix.org/conference/usenixsecurity16/technical-sessions/presentation/li.

Li, Frank, Ho, Grant, Kuan, Eric, Niu, Yuan, Ballard, Lucas, Thomas, Kurt, Bursztein, Elie, Paxson, Vern, 2016b. Remedying web hijacking: Notification effectiveness and webmaster comprehension. In: Proceedings of the 25th International Conference on World Wide Web. WWW '16, International World Wide Web Conferences Steering Committee, Montreal, Quebec, Canada, pp. 1009–1019. http://dx.doi.org/10.1145/2872427.2883039.

Li, Frank, Rogers, Lisa, Mathur, Arunesh, Malkin, Nathan, Chetty, Marshini, 2019. Keepers of the machines: Examining how system administrators manage software updates. In: Fifteenth Symposium on Usable Privacy and Security. SOUPS 2019, USENIX Association, Santa Clara, CA, pp. 273–288, https://www.usenix.org/conference/soups2019/presentation/li.

Lix, Lisa M., Keselman, Joanne C., Keselman, H.J., 1996. Consequences of assumption violations revisited: A quantitative review of alternatives to the one-way analysis of variance F test. Rev. Educ. Res. 66 (4), 579–619. http://dx.doi.org/10.3102/00346543066004579.

Lone, Qasim, Frik, Alisa, Luckie, Matthew, Korczyński, Maciej, Eeten, Michel van, Gañán, Carlos, 2022. Deployment of source address validation by network operators: A randomized control trial. In: 2022 IEEE Symposium on Security and Privacy. SP, SP, 00, 2361–2378.http://dx.doi.org/10.1109/sp46214.2022.9833701,

Maaß, Max, Clement, Marc-Pascal, Hollick, Matthias, 2021a. Snail mail beats email any day: On effective operator security notifications in the internet. In: The 16th International Conference on Availability, Reliability and Security. ARES 2021, ACM, New York, NY, USA, Vienna, Austria, pp. 1–13, https://dl.acm.org/doi/10.1145/3465481.3465743.

Maaß, Max, Pridöhl, Henning, Herrmann, Dominik, Hollick, Matthias, 2021b. Best practices for notification studies for security and privacy issues on the internet. In: The 16th International Conference on Availability, Reliability and Security. In: The 16th International Conference on Availability, Reliability and Security, Association for Computing Machinery, Vienna, Austria, pp. 1–10. http://dx.doi.org/10.1145/3465481.3470081.

Maaß, Max, Stöver, Alina, Pridöhl, Henning, Bretthauer, Sebastian, Herrmann, Dominik, Hollick, Matthias, Spiecker, Indra, 2021c. Effective notification campaigns on the web: A matter of trust, framing, and support. In: 30th USENIX Security Symposium. USENIX Security 21, USENIX Association, pp. 2489–2506, https://www.usenix.org/conference/usenixsecurity21/presentation/maass.

Martori, Art, 2020. Spamdexing: What is SEO spam and how to remove it. (last Accessed 28 October 2021). https://blog.sucuri.net/2020/02/spamdexing-seo-spam.html.

Mayer, Peter, Kunz, Alexandra, Volkamer, Melanie, 2017. Reliable behavioural factors in the information security context. In: Proceedings of the 12th International Conference on Availability, Reliability and Security (Reggio Calabria, Italy). ARES '17, Association for Computing Machinery, New York, NY, USA, p. 10. http://dx.doi.org/10.1145/3098954.3098986, Article 9.

Mayer, Peter, Zou, Yixin, Schaub, Florian, Aviv, Adam J., 2021. "Now I'm a bit angry:" Individuals' awareness, perception, and responses to data breaches that affected them. In: 30th USENIX Security Symposium. USENIX Security 21, USENIX Association, pp. 393–410, https://www.usenix.org/conference/usenixsecurity21/presentation/mayer.

mindUp Web & Intelligence GmbH, 2025a. Fake-online-shops - erkennung von fakeshops auf gehackten webseiten. (last Accessed 14 February 2025). https://www.mindup.de/data-scientists/anwendungsfaelle/fake-online-shops.

mindUp Web & Intelligence GmbH, 2025b. Gezieltes finden gehackter webseiten. (last Accessed 14 February 2025). https://www.mindup.de/nachrichten/artikel/gezieltes-finden-gehackter-webseiten.

Muniz, Caitlyn N., Fisher, Taylor, Smith, Katelyn, Ali, Roan, Howell, C. Jordan, Maimon, David, 2024. Hello, you've been hacked: a study of victim notification preferences. J. Crime Justice 1–17. http://dx.doi.org/10.1080/0735648X.2024.2340554.

Nosyk, Yevheniya, Korczyński, Maciej, Gañán, Carlos H., Król, Michał, Lone, Qasim, Duda, Andrzej, 2023. Don't get hijacked: Prevalence, mitigation, and impact of non-secure DNS dynamic updates. In: 2023 IEEE 22nd International Conference on Trust, Security and Privacy in Computing and Communications. TrustCom, Vol. 00, pp. 1480–1489. http://dx.doi.org/10.1109/trustcom60117.2023.00202.

O'Keefe, Daniel, 2015. Elaboration likelihood model. In: The Concise Encyclopedia of Communication.

Poteat, Tara, Li, Frank, 2021. Who you gonna call? an empirical evaluation of website security.txt deployment. In: Proceedings of the 21st ACM Internet Measurement Conference. IMC '21, pp. 526–532. http://dx.doi.org/10.1145/3487552.3487841.

2021. Princeton researcher apologizes for GDPR/CCPA email study . (last Accessed 18 April 2024). https://news.ycombinator.com/item?id=29650719.

Princeton University, 2021. Princeton-radboud study on privacy law implementation. (last Accessed 18 April 2024). https://privacystudy.cs.princeton.edu/.

Renaud, Karen, Searle, Rosalind, Dupuis, Marc, 2021a. Shame in cyber security: Effective behavior modification tool or counterproductive foil? New Secur. Parad. Work. 70–87. http://dx.doi.org/10.1145/3498891.3498896.

Renaud, Karen, Zimmermann, Verena, Schürmann, Tim, Böhm, Carlos, 2021b. Exploring cybersecurity-related emotions and finding that they are challenging to measure. Humanit. Soc. Sci. Commun. 8 (75), http://dx.doi.org/10.1057/s41599-021-00746-5.

Rodríguez, Elsa, Verstegen, Susanne, Noroozian, Arman, Inoue, Daisuke, Kasama, Takahiro, van Eeten, Michel, Gañán, Carlos H, 2021. User compliance and remediation success after IoT malware notifications. J. Cybersecur. 7 (1), tyab015. http://dx.doi.org/10.1093/cybsec/tyab015.

Rogers, Ronald W., 1975. A protection motivation theory of fear appeals and attitude change. J. Psychol. 1, 93–114. http://dx.doi.org/10.1080/00223980.1975.9915803.

Rostami, Asreen, Vigren, Minna, Raza, Shahid, Brown, Barry, 2022. Being hacked: Understanding victims' experiences of IoT hacking. In: Eighteenth Symposium on Usable Privacy and Security. SOUPS 2022, USENIX Association, Boston, MA, pp. 613–631, https://www.usenix.org/conference/soups2022/presentation/rostami.

Samarasinghe, Nayanamana, Mannan, Mohammad, 2021. On cloaking behaviors of malicious websites. Comput. Secur. 101, 102114. http://dx.doi.org/10.1016/j.cose.2020.102114, https://www.sciencedirect.com/science/article/pii/S0167404820303874.

Schmider, Emanuel, Ziegler, Matthias, Danay, Erik, Beyer, Luzi, Bühner, Markus, 2010. Is it really robust? Reinvestigating the robustness of ANOVA against violations of the normal distribution assumption. Methodology 6 (4), 147–151. http://dx.doi.org/10.1027/1614-2241/a000016.

Shafranovich, Y., Foudil, E., 2022. RFC 9116: A file format to aid in security vulnerability disclosure — datatracker.ietf.org. https://datatracker.ietf.org/doc/html/rfc9116. (last Accessed 12 January 2025).

SiteLock, 2022. Cybersecurity statistics report 2022. (last Accessed 10 October 2023). https://www.sitelock.com/resources/security-report/.

Soska, Kyle, Christin, Nicolas, 2014. Automatically detecting vulnerable websites before they turn malicious. In: 23rd USENIX Security Symposium. USENIX Security 14, USENIX Association, San Diego, CA, pp. 625–640, https://www.usenix.org/conference/usenixsecurity14/technical-sessions/presentation/soska.

Stock, Ben, Pellegrino, Giancarlo, Li, Frank, Backes, Michael, Rossow, Christian, 2018. Didn't you hear me? - Towards more successful web vulnerability notifications. In: Proceedings of the 25th Annual Symposium on Network and Distributed System Security. NDSS'18, pp. 1–15. http://dx.doi.org/10.14722/ndss.2018.23171, https://swag.cispa.saarland/papers/stock2018notification.pdf.

Stock, Ben, Pellegrino, Giancarlo, Rossow, Christian, Johns, Martin, Backes, Michael, 2016. Hey, you have a problem: On the feasibility of large-scale web vulnerability notification. In: 25th USENIX Security Symposium. USENIX Security 16, USENIX Association, Austin, TX, pp. 1015–1032, https://www.usenix.org/conference/usenixsecurity16/technical-sessions/presentation/stock.

StopBadware and Commtouch, 2012. Compromised websites: An owner's perspective. pp. 1–15, https://www.stopbadware.org/files/compromised-websites-an-owners-perspective.pdf.

Stöver, Alina, Gerber, Nina, Pridöhl, Henning, Maass, Max, Bretthauer, Sebastian, genannt Döhmann, Indra Spiecker, Hollick, Matthias, Herrmann, Dominik, 2023. How website owners face privacy issues: Thematic analysis of responses from a covert notification study reveals diverse circumstances and challenges. Proc. Priv. Enhancing Technol. 2023, 251–264, https://petsymposium.org/popets/2023/popets-2023-0051.php.

Sundaram, Karishma, 2022. Fix WordPress pharma hack and SEO. (last Accessed 19 April 2024). https://www.malcare.com/blog/what-is-pharma-hack-how-to-clean-it/.

Utz, Christine, Michels, Matthias, Degeling, Martin, Marnau, Ninja, Stock, Ben, 2023. Comparing large-scale privacy and security notifications. Proc. Priv. Enhancing Technol. 2023 (3), 173–193. http://dx.doi.org/10.56553/popets-2023-0076.

Vasek, Marie, Moore, Tyler, 2012. Do malware reports expedite cleanup? An experimental study. In: 5th Workshop on Cyber Security Experimentation and Test, CSET '12, Bellevue, WA, USA, August 6, 2012. USENIX Association, pp. 1–8, https://www.usenix.org/conference/cset12/workshop-program/presentation/vasek.

Wirth, Werner, Kühne, Rinaldo, 2013. Grundlagen der persuasionsforschung. Konzepte, theorien und zentrale einflussfaktoren. In: Handbuch Medienwirkungsforschung.

Wu, Qiushi, Lu, Kangjie, 2021. On the feasibility of stealthily introducing vulnerabilities in open-source software via hypocrite commits. https://api.semanticscholar.org/CorpusID:233479632.

Zeng, Eric, Li, Frank, Stark, Emily, Felt, Adrienne Porter, Tabriz, Parisa, 2019. Fixing HTTPS misconfigurations at scale: An experiment with security notifications. In: The 2019 Workshop on the Economics of Information Security (2019). Boston, MA, pp. 1–19, https://www.semanticscholar.org/paper/Fixing-HTTPS-Misconfigurations-at-Scale%3A-An-with-Zeng-Li/b22c522c6201f8545e1626deaf6ca43db52444d7.

Zhang, Jia, Duan, Haixin, Liu, Wu, Yao, Xingkun, 2017. How to notify a vulnerability to the right person? Case study: In an ISP scope. In: GLOBECOM 2017 - 2017 IEEE Global Communications Conference. pp. 1–7. http://dx.doi.org/10.1109/GLOCOM.2017.8253993.

Zou, Yixin, Danino, Shawn, Sun, Kaiwen, Schaub, Florian, 2019. You 'might' be affected: An empirical analysis of readability and usability issues in data breach notifications. In: Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (Glasgow, Scotland Uk). CHI '19, Association for Computing Machinery, New York, NY, USA, pp. 1–14. http://dx.doi.org/10.1145/3290605.3300424.

Zou, Yixin, Le, Khue, Mayer, Peter, Acquisti, Alessandro, Aviv, Adam J., Schaub, Florian, 2024. Encouraging users to change breached passwords using the protection motivation theory. ACM Trans. Comput. Hum. Interact. 31 (5), 45. http://dx.doi.org/10.1145/3689432, Article 63.