

Phishing in Organizations: Findings from a Large-Scale and Long-Term Study

Daniele Lain, Kari Kostiainen, and Srdjan Čapkun

Department of Computer Science

ETH Zurich, Switzerland

{daniele.lain, kari.kostiainen, srdjan.capkun} @inf.ethz.ch

Abstract—In this paper, we present findings from a large-scale and long-term phishing experiment that we conducted in collaboration with a partner company. Our experiment ran for 15 months during which time more than 14,000 study participants (employees of the company) received different simulated phishing emails in their normal working context. We also deployed a reporting button to the company’s email client which allowed the participants to report suspicious emails they received. We measured click rates for phishing emails, dangerous actions such as submitting credentials, and reported suspicious emails.

The results of our experiment provide three types of contributions. First, some of our findings support previous literature with improved ecological validity. One example of such results is good effectiveness of warnings on emails. Second, some of our results contradict prior literature and common industry practices. Surprisingly, we find that embedded training during simulated phishing exercises, as commonly deployed in the industry today, does not make employees more resilient to phishing, but instead it can have unexpected side effects that can make employees even more susceptible to phishing. And third, we report new findings. In particular, we are the first to demonstrate that using the employees as a collective phishing detection mechanism is practical in large organizations. Our results show that such crowd-sourcing allows fast detection of new phishing campaigns, the operational load for the organization is acceptable, and the employees remain active over long periods of time.

I. INTRODUCTION

Phishing remains a major problem on the Internet [1]. Deceptive emails that trick users to perform unsafe actions are getting increasingly sophisticated [2], [1] and during the last two decades phishing showed no sign of slowing down [3]. The job of cyber-criminals is made easy by the development of *phishing kits*, software capable of automatically creating deceptive copies of popular websites [4], [5], [6]. To make things even worse, the COVID-19 pandemic has shifted work, shopping and other activities online which in turn has created new phishing opportunities and increased phishing [7].

Researchers have studied phishing for decades (see [8], [9], [10], [11] for extensive reviews of early works) and proposed various defenses from email filters [10], [12], to detection of phishing websites [13], patterns of phishing campaigns [14], triggers that push people to fall for phishing [15], and ways to educate users [16], [11]. During the last decade, also an entire ecosystem of companies that provide phishing prevention products and services has emerged. Common commercial offerings include training and educational services [17], [18], [19], [20], databases of known URLs as well as emails used

by phishing attacks [21], [22], [23], and email filters powered by threat intelligence collected by specialists and reports from customers [24], [18], [19].

Our study and contributions. In this paper, we study phishing with a particular focus on phishing in *organizations*. We approach this topic through the following four questions – all related to human factors of phishing. First, we are interested to understand *which employees are the most vulnerable* to phishing in large organizations. We examine this through common aspects like employee demographics and job type. Second, we explore how the organization’s *phishing vulnerability evolves over time*. For instance, we study how many employees will eventually fall for phishing in continued exposure to phishing. Third, we study *how organizations can help their employees* in phishing prevention. In particular, we analyze the benefits of currently popular tools such as embedded phishing training and warnings on top of suspicious emails. And fourth, we explore whether *the employees can collectively help the organization* in phishing prevention. Regarding this question, we focus on using the employees as a collective phishing detection sensor – an idea that has been previously suggested [25], [23], but prior to our work, its effectiveness and feasibility has not been publicly evaluated in a real large organization.

To answer these questions, we designed and conducted a large-scale and long-term phishing study in collaboration with a partner company. Our study ran for 15 months (July 2019–October 2020) and during it 14,773 employees of the company became participants in our experiment. Our study involved sending simulated phishing emails to the participants, who received them as part of their normal work flow and context. We measured their click rates, submission of credentials, and enabling macros on attachments. We also deployed a reporting button to the corporate email client which allowed our study participants to easily report emails that they found suspicious, and analyzed the reported emails.

To the best of our knowledge, our experiment is the first study of phishing in organizations that is at the same time large-scale (14k participants), long-term (15 months), realistic (we measure real employees’ phishing behavior in their actual working context), and diverse (including participants across various corporate departments and job roles). All comparable, previous studies are either smaller [26], [27], [28], [29], [30], [31], [32], [33], [34], [35], [36], [37], [38], [16],

[39], shorter [36], [35], [26], [16], [33], [27], based on role-play [40], [41], [29], [32], [31], or less diverse [41], [34], [29], [31], [28], [27], as we will elaborate in Section II.

The results of our experiment provide three types of contributions. First, we report several results that *support* previous literature with increased ecological validity (e.g., more study participants, longer study duration, or more realistic study setting). Among others, we find email warnings are effective and observe many “repeated clickers” [42] in our experiment.

Second, our study uncovers a few findings that *contradict* both the conclusions of previous academic studies and common industry practices. In particular, we find that embedded phishing training, as commonly used in the industry today, can lead to unexpected side effects and even be detrimental to phishing prevention. This is a significant finding, due to wide use of this practice in the industry.

And third, our results provide *new insights* to phishing in organizations. In particular, as one of the main contributions of this paper, our experiment is the first to demonstrate that crowd-sourced phishing detection can be effective, fast, and sustainable over long periods of time. During our experiment, the employees reported thousands of suspicious emails which represented hundreds of real and previously unseen phishing campaigns. The reporting speed of our simulated phishing emails indicates that new campaigns can be detected within few minutes from their launch. We designed a simple processing pipeline that combined automated and manual analysis for the reported emails. Our experiment shows that through such techniques, the operational load of handling all the reported emails can be made low even in large organizations. Our experiment also demonstrates that large employee bases can collectively retain sufficiently high reporting rates over long periods of time. In summary, this paper is the first to demonstrate that crowd-sourced phishing detection is a practical and effective option for many organizations.

To summarize, this paper makes the following contributions:

- 1) *Extensive measurement study* on human factors of phishing and phishing prevention in large organizations.
- 2) *Supportive results* for several previous research findings with improved ecological validity.
- 3) *Contradicting findings* that challenge the conclusions of previous research studies and popular industry practices.
- 4) *Large-scale evaluation of crowd-sourced phishing reporting* that shows fast detection, small operational overhead, and sustained employee reporting activity.

Paper outline. This paper is organized as follows. In Section II, we define our research questions and provide an overview of our findings. We describe our experimental setup in Section III. We report results related to employee demographics in Section IV. Section V shows how phishing vulnerability evolved over time in our study. Section VI explains our results related to warnings and embedded training. Section VII analyzes crowd-sourced phishing detection. In Section VIII we discuss validity of our study. Section IX reviews related work and Section X concludes the paper.

II. RESEARCH QUESTIONS AND MAIN FINDINGS

In this section, we first define the research questions that our study was designed to answer and then provide a summary of our main findings, both summarized in Table I.

A. Research Questions

RQ1: Which employees fall for phishing? The first goal of our experiment was to understand which employees in a large organization are the most likely to fall for phishing. In particular, we wanted to understand how employee characteristics that are easily available to organizations, such as age, gender, and the assumed level of computer use in one’s job type, correlate with phishing susceptibility.

RQ2: How does organization’s vulnerability to phishing evolve over time? The second goal of our experiment was to understand how the continued presence of phishing affects organizations over time. We examine topics like how large is the fraction of the employee base that will eventually fall for phishing, and how many individuals repeatedly fall for phishing [42].

RQ3: How effective are phishing warnings and training? Our third goal was to understand how large organizations can help their employees to recognize phishing emails and thus defend themselves against phishing. Today, organizations can choose from a range of tools and educational measures designed for this purpose. In our study, we focused on evaluating tools that can be deployed to a large employee base with a moderate cost, as such tools are commonly used in practice.

The first tool whose effectiveness we decided to examine was *warnings* on top of suspicious emails. Warnings are used in many popular email clients and services such as Gmail [43]: they are shown on top of the emails where automated phishing detection mechanism has identified some risky or suspicious features in the email, but it could not label the email as phishing with sufficiently high confidence (often email filters are tuned to be permissive to avoid too many false positives).

The second tool we wanted to test was *simulated phishing exercises* [32], [11] in combination with *embedded training* [33]. During the last decade, simulated phishing exercises have become a common industry practice [18], [19], [17], [20]. In a simulated phishing exercise, the organization sends emails that mimic real phishing emails to their employees and then track which employees perform unsafe actions such as clicking links or disclosing credentials to a web page. Often such exercises are combined with embedded training (sometimes also called contextual training), where employees that fail the exercise (e.g., by clicking on a link or disclosing their credentials) are forwarded to an information resource like a web page that provides educational material about phishing.

RQ4: Can employees help the organization in phishing detection? The fourth goal of our experiment was to understand if a large employee base can collectively help the organization in phishing prevention. More precisely, we wanted to understand whether using employees as a *crowd-sourced phishing*

TABLE I: **Summary of our research questions and main results** that include findings that support prior literature, findings that contradict previous studies, and new insights. The two most significant contributions of this paper are **marked in bold**.

	Findings that support previous studies in literature	Findings that contradict previous studies in literature	New findings on phishing in organizations
RQ1: Which employees fall for phishing? (Section IV)	Age and computer skills correlate with phishing susceptibility [40], [41], [34], [36], [35], [29], [11], [26]	Gender does not correlate with phishing susceptibility (contradicts [36], [30], [40], [41])	Type of computer use is more predictive for phishing vulnerability than amount of computer use
RQ2: How does organization's vulnerability to phishing evolve over time? (Section V)	There are several "repeated clickers" in a large organization [42]		Many employees will eventually fall for phishing if continuously exposed
RQ3: How effective are phishing warnings and training? (Section VI)	Warnings on top of suspicious emails are effective [38], [37], [39]	Voluntary embedded training in simulated phishing exercises is not effective (contradicts [33], [32], [34])	More detailed warnings are not more effective than simple ones
RQ4: Can employees help the organization in phishing detection? (Section VII)			Crowd-sourcing phishing email detection is both effective and feasible

detection mechanism is efficient (can phishing campaigns be detected fast enough?), practical (does the administrative load of processing the reported emails remain acceptable?), and sustainable (will employees continue to report emails over time?) in a large organization. Additionally, our aim was to understand if the presence of a feedback mechanism encourages employees to report suspicious emails more.

B. Summary of Main Findings

Next, we provide an overview of the main findings of our experiment, and briefly discuss how these results relate to prior research literature (a more detailed survey of related work is given in Section IX). Our findings *support* claims of previous studies with improved ecological validity; *contradict* prior conclusions and common industry practices; and provide *new insights* related to phishing in large organizations.

Findings related to RQ1. The results of our experiment support previous work which has showed that age [40], [41], [34], [36], [35], [29] and computer skills [11], [26] both correlate with phishing vulnerability. Similar to previous studies, we also find that older and younger employees are more at risk, as well as people with lower computer skills. Our experiment improves the ecological validity of these studies which were either smaller [36], [35], [30], [34], [29], [26], shorter in duration [36], [35], [26], featured populations with less diversity (e.g., mostly university students and employees [41], [34], [29], skewed in age [35]), or featured role-play or quiz-style studies only [40], [41], [29]. Opposed to previous literature [36], [30], [40], [41], we do not find gender to correlate with phishing susceptibility. The correlation that we observe is explained much better by skewed distribution of the different types of jobs among genders. We improve on these studies by reporting on a larger and more diverse population in their day-to-day job environment.

As a new finding, our study shows that the most vulnerable employees are those who use computers daily for repetitive

task with a specialized software only, rather than those employees who do not need computers in their day-to-day job. That is, in our experiment, the *type* of computer use is more predictive for phishing vulnerability than the *amount*. We discuss these topics more in Section IV.

Findings related to RQ2. Similar to previous studies, we find several "repeated clickers" who fail simulated phishing exercises multiple times [42]. We also find that if exposure to phishing continues in an organization, eventually a significant fraction of employees will fall for phishing. We elaborate on these results in Section V.

Findings related to RQ3. Our results support the previous studies that find contextual warnings effective [38], [37], [39] and the common industry practice of using such warnings [43]. We improve these studies thanks to a larger, more general population with a rigorous control group.

Interestingly, contradicting prior research results [33], [32], [34] and a common industry practice [19], [17], [20], [18], we found that the combination of simulated phishing exercises and voluntary embedded training (i.e., employees were not required to complete the training) not only failed to improve employee's phishing resilience, but it actually even made employees more susceptible to phishing. Compared to our experiment, previous studies featured less participants [33], [31], [28], [16], [32], [27], were shorter in time [16], [33], [27], had populations with little diversity [31], [28], [27] or tested a role-playing setting only [32], [31]. Our results suggest caution in the design of embedded training: we discuss the possible reasons (such as false sense of corporate IT security) and practical implications of this somewhat surprising and non-intuitive finding at length in Section VI.

Another novel finding of our study is that adding more details to contextual warnings (e.g., explaining the reasons the email was flagged as suspicious) does not reduce phishing effectiveness significantly.

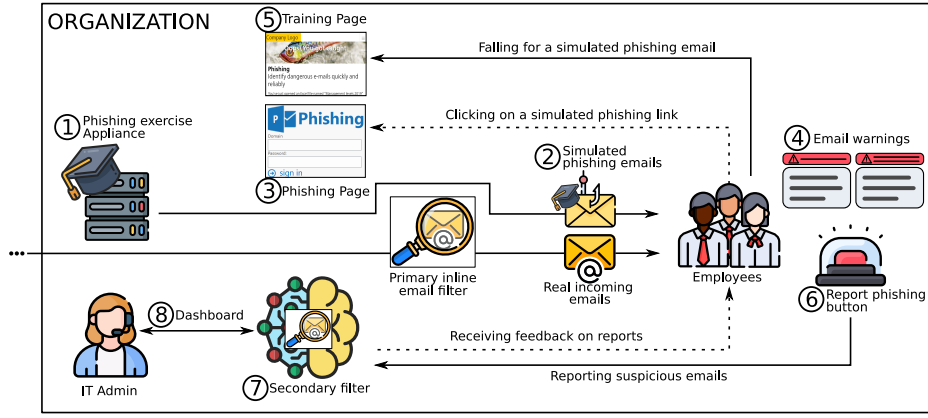


Fig. 1: Overview of the measurement infrastructure that we deployed in the partner company.

Findings related to RQ4. One of the main contributions of this paper is that we demonstrate experimentally that crowd-sourced phishing detection can be efficient and sustainable in large organizations. The idea of crowd-sourcing phishing detection to employees has been suggested in previous papers [25], [23]. Our contribution is that we are the first to evaluate this idea over a long period of time in the context of a real large organization.¹ Our experiment shows that crowd-sourced phishing detection enables organizations to detect a large number of previously unseen real phishing campaigns with a short delay from the start of the campaign. The processing pipeline that we developed as part of our experiment also shows that the operational load of phishing report processing can be kept small, even in large organizations. Our study also demonstrates that a sufficiently high number of employees report suspicious emails actively over long periods of time. In summary, we show that crowd-sourced phishing detection provides a viable option for many organizations. Section VII provides full discussion of this topic.

III. EXPERIMENTAL SETUP

In this section we explain how we performed this study in collaboration with a partner company.

A. Study Organization

Partner company. For this study, we collaborated with a company which employs more than 56,000 people of diverse technical skills, age groups, and jobs. Our partner company is a large public company, dealing in logistics, finance, transport, and IT services. They employ people with different duties: field workers, branch workers that work in front-end stores in contact with the general public, and office workers of different qualifications, from IT to marketing and accounting.

¹Phishing reporting by users is a widely-used industry practice. For example, there are service providers who aggregate data from many of their business customers [24], [19] and large email providers who collect reports to feed machine learning models [44]. However, prior to our work, it has not been publicly evaluated whether the employee base of a single organization can be effectively leveraged as a phishing detection mechanism.

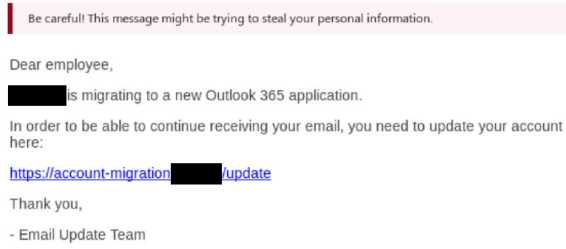
As is common industry practice [19], [17], [20], [18], at the time of study planning our partner company was already running a phishing awareness campaign which included simulated phishing emails and contextual (embedded) training.

Our role. For this study, we leveraged the already existing phishing awareness campaign as a testbed for our research questions. More precisely, we collaborated with our partner company in two ways. First, as a scientific advisor who helped in the design of the experiment. By deploying different tools and conditions to different employees, we were able to use the existing campaign to address our research questions. At the end of the study, we received anonymized data from the company in bulk and analyzed it. Second, we took an active part by administering a questionnaire to randomly-selected employees and analyzing the responses.

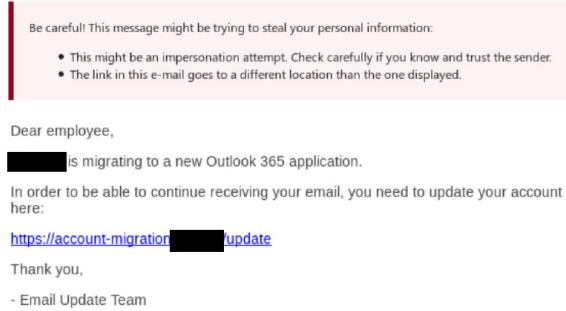
Company's role. The role of the company was three-fold: it designed all the simulated phishing emails; provided the infrastructure for sending the simulated phishing emails and measuring dangerous actions such as clicks; and hosted the embedded training resources (an educational webpage that was shown to those employees who performed a dangerous action). The company had a pre-existing collaboration with an external service provider that specializes in phishing awareness and education. This service provider assisted the company in phishing email and contextual training page design. The study was initiated and approved by the CISO of the company.

B. Measurement Infrastructure

Phishing exercise component. Our partner company deployed a phishing exercise component, shown as ① in Figure 1, implemented by the service provider who specializes in phishing awareness and training, that sent simulated phishing emails ② crafted by human experts. These emails could either link to a deceptive website (hosted by component ③) or have a malicious file attached, with the goal of deceiving the participant to do a *dangerous action*, such as submitting their credentials or enabling macros on an attachment.



(a) Short warnings.




(b) Detailed warnings.

Fig. 2: Warnings that we added to selected participants’ email clients on top of simulated phishing emails.

Deployed warnings. Based on our recommendation, the company deployed two types of warnings ④ that could be triggered to appear on top of the simulated phishing emails on the employees’ email client (Outlook). As a baseline, we deployed *short* warnings (Figure 2a), visually identical to the standard Outlook warnings that employees are used to, containing a similar generic sentence, warning the recipient to be careful because the email “looks suspicious”.

We also developed and deployed *detailed* warnings, shown in Figure 2b, again visually identical to Outlook warnings, but adding a list of reasons why the email might be suspicious, e.g., mismatches between the email of the sender and the displayed name, or mismatches between the displayed link and the pointed domain. Such information could be generated automatically in a deployment that adds warnings to emails that seem suspicious, when there is not sufficient certainty to block the email.

Deployed training. The phishing exercise component also hosted a training web page on phishing ⑤ shown after someone performed the dangerous action of a simulated phishing email. This internal corporate web page (part of it shown in Figure 3) explained to the employee what happened in detail (i.e., that they failed a phishing exercise from their organization), specific cues one should have paid attention to in the email, tips to avoid phishing in the future, an instructional video, and further quizzes and learning material on phishing. The training page was developed according to the best practices in academia [45], [31] and industry [19], [17] by the external service provider; we provide an excerpt in Appendix A. The training page was delivered to employees such that there was no enforcement that the employee has to



Phishing
Identify dangerous e-mails quickly and reliably

Fig. 3: Header of the contextual training awareness web page that was displayed after falling for a simulated phishing email.




Fig. 4: Menu bar of the company’s email client (Outlook), modified to include a button to report suspicious emails.

read the whole webpage or take the quizzes.

Reporting button. Our partner company deployed a button ⑥ for reporting suspicious emails. This button was introduced in the Outlook client, as shown in Figure 4, and it was advertised in the internal news of the company before the start of the experiment. When reporting a suspicious email, employees could toggle a checkbox to report that they also opened the attachment or visited the link in the email, to notify the IT department about a possible incident.

Reported email processing. All emails that were reported by our study participants were triaged by a commercial anti-phishing appliance ⑦ that ran a more-detailed secondary analysis. The secondary analysis performed on the reported emails differed from the company’s primary inline filter in two ways: (a) more time consuming checks, such as following links, were performed, and (b) the analysis settings were tuned to be more aggressive, as at this point we did not need to avoid too many false positives. The results of the secondary analysis were presented to the company’s IT department via a dashboard, where the appliance verdicts could be either confirmed or subverted based on manual analysis ⑧. Further, the appliance could return feedback to the employees, indicating whether the reported email was indeed malicious or not.

C. Study Participants

The company enrolled 14,733 employees to be part of the experiment: we refer to them as *participants*. Participants were selected uniformly at random from the whole company’s employee base, comprising many different job types, from accountants, IT, marketing, and managerial roles, to less technical jobs (e.g., in logistics, or working in retail shops). In total, participants spanned 28 organizational groups of the company, and represented 3,827 different teams.

For the purposes of our study, we classified the participants in terms of their age, gender and computer use in the day-

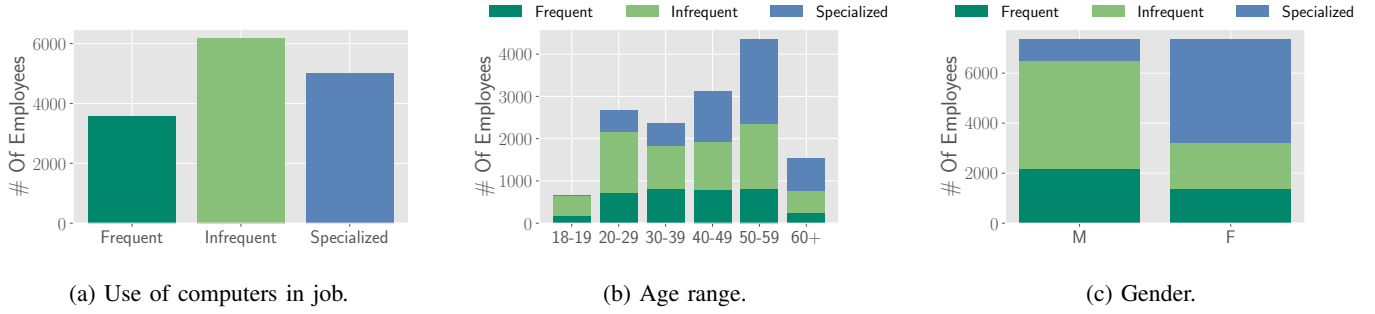


Fig. 5: Demographics information of the study participants. Age and gender are further divided by computer use in their job.

to-day job. Participants were divided by computer use (shown in Figure 5a) in three different categories: (i) office workers using computers daily, either in IT-related jobs or in jobs that use computers such as marketing and accounting (*Frequent use*); (ii) employment roles such as retail shop workers in contact with the general public, that mostly use a point-of-sales software and configure services from it (*Specialized use*); and (iii) roles such as team leaders of field workers in logistics who have a corporate email account but rarely use computers in their duties (*Infrequent use*). Study participants spanned an age range 18-73, with all age groups being well represented in our set of participants (Figure 5b). The gender distribution (Figure 5c) of our participants was balanced: 7,377 were male, and 7,356 were female, similar to the distribution of the company’s employee base. We observe an imbalance in the use of computers for gender and age: the majority of the branch workers that mostly use one specialized program are female, and are skewed to older ages, while users that work with or without computers are more uniformly distributed.

D. Study Group Sampling

Following on our advice, the company assigned each of the 14,733 participants to one of 12 different *user groups* generated by combining the settings administered for different tested tools and mechanisms:

- **Warnings (3 settings):** every participant received one of three possible settings on their simulated phishing emails: a simple warning; a detailed warning; or no warning, as a control setting.
- **Training (2 settings):** every participant that failed one simulated phishing attack by performing the dangerous action could either be redirected to the training page; or receive no such training, as a control setting.
- **Report feedback (2 settings):** after reporting a suspicious email, participants could always receive the result of their report as feedback; or they could receive the result only when they reported a legitimate email, as a control setting.

For example, Group 1 was administered simple warnings, training, and no feedback after correctly reporting phishing, while Group 2 had the same configuration except receiving complex warnings, and so on. Each participant was randomly

assigned to one of the $3 \times 2 \times 2 = 12$ groups, so that they were approximately of the same size: from 1,223 participants on the smaller to 1,231 participants on the larger.

E. Experiment Execution

From July 2019 to October 2020, the company sent 8 different simulated phishing emails to each of the 14,733 participants. The participants received the first 6 emails in random order and at random time intervals [16] during the first 12 months of the experiment (July 2019–July 2020). They received the last two phishing emails² from August 2020 to October 2020, again in random order and at random time intervals. Participants were not aware of our study specifically, to not modify their behavior [47], [48]; however, they were aware that the company may occasionally send phishing exercises to their employees.

The 8 different email campaigns, of varying difficulty, were designed to simulate broad phishing campaigns targeted to the organization as a whole, rather than sophisticated, individually-crafted spear phishing. Each different email represented a typical phishing scenario, such as prompts to check their corporate credentials, migrate their email accounts to a new system, or parcel delivery notes as attachments, and used different triggers, such as a sense of authority or urgency or leveraging people’s curiosity [27]. Five emails contained a link to a phishing website, while three had an attached file. We provide the English version of selected emails in Appendix B.

During the experiment, our partner company recorded the following interactions with the simulated emails:

- **Clicks** on the links contained in the email;
- **Dangerous actions:** further falling for the phish by, e.g., submitting credentials to the linked website, or enabling macros on the attached document.

The company also recorded participants’ reports of suspicious emails. For each reported email, they stored whether it was one of our simulated phishing emails, the result of the secondary analysis by the anti-phishing appliance, and whether any employee from the IT department looked at such result and confirmed or subverted its verdict. During the last 5 months of

²The last emails were supposed to be three; however, a simulated *CEO fraud* phishing attack [46] caused unwanted confusion inside the company and this specific simulated phishing email was canceled.

the experiment, the company also recorded how many inbound emails were similar to a reported one in a 20-days window around the date of the report.

At the end of the experiment, we administered a questionnaire with 27 closed-ended questions to 1000 randomly selected participants. Participants that accepted to respond were informed that their replies were recorded anonymously and would further not be shared with their employer, to encourage honest answers. The first questions asked participants about knowledge of phishing and other email threats, questions about email warnings, the button to report phishing, contextual training, and whether they recalled falling for phishing. We report selected questions from the questionnaire in Appendix C. We received 151 complete answers.

F. Ethics and Safety

Study approval. This study was initiated and approved by the CISO of our partner company. During the study, we never had access to any PII, and were only given access to anonymized data after collection by the company (see Section III-A). Since the analysis of anonymized data does not require IRB approval according to our institution’s guidelines, we did not submit a formal request.

Risks to participants. Our partner company informs its employees about their phishing awareness campaign, that includes phishing exercises. Thus, our study participants were generally aware that the company may send them simulated phishing emails. The company did not specifically inform participants about the simulated phishing emails that were sent as part of this study (i.e., no informed consent or debriefing). Participants not in the embedded training group were not specifically informed about the simulated phishing emails, while participants in the embedded training group were informed that the email was phishing if they fell for it.

Our participants were subject to minimal risk as part of this experiment: they were not exposed to greater risk than what they encounter as part of their normal daily life [49], because they receive real phishing and other malicious emails regularly. Experiments such as the one we conducted here can have negative impacts such as wasting employees’ time or creating distrust towards the company [49]. This experiment took place as part of the company’s existing training program; given this context, we felt that the scientific impact of our experiment merited these potential negative impacts

Data collection and protection. During the study, our partner company collected data regarding clicks and dangerous actions, and data on emails reported as phishing by participants. If a study participant entered their password on the simulated phishing web page, our partner company did not record the entered credentials nor checked if they were correct. The collected dataset was accessible to a small number of employees working in the IT security department of our partner company and protected with two-factor authentication.

The collected dataset was provided to us in anonymized format such that only attributes like gender, age, and level

of computer use were preserved. Our partner company used the dataset internally to assess its overall exposure to phishing threats, and ensured us that the dataset will not be used for any other purpose, such as employee performance assessment.

The reported emails did not carry any PII: every report recorded whether the reported email was a simulated one or not, and scores and verdicts of the anti-phishing appliance. None of these information can link to the original sender, subject, or content of the message.

G. Experiment Statistics

Overall, the study participants clicked on 6,680 out of 117,864 simulated phishing emails (5.67%). During the 15 months, 4,729/14,733 participants (32.10%) clicked on at least one phishing email. The trend for dangerous actions is similar, with the numbers slightly lower: participants fell for 4,885 simulated phishing emails (4.14% of the total sent emails, and 73.13% of all the clicked simulated phishing emails), and 3,747/14,733 participants (25.43%) users did at least one dangerous action.

There were 4,260 study participants that reported at least one email. In total, the participants reported 14,401 emails, of which 11,035 were our simulated emails. The button to report phishing was also deployed to 6300 employees that were not part of the experiment but could report phishing: 1,543 of them reported at least one suspicious email, and they reported 4,075 emails. Thus, the total number of reported emails we received during the 15 months was 18,476.

IV. WHICH EMPLOYEES FALL FOR PHISHING?

In this section, we analyze the experiment data to understand which employees are the most likely to fall for phishing (RQ1). Recall from Section III that we classify participants based on Frequent, Infrequent and Specialized use. We count the number of clicked links and dangerous actions based on demographics and job categories (see Figure 6). For our following analysis, we define the following three hypotheses:

- H1: *Employees’ use of computers in their job correlates to falling for phishing.*
- H2: *Employees’ age correlates to falling for phishing.*
- H3: *Employees’ gender correlates to falling for phishing.*

To analyze the measured numbers, we fit a linear model with Type III sum of squares to analyze both the demographic properties by themselves, and to capture the interactions among them. This statistical tool allows us to measure the impact of the independent variables (i.e., the demographic properties) on the dependent variables: number of clicked links and dangerous actions, that we use as proxies for phishing susceptibility. We fit the model with all the combinations of demographic properties, and exclude the non-significant factors until we obtain a final model with following results.

The results support H1: correlation with computer use. As can be seen from Figure 6a, participants whose job type involve Specialized computer use (e.g., branch workers who mostly use a single dedicated program) clicked on more links in phishing emails and performed more dangerous actions




Fig. 6: Percentage of dangerous actions performed out of all phishing emails sent, divided by different demographics. Frequent use of computers but in a very specialized setting, and young and older age all influence the susceptibility to phishing.

than participants in the other comparable groups (Frequent and Infrequent use). Our fitted model shows that computer use is significant (clicks: $F(2, 14710) = 11.01, p < 0.001$; dangerous actions: $F(2, 14710) = 9.45, p < 0.001$) and a Tukey HSD post-hoc test confirms that the difference between Specialized use and the other two groups is significant both for clicks and dangerous actions. However, the difference between Frequent and Infrequent use is not significant. Thus, while we support previous work that showed relationship between phishing susceptibility and knowledge of technology [11], this last observation invites to caution, as this relationship seems more nuanced. While it is common to leverage the *amount* of computer use in participants’ jobs as a proxy for technological skills, our results suggest that the *type* of computer use and the expectations in one’s job might also influence phishing susceptibility. For example, Specialized use participants in our partner organization may be expected to interact with emails more than Infrequent use participants, who may, therefore, be more suspicious of incoming emails.

The results support H2: correlation with age. The youngest employees clicked more and performed more dangerous actions. Our model confirms the interaction between age and phishing susceptibility (click rate $F(5, 14710) = 4.70, p < 0.001$; dangerous action rate ($F(5, 14710) = 3.84, p < 0.001$). We ran a Tukey HSD test to analyze which groups were more at-risk and confirm what Figure 6b shows: participants aged 18–19 were much more likely to click on phishing links and perform the dangerous action than any other age group; participants in the 50–59 age range were also more at risk than the top performers aged 20–29 and 60+. This result supports previous literature [40], [41], [34].

The results do not support H3: correlation with gender. Our participants’ computer use w.r.t. their gender is not uniform (recall Figure 5c). Thus, further dividing interactions of both genders by use of computers shows a large difference among the same gender, shown in Figure 6c and confirmed by our model: the combination of gender and computer use is significant (click rate $F(2, 14710) = 13.06, p < 0.001$), but gender by itself is not (click rate $F(2, 14710) = 0.23, p = 0.63$). Indeed, Figure 6c shows us that while Frequent use females were




Fig. 7: Number of simulated phishing emails that participants clicked or performed the dangerous action (8 being maximum; a missing bar denotes zero participants).

more susceptible than Frequent males, Specialized use males were more susceptible than their female counterpart. Thus, phishing susceptibility of participants can be better explained by considering the imbalance in job types, contradicting some previous studies [36], [30].

V. PHISHING VULNERABILITY OVER TIME

In this section, we leverage our 15-month study to analyze how the phishing susceptibility of the organization evolves over time (RQ2). To do so, we analyze trends of clicks and dangerous actions over time: how many times (out of maximum 8) participants interacted with the phishes, and how many participants over time eventually did so at least once.

Repeated clickers. We report in Figure 7 the histogram of how many participants clicked or performed the dangerous action on the simulations a given amount of times. A total of 1,448 (30.62%) participants clicked on two or more phishes, and 896 (23.91%) performed the dangerous action on two or more—one participant even fell for 6 out of 8 simulations. Thus, we observe that there will be a small number of employees that will click or fall for phishing emails multiple times, supporting a previous preliminary study [42]. Similarly to the raw amount of clicks and dangerous actions, we observe a correlation between age groups and clicking (Welch-corrected ANOVA

$F(5, 4199) = 5.72, p < 0.001$) or performing the dangerous action ($F(5, 4186) = 3.66, p = 0.002$) on more than one simulated phishing emails. In both cases, a Tukey HSD test shows that the younger group of participants aged 18-19 stands out as the one more likely to click more than once.

Many employees will eventually fall for phishing if continuously exposed. In our experiment 4,729 out of 14,733 (32.10%) participants clicked on at least one link or attachment in our simulated phishing emails. A similar high number applies to dangerous actions: 3,747 out of 14,733 (25.43%) performed at least one. These results indicate that a rather large fraction of the entire employee base will be vulnerable to phishing when exposed to phishing emails for a sufficiently long time. We are the first to show such result at scale.

VI. EFFECTIVENESS OF WARNINGS AND TRAINING

In this section we analyze the data collected from our experiment to answer RQ3 related to the effectiveness of phishing warnings and training.

A. Effectiveness of Warnings

Recall from Section III that we experimented with two types of warnings (short and detailed), and a control group that did not see any warnings. To analyze the effectiveness of these two warning types, we use the following hypotheses:

- H4: *Adding warnings on top of suspicious emails helps users in detecting phishing.*
- H5: *Detailed warnings are more effective than short ones.*

The results support H4: warnings help users. Figure 8 shows click and dangerous action rate for the different warning configurations. We observe that both types of warnings greatly helped participants both in avoiding clicking on links in our simulated phishing emails and not falling for the phish by performing the dangerous action. Considering click rate, the group with no warnings clicked 3,964 times, compared to the lower 1,427 clicks for short and 1,289 clicks for long warnings (Welch-corrected ANOVA $F(2, 7485) = 564.71, p < 0.001$). Dangerous actions rate is similar: 2,994 dangerous actions for no warnings, compared to 998 and 893 dangerous actions, respectively ($F(2, 7461) = 392.58, p < 0.001$). Figure 9 shows the histogram of how many participants clicked on a simulated phish a given amount of times. We observe a strong correlation between receiving any of the warnings and not clicking or performing the dangerous action more than once (clicks: $F(2, 9287) = 358.88, p < 0.001$, dangerous actions: $F(2, 9194) = 239.68, p < 0.001$). Our results support this widespread industry practice [43].

The results do not support H5: detailed warnings are not more effective than short ones. To check whether there is any difference between short and detailed warnings, we ran a Tukey HSD test between all groups and observed that, while both warnings correlate with lower total clicks and dangerous actions, there is no significant difference between short and detailed warnings. Thus, the way we provided additional information to users (by mimicking the current industry practices,




Fig. 8: Dangerous actions by administered warning. Both warning types helped the participants significantly.




Fig. 9: Number of different phishing emails that participants clicked on, by administered warning. Missing bars denote a 0.

rather than making radical changes to email warnings [37]) does not seem to provide better phishing protection.

B. Effectiveness of Contextual Training

Recall from Section III that we tested the effectiveness of contextual training after falling for a simulated phishing email by administering it only to half of the participants—the other half was a control group for training and did not see the webpage. We formulate the following hypothesis:

- H6: *Receiving contextual training helps users improve in future phishing detection.*

We analyze both the frequency at which participants clicked or performed the dangerous action, and the correlation between training and doing more than one click or dangerous action.

The results do not support H6: voluntary contextual training does not improve future phishing detection. Surprisingly, we observe that both click and dangerous actions rates are higher for participants that received contextual training (i.e., participants who were forwarded to a training page) after falling for simulated phishes: for clicks, 3,087 versus 3,593; for dangerous actions, 2,155 versus 2,730. Figure 10 shows the histogram of how many different phishing emails participants performed the dangerous action on. As expected, the number of participants that did not fall for any simulated phish, or fell only once, is similar among the two groups: such participants

that were in the training group either never saw the training page, or saw it after performing their only dangerous action. However, if we focus on participants that fell two or more times (and thus, on participants in the training group that fell again for phishing after being shown the training page), we see that the distribution is more skewed to the right for participants in the training group. Indeed, participants that clicked on two or more phishing emails were 647 without training, and 801 with training. This shows a strong correlation between the provided training page and clicking on phishing emails or even performing the dangerous action more than once (Welch-corrected ANOVA for clicking: $F(1, 14592) = 18.37, p < 0.001$; dangerous actions: $F(1, 14279) = 33.80, p < 0.001$).

This perhaps surprising result requires a careful interpretation. What our experiment showed is that *this particular* way of delivering voluntary training does not work. Instead, such training method may cause unexpected and negative side effects, such as increased susceptibility to phishing. This finding is significant, because the tested phishing training delivery method is a common industry practice [19], [17], [20], [18], and the training material (refer to Section III) was designed by a specialized company according to known guidelines and best practices from previous work [31], [34], [45]. It would be interesting to study whether other possible ways to deliver contextual training (e.g., ones where interaction with the provided training material is enforced) would work better. Our study did not test the effectiveness of mandatory training.

To gain some insights on why susceptibility to phishing *increased* among those participants who were forwarded to the training page, we analyzed the answers to our post-experiment questionnaire. One possible explanation that emerges from the questionnaire responses was a false sense of security that is related to the deployed training method: out of the respondents who remembered seeing the training page, 43% selected the option “seeing the training web page made me feel safe”, and 40% selected the option “the company is protecting me from bad emails”. It remains an open question for future work to explore whether this is due to a misinterpretation of the training page (i.e., whether the participants thought they were protected from a real attack), or if this is because of overconfidence in the organization’s IT measures in general, as observed in similar settings in the past [39], [50], [51].

Ultimately, our result shows that organizations need to be careful when using this training method, and aware of possible unintended side effects.

VII. CAN EMPLOYEES HELP THE ORGANIZATION?

We now analyze the data collected from our experiment to answer RQ4, related to crowd-sourced phishing detection in an organization. Such method needs to fulfill the following requirements to be useful:

- **Sustainability:** employees need to keep reporting suspicious emails over long period of time.
- **Effectiveness:** employees’ reports need to be sufficiently accurate and timely so that the organization can stop new campaigns quickly enough.




Fig. 10: Number of different simulated phishing emails participants performed the dangerous action on, by administered contextual training. Missing bars denote a 0.

- **Practicality:** the operational workload to process all the reported emails needs to remain acceptable.

A. Reporting Sustainability for Employees

Recall from Section III that we decided to experiment with two types of feedback: (i) always receive the result of their report; or (ii) receive the result only when (erroneously) reporting a legitimate (non-phishing) email. To investigate reporting sustainability, we examine how the employees’ activity in reporting suspicious emails evolved over time, and whether the tested method of encouraging reporting worked. We examine these questions using the following two hypotheses:

- H7: *Employees keep reporting over time at a steady rate*
- H8: *Providing feedback to reports encourages to report again in the future*

We count all reports and analyze their rate over time, and compare the number of participants that reported more emails after receiving the two different types of feedback.

The results support H7: employees continue reporting emails. Figure 11 shows the number of suspicious emails reported over the duration of the entire experiment.³ We observe a steady income of reports that does not slow down (and even increased when the two new phishing emails were released in August 2020), as shown by the constant fraction of simulated emails reported daily. We further analyze the distribution of frequency of reports that is shown in Figure 12. While 90% of the employees that reported suspicious emails reported 6 or less, there is a non-negligible amount of very active users. We conclude that in our experiment of 15 months, there was no significant “reporting fatigue” suggesting that, if reporting is made easy, employees can actively keep on reporting suspicious emails for long periods of time.

Additionally, we examined whether any demographic influences the quantity of reports by fitting a linear model with Type III sum of squares. Similarly to phishing susceptibility, the combinations of age and computer use in job, and gender and computer use in job are significant (age and computer

³These numbers include all 21,000 employees that received the button.




Fig. 11: Cumulative email reports over time. The dashed red line shows the percentage of simulated emails reported daily.

use $F(10, 14710) = 6.49, p < 0.001$; gender and computer use $F(2, 14710) = 11.35, p < 0.001$). Considering the skewed distribution of computer use, we assume it is the main contributing factor. Indeed, we find that Frequent computer use participants reported a very encouraging 22% of all the simulated emails that they received, while Infrequent use participants reported only 10.20% and Specialized use 7.60%. We conclude that, quite intuitively, employees with the best expected computer skills are also the most active reporters. However, interestingly, Infrequent use participants were more active than the Specialized ones.

The results support H8: positive feedback encourages employees to report more. We find a significant interaction between the type of administered feedback type and the amount of reported emails. To measure this, we first exclude all the participants that never reported any email. Then we count how many emails were reported by the group that actually received positive feedback and by the one that only received feedback about false reports. The former (2,046 participants) is composed by participants in groups that always received feedback and that reported at least one malicious or simulated email (thus receiving the positive feedback). The latter (2,201 participants) is formed by those in groups that did not receive positive feedback, and by those in a group that could receive positive feedback but only reported legitimate emails (thus never receiving the positive feedback). We ran a Welch-corrected ANOVA ($F(1, 3224) = 31.62, p < 0.001$) confirming that participants that saw the positive feedback were more likely to report more emails.

B. Effectiveness of Crowd-Sourced Phishing Detection

To analyze the effectiveness of crowd-sourced phishing detection mechanism as a whole, we analyze *timeliness* and *accuracy* of reports. In addition to sufficiently high reporting activity, organizations need both quick and sufficiently accurate reporting to be able to detect and stop novel phishing campaigns that are often short-lived [14].

We note that since we did not send thousands of copies of the same phishing email at the same time, we cannot directly measure how fast such mass phishing campaigns are reported




Fig. 12: Distribution of the number of reports per user.




Fig. 13: Average per-user-group Cumulative Distribution Function of reported simulated emails as a function of the interval between email reception and report, with standard deviation as a colored region.

and thus detected. Instead, we measured how fast our randomly timed simulated phishing emails were reported by participants. Based on these numbers, we can then estimate how quickly and accurately real mass campaigns could be detected.

Timeliness. We show in Figure 13 the percentage of reports of our simulated emails that arrived shortly after their delivery. We can observe that the reaction time of the employee base as a whole is fast: on average around 10% of the reports arrived within 5 minutes; 20% within 15; and 30% to 40% within 30 minutes. We observe no significant difference between the reporting times of different simulated email campaigns: despite all having different number of total reports (from 2,538 reports of the most reported simulation, down to 832 reports for the least reported), all consistently see a similar amount of reports incoming within the first 30 minutes.

To apply these numbers to a hypothetical company of 1,000 employees where 100 of them are targeted by a phishing campaign, we would have between 8 and 25 reports of the email by employees—of which one within 5 minutes with high probability, and a larger number within 30 minutes.

Accuracy. The average accuracy of reports was good: 68%, up to 79% if spam emails should be reported as well.⁴ We observe

⁴As ground truth, we consider here the outcome of the secondary antiphishing appliance, corrected and validated by the IT department of the company.

that the distribution of employees’ accuracy in reporting is wide: while over 60% of the reporting employees have an accuracy of 80% or more, there is a non-trivial fraction that was always wrong (13% if spam should be reported; 22% otherwise)—however, it mostly comprises employees who reported only a single email. The accuracy of the top 10% of very active employees that reported 6 or more emails (recall Section VII-A) is around 5% higher than considering all employees. We further note that very high reporting accuracy is not crucial. If using a secondary anti-phishing appliance to triage reports, as done in our experiment, employees can be encouraged to be overly-cautious and report emails when in doubt (not only when absolutely sure), as the appliance can serve as a first check on the email, and keep the operational workload acceptable, as discussed below.

Incident awareness. Additionally, we analyzed the employees’ awareness of being the victims of a security incident. We start by noticing that 6% of the participants who performed the dangerous action on our simulated emails immediately reported the email, thus realizing they were victims of a phishing attack.⁵ Only 3.7% of these participants did not toggle the checkbox of the report button that allowed employees to report whether they visited the link contained in the email, or opened its attachments. Interestingly, we observe that some participants were overly-cautious: 13% of the reports of simulated phishing emails stated that they opened the link or attachment, despite not having done it.

C. Practicality for the Organization

We observe that a secondary appliance triaging the reports makes the added workload reasonable: out of the 7,191 non-simulated reported emails in 15 months, only 689 (9%) of the decisions were taken by human administrators, and actually overturned the decision taken by the appliance only 50 times (7% of the total handled cases). The main goal of this secondary appliance is to filter out the reports of clear benign emails or minor threats such as spam, which include the majority of our collected reports: out of 7,191 reports of emails not part of our exercise, 3,531 were benign, and 2,371 were spam or unwanted newsletters. Thus, only roughly 1.5 emails per day needed manual handling from the IT department—a clearly acceptable workload for a large organization that was collecting reports from over 21,000 users.

D. Finding Real Phishing Campaigns

We further validate our crowd-sourced phishing detection approach by analyzing whether we caught any *real* phishing campaigns delivered to employees of the company (in addition to our simulated phishing emails). We use the verdicts of the secondary filter and manual inspection by IT specialists to find reported phishing and other malicious emails. We observed 918 reports of real phishing emails during the last 5

⁵We only measure participants that did not receive training after falling for a simulated phish, because they had to understand by themselves what happened—the training material stated clearly that it was a simulated phishing attack.

months of our deployment. With email similarity techniques, we measured how many emails similar to the reported ones were incoming and found 252 large-scale phishing campaigns comprising 28,830 emails, and 1,534 emails with malware attached that our crowd-sourced approach would have detected in a short time span from their beginning.

VIII. STUDY VALIDITY

Simulated emails limitations. Recall that 3 of our 8 emails had a malicious attachment where the dangerous action was to enable macros. While the company could monitor when macros were enabled, with a network call to the monitoring infrastructure, it could not know when participants only *clicked*, i.e., simply opened and closed the attachment without enabling macros. Thus, for attachments we underestimate the number of clicks by setting it to the number of dangerous actions.

The company did not record when a simulated email was opened, thus we do not know the conversion rate from opening the email to clicking and dangerous actions.

Due to data protection concerns, it was not recorded whether employees submitted their valid credentials to the simulated phishing websites. Therefore, the amount of employees that we recorded performing some of the dangerous actions (e.g., submitting credentials) could be overestimated, because we cannot filter out employees who submitted bogus credentials.

Email warnings limitations. Our partner company added warnings on top of simulated phishing emails, but not on top of emails that the inline filtering solution in use deemed suspicious but let through anyway. This could lead some participants to fall once for our simulated phishing email, and subsequently associate the presence of warnings to surely suspicious emails, or even worse, to training exercises. Further studies on this promising type of warnings when added on top of legitimate but suspicious-looking emails are needed, to remove the potential bias.

Campaigns success rates. As shown in Figure 14, the different simulated phishing email campaigns had different success rates. Such differences do not influence our analysis of RQ1-RQ4 due to two reasons. First, we always count the total number of clicks or dangerous actions for all campaigns and compare total counts. Second, the order of administered campaigns was randomized for each participant in large groups.

Applicability to different companies. Our partner company operates in numerous different sectors, has a diverse workforce, and large size. Thus, we believe that our results can generalize to various similar-sized (large) companies. It is unclear whether our results generalize to companies with very specialized IT workers, e.g., software engineering companies, or to very small organizations.

IX. RELATED WORK

Phishing and demographics. Age is one of the most analyzed factors of phishing, as it intuitively often correlates with technological skills. Studies showed that very young [40],

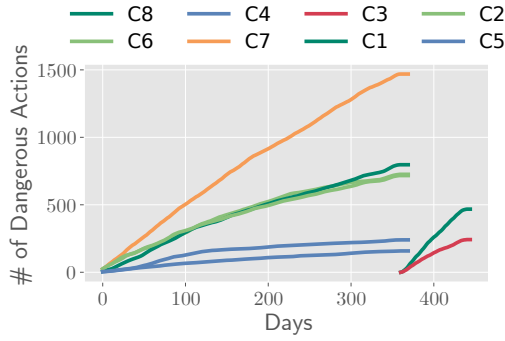


Fig. 14: Cumulative count of dangerous actions per campaign during the experiment duration.

[41], [34], and older people [36], [35], [29] are more at-risk for phishing. Preliminary studies show that aging increases susceptibility to phishing [35], but only the two extremes (very young and senior persons) were tested, not the full age spectrum. Further, different age ranges are susceptible to different types of phishing emails [15], [36]. Gender is a more divisive demographic, according to a recent literature survey [11], but the studies that do find an impact show that women are more vulnerable [36], and can detect less phishing attempts [30]. Experience with computers [11], experience of previous phishing attempts [29], and seniority at an organization [26] also positively influence phishing immunity.

Phishing at workplace. Some previous studies show that within the organization’s boundaries, employees feel safer and generally trust their company’s measures, thus lowering their attention [39], [50], and the existence of “repeated clickers” that are extremely at risk of being phished [42]. Other studies find that helping employees in phishing prevention is made hard by the fact they struggle to comply with corporate security policies and often ignore them [52], [53].

Phishing training and warnings. There is broad consensus that training should be active, e.g., with security games [54]. A popular mechanism is running simulated phishing exercises, an approach adopted by several companies [19], [18], [17], [20] following promising research results [32], [33], [34], where (possibly unaware [47], [48]) employees receive simulated phishing emails over time, ideally at random intervals [16]. This practice is often combined with embedded training: immediately redirecting the employee that fell for phishing to a dedicated web page explaining the simulated attack they just fell for and providing information about phishing [34].

Previous studies suggest that training should be continuous, as knowledge retention spans from a few days [33], [45] to a few months at most [34]. However, research effort has unclear external validity, because most work employed small populations [33], [31], [28], [16], [32], [27], was shorter in time [16], [33], [27], had populations with little diversity [31], [28], [27] or tested a role-playing setting only [32], [31]—and a recent study questioned whether these results transfer to a

corporate setting [55]. Further, the business ecosystem that emerged around embedded phishing training [19], [18], [17], [20] claims improvements due to the benefits of training in a recent collaborative report [2], but does not report results of experiment in controlled settings [56].

Phishing warnings have been studied extensively in the context of browsers (e.g., [57], [58]). Some recent works [37] also evaluate different kinds of warnings shown on the email client. While too frequent warnings can susceptible to habituation and lose some of their effectiveness over time [59], the literature agrees that carefully-timed warnings are in general effective.

Crowd-sourced phishing detection. Several companies already provide tools to report phishing emails, to quickly detect new attacks using aggregate information across multiple customers [19], [24]. The same companies report that users are improving at reporting phishing attempts over time [2], [1], however, other work showed that users are reticent to report phishing to the IT because of the lack of transparency in the process [60] and lack of fast responses from the system [39]. Prior to our work, it was not known if employees as a crowd-sourcing mechanism in a closed scenario, such as a corporation that manages reported phishing in-house, works effectively with acceptable operational workload. Few recent works also suggest this concept, but do not evaluate it [25], [61].

X. CONCLUSIONS AND FUTURE WORK

Thanks to our long-term and large-scale experiment, in this paper we supported several prior findings such as effectiveness of warnings with increased ecological validity. Further, we found that embedded phishing training, as commonly deployed in the industry today, is not effective and can in fact have negative side effects. In this regard, our results contradict prior literature and common industry practices. Finally, we are the first to experimentally demonstrate that crowd-sourced phishing detection is effective and practical in a single organization.

Based on these results, we encourage organizations to adopt phishing prevention tools like warnings that have been extensively studied and where the available literature supports their effectiveness overwhelmingly. We call for caution in the deployment of methods like embedded phishing exercises and training, where the existing literature is less unanimous about their effectiveness, and our research discovers potential negative side effects. We recommend organizations to consider crowd-sourced phishing detection as a new and complementary way to improve the overall phishing prevention capabilities of the organization, since its effectiveness looks promising and operational workload remains low.

Our work also identifies topics where more research is needed. Our research shows that the effectiveness of phishing exercise and training has not been sufficiently measured, and it remains unknown what is the most effective way to deliver embedded phishing training. More research is also needed to better understand the (psychological) effects of phishing exercises and training that is embedded into the normal working context of employees, and how such effects may influence the employees’ future handling of real phishing emails.

ACKNOWLEDGMENTS

This research has been partially supported by the Zurich Information Security and Privacy Center (ZISC).

REFERENCES

- [1] Verizon, “2012 Data Breach Investigations Report,” <https://enterprise.verizon.com/resources/reports/2020-data-breach-investigations-report.pdf>, 2020, [Online; accessed 20 March 2021].
- [2] —, “2019 Data Breach Investigations Report,” <https://enterprise.verizon.com/resources/reports/2019/2019-data-breach-investigations-report.pdf>, 2019, [Online; accessed 20 March 2021].
- [3] J. Hong, “The state of phishing attacks,” *Communications of the ACM*, vol. 55, no. 1, pp. 74–81, 2012.
- [4] A. Oest, Y. Safai, A. Doupe, G.-J. Ahn, B. Wardman, and G. Warner, “Inside a phisher’s mind: Understanding the anti-phishing ecosystem through phishing kit analysis,” in *2018 APWG Symposium on Electronic Crime Research (eCrime)*. IEEE, 2018, pp. 1–12.
- [5] M. Cova, C. Kruegel, and G. Vigna, “There is no free phish: An analysis of ‘free’ and live phishing kits,” *WOOT*, vol. 8, pp. 1–8, 2008.
- [6] X. Han, N. Kheir, and D. Balzarotti, “Phisheye: Live monitoring of sandboxed phishing kits,” in *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, 2016, pp. 1402–1413.
- [7] M. Bitaab, H. Cho, A. Oest, P. Zhang, Z. Sun, R. Pourmohamad, D. Kim, T. Bao, R. Wang, Y. Shoshitaishvili, A. Doupe, and G.-J. Ahn, “Scam pandemic: How attackers exploit public fear through phishing,” in *eCrime Symposium on Electronic Crime Research*, 2020.
- [8] M. Khonji, Y. Iraqi, and A. Jones, “Phishing detection: a literature survey,” *IEEE Communications Surveys & Tutorials*, vol. 15, no. 4, pp. 2091–2121, 2013.
- [9] A. Aleroud and L. Zhou, “Phishing environments, techniques, and countermeasures: A survey,” *Computers & Security*, vol. 68, pp. 160–196, 2017.
- [10] A. Almomani, B. B. Gupta, S. Atawneh, A. Meulenberg, and E. Almomani, “A survey of phishing email filtering techniques,” *IEEE communications surveys & tutorials*, vol. 15, no. 4, pp. 2070–2090, 2013.
- [11] D. Jampen, G. Gür, T. Sutter, and B. Tellenbach, “Don’t click: towards an effective anti-phishing training, a comparative literature review,” *Human-centric Computing and Information Sciences*, vol. 10, no. 1, pp. 1–41, 2020.
- [12] I. Fette, N. Sadeh, and A. Tomasic, “Learning to detect phishing emails,” in *Proceedings of the 16th international conference on World Wide Web*, 2007, pp. 649–656.
- [13] G. Xiang, J. Hong, C. P. Rose, and L. Cranor, “Cantina+ a feature-rich machine learning framework for detecting phishing web sites,” *ACM Transactions on Information and System Security (TISSEC)*, vol. 14, no. 2, pp. 1–28, 2011.
- [14] A. Oest, P. Zhang, B. Wardman, E. Nunes, J. Burgis, A. Zand, K. Thomas, A. Doupe, and G.-J. Ahn, “Sunrise to sunset: Analyzing the end-to-end life cycle and effectiveness of phishing attacks at scale,” in *29th {USENIX} Security Symposium ({USENIX} Security 20)*, 2020, pp. 361–377.
- [15] J. Wang, T. Herath, R. Chen, A. Vishwanath, and H. R. Rao, “Research article phishing susceptibility: An investigation into the processing of a targeted spear phishing email,” *IEEE transactions on professional communication*, vol. 55, no. 4, pp. 345–362, 2012.
- [16] R. Wash and M. M. Cooper, “Who provides phishing training? facts, stories, and people like me,” in *Proceedings of the 2018 chi conference on human factors in computing systems*, 2018, pp. 1–12.
- [17] Proofpoint, “Proofpoint Security Awareness Training,” <https://www.proofpoint.com/us/products/security-awareness-training>, 2021, [Online; accessed 20 March 2021].
- [18] Cofense, “Cofense Phishing Solutions and Products,” <https://cofense.com/product-overview/>, 2021, [Online; accessed 20 March 2021].
- [19] Rapid7, “Phishing Awareness Training,” <https://www.rapid7.com/solutions/phishing-awareness-training/>, 2021, [Online; accessed 20 March 2021].
- [20] KnowBe4, “Phishing,” <https://www.knowbe4.com/phishing>, 2021, [Online; accessed 20 March 2021].
- [21] PhishTank, “Join the fight against phishing,” <http://phishtank.org/>, 2021, [Online; accessed 20 March 2021].
- [22] OpenPhish, “Phishing Intelligence,” <https://openphish.com/>, 2021, [Online; accessed 20 March 2021].
- [23] T. Moore and R. Clayton, “Evaluating the wisdom of crowds in assessing phishing websites,” in *International Conference on Financial Cryptography and Data Security*. Springer, 2008, pp. 16–30.
- [24] Proofpoint, “Introducing PhishAlarm, Wombat’s One-Click Email Reporting Button,” <https://www.proofpoint.com/us/security-awareness/post/introducing-phishalarm-wombats-one-click-email-reporting-button>, 2015, [Online; accessed 20 March 2021].
- [25] P. Burda, L. Allodi, and N. Zannone, “Don’t forget the human: a crowdsourced approach to automate response and containment against spear phishing attacks,” in *2020 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*. IEEE, 2020, pp. 471–476.
- [26] P. Burda, T. Chotza, L. Allodi, and N. Zannone, “Testing the effectiveness of tailored phishing techniques in industry and academia: a field experiment,” in *Proceedings of the 15th International Conference on Availability, Reliability and Security*, 2020, pp. 1–10.
- [27] A. Burns, M. E. Johnson, and D. D. Caputo, “Spear phishing in a barrel: Insights from a targeted phishing campaign,” *Journal of Organizational Computing and Electronic Commerce*, vol. 29, no. 1, pp. 24–39, 2019.
- [28] A. Carella, M. Kotsoev, and T. M. Truta, “Impact of security awareness training on phishing click-through rates,” in *2017 IEEE international conference on Big Data (Big Data)*. IEEE, 2017, pp. 4458–4466.
- [29] B. E. Gavett, R. Zhao, S. E. John, C. A. Bussell, J. R. Roberts, and C. Yue, “Phishing suspiciousness in older and younger adults: The role of executive functioning,” *PLoS one*, vol. 12, no. 2, p. e0171620, 2017.
- [30] C. Iuga, J. R. Nurse, and A. Erola, “Baiting the hook: factors impacting susceptibility to phishing attacks,” *Human-centric Computing and Information Sciences*, vol. 6, no. 1, pp. 1–20, 2016.
- [31] P. Kumaraguru, Y. Rhee, S. Sheng, S. Hasan, A. Acquisti, L. F. Cranor, and J. Hong, “Getting users to pay attention to anti-phishing education: evaluation of retention and transfer,” in *Proceedings of the anti-phishing working groups 2nd annual eCrime researchers summit*, 2007, pp. 70–81.
- [32] P. Kumaraguru, Y. Rhee, A. Acquisti, L. F. Cranor, J. Hong, and E. Nunge, “Protecting people from phishing: the design and evaluation of an embedded training email system,” in *Proceedings of the SIGCHI conference on Human factors in computing systems*, 2007, pp. 905–914.
- [33] P. Kumaraguru, S. Sheng, A. Acquisti, L. F. Cranor, and J. Hong, “Lessons from a real world evaluation of anti-phishing training,” in *2008 eCrime Researchers Summit*. IEEE, 2008, pp. 1–12.
- [34] P. Kumaraguru, J. Cranshaw, A. Acquisti, L. Cranor, J. Hong, M. A. Blair, and T. Pham, “School of phish: a real-world evaluation of anti-phishing training,” in *Proceedings of the 5th Symposium on Usable Privacy and Security*, 2009, pp. 1–12.
- [35] T. Lin, D. E. Capecci, D. M. Ellis, H. A. Rocha, S. Dommaraju, D. S. Oliveira, and N. C. Ebner, “Susceptibility to spear-phishing emails: Effects of internet user demographics and email content,” *ACM Transactions on Computer-Human Interaction (TOCHI)*, vol. 26, no. 5, pp. 1–28, 2019.
- [36] D. Oliveira, H. Rocha, H. Yang, D. Ellis, S. Dommaraju, M. Muradoglu, D. Weir, A. Soliman, T. Lin, and N. Ebner, “Dissecting spear phishing emails for older vs young adults: On the interplay of weapons of influence and life domains in predicting susceptibility to phishing,” in *Proceedings of the 2017 chi conference on human factors in computing systems*, 2017, pp. 6412–6424.
- [37] J. Petelka, Y. Zou, and F. Schaub, “Put your warning where your link is: Improving and evaluating email phishing warnings,” in *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, 2019, pp. 1–15.
- [38] M. Volkamer, K. Renaud, B. Reinheimer, and A. Kunz, “User experiences of torpedo: Tooltip-powered phishing email detection,” *Computers & Security*, vol. 71, pp. 100–113, 2017.
- [39] E. J. Williams, J. Hinds, and A. N. Joinson, “Exploring susceptibility to phishing in the workplace,” *International Journal of Human-Computer Studies*, vol. 120, pp. 1–13, 2018.
- [40] S. Sheng, M. Holbrook, P. Kumaraguru, L. F. Cranor, and J. Downs, “Who falls for phish? a demographic analysis of phishing susceptibility and effectiveness of interventions,” in *Proceedings of the SIGCHI conference on human factors in computing systems*, 2010, pp. 373–382.
- [41] M. Blythe, H. Petrie, and J. A. Clark, “F for fake: four studies on how we fall for phish,” in *Proceedings of the SIGCHI conference on human factors in computing systems*, 2011, pp. 3469–3478.

- [42] M. Canham, M. Constantino, I. Hudson, S. M. Fiore, B. Caulkins, and L. Reinerman-Jones, "The enduring mystery of the repeat clickers," in *Fifteenth Symposium on Usable Privacy and Security (SOUPS 2019)*. USENIX Advanced Computing Systems Association, 2019.
- [43] G. Workspace, "Advanced phishing and malware protection," <https://support.google.com/a/answer/9157861>, 2021, [Online; accessed 20 March 2021].
- [44] Gmail, "Avoid and report phishing emails," <https://support.google.com/mail/answer/8253>, 2021, [Online; accessed 20 March 2021].
- [45] P. Kumaraguru, S. Sheng, A. Acquisti, L. F. Cranor, and J. Hong, "Teaching johnny not to fall for phish," *ACM Transactions on Internet Technology (TOIT)*, vol. 10, no. 2, pp. 1–31, 2010.
- [46] S. Mansfield-Devine, "The imitation game: how business email compromise scams are robbing organisations," *Computer Fraud & Security*, vol. 2016, no. 11, pp. 5–10, 2016.
- [47] K. Parsons, A. McCormac, M. Pattinson, M. Butavicius, and C. Jerram, "The design of phishing studies: Challenges for researchers," *Computers & Security*, vol. 52, pp. 194–206, 2015.
- [48] R. C. Dodge Jr, C. Carver, and A. J. Ferguson, "Phishing for user security awareness," *computers & security*, vol. 26, no. 1, pp. 73–80, 2007.
- [49] P. Finn and M. Jakobsson, "Designing ethical phishing experiments," *IEEE Technology and Society Magazine*, vol. 26, no. 1, pp. 46–58, 2007.
- [50] K. K. Greene, M. Steves, M. Theofanos, and J. Kostick, "User context: an explanatory variable in phishing susceptibility," in *Proc. 2018 Workshop Usable Security*, 2018.
- [51] D. Conway, R. Taib, M. Harris, K. Yu, S. Berkovsky, and F. Chen, "A qualitative investigation of bank employee experiences of information security and phishing," in *Thirteenth Symposium on Usable Privacy and Security ({SOUPS} 2017)*, 2017, pp. 115–129.
- [52] H.-y. S. Tsai, M. Jiang, S. Alhabash, R. LaRose, N. J. Rifon, and S. R. Cotten, "Understanding online safety behaviors: A protection motivation theory perspective," *Computers & Security*, vol. 59, pp. 138–150, 2016.
- [53] M. Siponen, M. A. Mahmood, and S. Pahlila, "Employees' adherence to information security policies: An exploratory field study," *Information & management*, vol. 51, no. 2, pp. 217–224, 2014.
- [54] S. Sheng, B. Magnien, P. Kumaraguru, A. Acquisti, L. F. Cranor, J. Hong, and E. Nunge, "Anti-phishing phil: the design and evaluation of a game that teaches people not to fall for phish," in *Proceedings of the 3rd symposium on Usable privacy and security*, 2007, pp. 88–99.
- [55] D. D. Caputo, S. L. Pfleeger, J. D. Freeman, and M. E. Johnson, "Going spear phishing: Exploring embedded training and awareness," *IEEE Security & Privacy*, vol. 12, no. 1, pp. 28–38, 2013.
- [56] H. Siadati, S. Palka, A. Siegel, and D. McCoy, "Measuring the effectiveness of embedded phishing exercises," in *10th {USENIX} Workshop on Cyber Security Experimentation and Test ({CSET} 17)*, 2017.
- [57] S. Egelman, L. F. Cranor, and J. Hong, "You've been warned: an empirical study of the effectiveness of web browser phishing warnings," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 2008, pp. 1065–1074.
- [58] D. Akhawe and A. P. Felt, "Alice in warningland: A large-scale field study of browser security warning effectiveness," in *22nd {USENIX} Security Symposium ({USENIX} Security 13)*, 2013, pp. 257–272.
- [59] A. Vance, B. Kirwan, D. Bjornn, J. Jenkins, and B. B. Anderson, "What do we really know about how habituation to warnings occurs over time? a longitudinal fmri study of habituation and polymorphic warnings," in *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, 2017, pp. 2215–2227.
- [60] Y. Kwak, S. Lee, A. Damiano, and A. Vishwanath, "Why do users not report spear phishing emails?" *Telematics and Informatics*, vol. 48, p. 101343, 2020.
- [61] C. Nguyen, M. L. Jensen, A. Durcikova, and R. T. Wright, "A comparison of features in a crowdsourced phishing warning system," *Information Systems Journal*, 2021.

APPENDIX

We report here some of the materials we used in our experiment: the simulated phishing emails, the embedded training webpage that participants viewed when performing the dangerous action on the emails, and the questionnaire we administered at the end of the study.

Due to space limitations, we only report a sample of each material: one embedded training webpage tailored for a specific simulated email; four simulated phishing emails, and the questionnaire questions that gave us some insights that we reported in this paper. We believe this sample is sufficient to understand our design in full details.

A. Embedded Training Webpage

We show in Figure 15 the contextual training webpage displayed when employees performed the dangerous action of one simulated phishing email. It contains tailored information, explanation about the awareness campaign, and the tabs further contain information about email threats and an instructional video.

B. Simulated Phishing Emails

Figure 16 shows four of the simulated phishing emails we sent to participants. The reported emails either aim to get the participants to click a link to a malicious webpage, which in turn aims to get the participants to do the unsafe action (e.g., submit their credentials), or aim to get the participant to download an attachment, e.g., a document that prompts to enable macros. Each email used different triggers to urge participants to click, such as curiosity or scare of consequences.

C. Questionnaire

The questionnaire consisted of 27 closed-ended questions, such as yes/no questions, and multiple choice questions where participants could select more than one answer. Every question offered an answer such as not knowing or not remembering, and could also be left unanswered by the respondent. We can group the questions in five main categories:

- Familiarity with computer and email security threats.
- Email warnings.
- The button to report phishing in the email client.
- Remembering suspicious emails and security incidents.
- The contextual training webpage.

The groups of questions about email warnings, the report button, and the contextual training webpage were preceded by a recall about the tool, e.g., we displayed a screenshot of the training webpage immediately before its questions.

The questions about the deployed tools (email warnings, report button, contextual training webpage) were preceded by recalling it to respondents, e.g., we displayed a screenshot of the training webpage immediately before asking questions about it. They asked the respondent if they remembered noticing the tool in the past 12 months and using it, and what they thought about the tool.

We report as a sample the questions about the training page in the following.


- Q22: Do you remember seeing this training page during the past 12 months?
 - Yes; No; I'm not sure
- Q23: How did you feel when you saw the training page?
 - Embarrassed. I understood that I had made a mistake.

- *Concerned. I realized I had endangered my own and my company's online security.*
- *Safe. I felt that the organization is protecting me online.*
- *Uninterested. Seeing the training page did not trigger any emotional reaction.*
- *I don't remember.*
- Q24: How much time would you estimate you spent on the training page?
 - *More than a minute. I read the whole page carefully.*
 - *Less than a minute. I briefly skimmed the provided information.*
 - *Few seconds. I opened the page but did not read its contents.*
 - *I don't remember.*
- Q25: Did you find the content of the training page trustworthy?
 - *Yes. I thought it was from a legitimate source like the IT department of the organization.*
 - *No. The training page looked suspicious to me (perhaps a scam).*
 - *I don't remember.*
- Q26: Did you find the content of the training page useful?
 - *Yes. I found it a good reminder of threats of malicious emails.*
 - *No. The provided information was not helpful to me.*
 - *Not sure.*
- Q27: After visiting the training page, did your attitude towards suspicious emails change?
 - *Yes. I learned more about how to check suspicious emails.*
 - *Yes. I realized that suspicious emails can be part of a corporate training campaign.*
 - *Yes. I felt that the organization is protecting me from bad emails.*
 - *No. I already knew the information in the webpage.*
 - *No. The content of the page was not clear or informative.*
 - *Don't remember.*

Company Logo

Privacy PolicyContact

En



Phishing

Identify dangerous e-mails quickly and reliably

You've just opened an Excel file named "Management levels 2019" and enabled the macros included in the document by clicking "Enable editing" and "Activate content" in the status bar. When you enabled editing, your computer could have been infected with malware (malicious software) in the worst-case scenario.

The e-mail appeared to originate from within the company and prompted you to open an attachment claiming to be a new management list following internal reorganization. You were prompted to activate the macros in the attached document. You did as you were asked without noticing the alarm signals.

You could have seen through this particular "management list" attack. To begin with, the sender looked suspicious and was indeed fake (██████████). The fact that such sensitive information was sent by e-mail, without encryption and using an incorrect address should have led you to doubt its authenticity. When you were then asked to activate the macros, you should definitely have become suspicious.

Do not open any e-mail attachments if you are not sure what they contain!

Note: This attack was part of a ██████████ awareness campaign. No data was transferred and no malware was installed.

Tips

Information

Video

Preventing the attack

Never enable active content, such as the Word macros in this case, if you are not sure that the document comes from a reliable source. In this attack, just opening the Word file would not have caused any damage.

Be wary of all e-mails that arrive in your inbox:

- How does the sender know your address and why are you receiving the e-mail?
- Do you know the sender, does the content make sense and does the language sound like the sender?
- Are you being pressured into doing something or is something being offered to you in a pushy manner?

Strengthen your knowledge with our phishing exercise.

Start the exercise now

Fig. 15: Sample contextual training webpage displayed to employees after dangerous actions on a simulated phishing email.

Password change



Dear employee

Your password expires in two days!

Last password change: 6th January 2019


Please click [here](#) to change your password. You could lose access to important systems if you do not change your password.

This is an automated email.

Best regards
User Help Desk
[redacted]

(a) Email prompting to change the organization's password.

Assignments to management levels



Download

To all department heads

As announced in the newsletter of April 5, 2021, as part of the reorganization, the management levels must be adjusted. Since some employees will be downgraded, this will probably lead to considerable discussions.

Attached is a list of all changes. You can filter by department, etc., simply activate the macros.

Many greetings
[redacted]

(b) Email with attachment with malicious macros.

Private data found - Your action is required!



Private Data

We found a lot of private data on user drives on the [redacted] infrastructure. This consumes too much disk space on servers and backup systems get overloaded.

IMPORTANT: Every file identified as private data will be deleted on April 19, 2021 at 3pm.

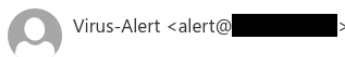
You may choose to view your private data and save it to a local drive if you wish. Please log on with your personal account to do so:

[http://\[redacted\]](http://[redacted])

Best regards
Service Desk

(c) Email prompting to check one's files in a corporate drive.

IMPORTANT: Virus found



A virus was found on your computer:

K8Stba-trojan.vbs

The virus can not be deleted automatically due to insufficient user rights. Please start the scan manually to remove the virus.

[http://\[redacted\]](http://[redacted])

Best regards
Virusscan [redacted]

(d) Email alerting of presumed malware.

Fig. 16: Sample of the simulated phishing emails. We report three emails containing a link to a malicious webpage (e.g., that asked for credentials, or prompted a download), and one with a malicious attachment.