“What was that site doing with my Facebook password?” Designing Password-Reuse Notifications

Maximilian Golla
Ruhr University Bochum
maximilian.golla@rub.de

Lydia Filipe
University of Chicago
lydiafilipe@uchicago.edu

Miranda Wei
University of Chicago
weim@uchicago.edu

Markus Dürmuth
Ruhr University Bochum
markus.duermuth@rub.de

Juliette Hainline
University of Chicago
juliettehainline@uchicago.edu

Elissa Redmiles
University of Maryland
eredmiles@cs.umd.edu

Blase Ur
University of Chicago
blase@uchicago.edu

ABSTRACT

Password reuse is widespread, so a breach of one provider's password database threatens accounts on other providers. When companies find stolen credentials on the black market and notice potential password reuse, they may require a password reset and send affected users a notification. Through two user studies, we provide insight into such notifications. In Study 1, 180 respondents saw one of six representative notifications used by companies in situations potentially involving password reuse. Respondents answered questions about their reactions and understanding of the situation. Notifications differed in the concern they elicited and intended actions they inspired. Concerningly, less than a third of respondents reported intentions to change any passwords. In Study 2, 588 respondents saw one of 15 variations on a model notification synthesizing results from Study 1. While the variations' impact differed in small ways, respondents' intended actions across all notifications would leave them vulnerable to future password-reuse attacks. We discuss best practices for password-reuse notifications and how notifications alone appear insufficient in solving password reuse.

CCS CONCEPTS

• Security and privacy → Usability in security and privacy;

KEYWORDS

Notifications; Password Reuse; Data Breaches; Usable Security

ACM Reference Format:

Maximilian Golla, Miranda Wei, Juliette Hainline, Lydia Filipe, Markus Dürmuth, Elissa Redmiles, and Blase Ur. 2018. “What was that site doing with my Facebook password?” Designing Password-Reuse Notifications. In *CCS '18: 2018 ACM SIGSAC Conference on Computer & Communications Security*, Oct. 15–19, 2018, Toronto, ON, Canada. ACM, New York, NY, USA, 18 pages. <https://doi.org/10.1145/3243734.3243767>

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

CCS '18, October 15–19, 2018, Toronto, ON, Canada

© 2018 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-5693-0/18/10.

<https://doi.org/10.1145/3243734.3243767>

1 INTRODUCTION

People reuse passwords [10, 16, 20, 32, 44, 46, 65]. An average user may have hundreds of different online accounts [16, 44, 65], and passwords are unlikely to be completely replaced anytime soon [4]. As password managers [13, 37] and single sign-on systems [3, 55] have low adoption, password reuse is a common coping strategy.

Password reuse has major ramifications for the security of online accounts. A breach of one account provider's password database puts at risk accounts on other services where login credentials are the same as, or even just similar to [63], the breached accounts. Attackers that target large leaks of passwords stored using computationally expensive hash functions (e.g., *bcrypt*) exploit this password reuse in offline guessing [22]. Attackers try to match identifiers like usernames and email addresses to previously cracked credentials. They then transform the already known passwords to increase their likelihood of correctly guessing passwords [9, 26, 62].

Unfortunately, password breaches are common. The website *haveibeenpwned.com* counts billions of compromised account credentials due to data breaches, including from high-profile services like Yahoo!, LinkedIn, MySpace, and Dropbox [31, 45]. Thomas et al. estimated that 7–25 % of passwords traded on black-market forums match high-value targets like Google accounts [57].

Account providers send a variety of notifications about situations potentially caused by password reuse. We refer to all such notifications as *password-reuse notifications*, regardless of whether password reuse is explicitly mentioned. To protect their users, some providers proactively monitor black-market sources for passwords stolen from other sites, searching for matches in their own password database [39]. Once aware of such situations, these providers send notifications to affected users, encouraging them to change their password. Password-reuse notifications also include notifications about suspicious login attempts, which may have been triggered by a password-reuse attack, or notifications requiring a password reset after a data breach. In a recent example, Twitter asked users to change not only their Twitter passwords, but also passwords on services where they had reused their Twitter password [1].

Surprisingly little is known about how users interpret or respond to password-reuse notifications, and how the design of such notifications impacts users' understanding and risk perception.

Current password-reuse notifications vary widely, and despite the frequency with which such notifications are sent, no best practices have been outlined. This paucity of knowledge contrasts with the large and rich literature investigating the design of warnings and notifications about other security-critical tasks, including detecting phishing [11, 53], TLS-protected browsing [2, 15], malware [6, 7], and two-factor authentication (2FA) [48]. Many studies have aimed to help users make better passwords [12, 40, 59] or measured the prevalence of password reuse [10, 32, 46, 57]. This paper is the first to explore how to inform users about situations caused by password reuse and help them recover from the resultant consequences.

Password-reuse notifications face the herculean task of helping users understand and respond to a convoluted situation. Users have posted on Twitter about their confusion about receiving such notifications. For example, one tweet about a notification asked, “What was another site doing with my Facebook password in the first place?” This may be because understanding the risks of password reuse requires knowledge of how attackers leverage password breaches to compromise accounts on other services. Password-reuse notifications must address this underlying complexity to convince users to replace reused passwords across all sites with a new, unique password for each account. We explain the complexity of these issues from the perspective of a fictitious company, AcmeCo, which we adopt for the remainder of this paper.

We conducted two complementary user studies about password-reuse notifications. First, we sought to understand how users understand and perceive existing notifications. We collected 24 notifications sent by real companies in situations that may have been caused by password reuse. We chose six notifications whose characteristics were representative of the full 24. In Study 1, we conducted a scenario-based online survey in which 180 Mechanical Turk workers saw one of these six notifications (Sec. 3). We asked respondents why they might have received such a notification, the feelings the notification elicits, and what actions they might take in response.

Respondents reported they would be alarmed and confused, and that they would intend to take action in response to receiving these notifications (Sec. 4). Some notifications were more effective than others at encouraging a response. Ultimately, though, participants’ responses misattributed the potential root cause of receiving these (real, previously deployed) notifications. Only 20.6 % mentioned the breach of another company’s password database as a potential cause, and only 18.8 % mentioned password reuse as a factor.

Based on respondents’ perceptions and responses (Sec. 5), we identified five design goals for password-reuse notifications that integrated characteristics of notifications that were effective in Study 1 and improved upon characteristics that were less effective. We then conducted a follow-up study to analyze a model notification we believed achieved all five design goals (Sec. 6). This notification explicitly describes password reuse and the breach of another provider as the cause of the notification. Additionally, it forces a password reset, encourages other beneficial security actions, and is delivered through multiple mediums. Study 2 was again a scenario-based survey in which 588 Mechanical Turk workers saw one of 15 variants of this model notification.

While Study 2 respondents perceived our model notification as official and urgent, they nonetheless misattributed the root cause

of the notification (Sec. 7). Many respondents did not perceive password reuse as a potential cause of the situation. Additionally, although nearly all respondents stated intentions to change one or more passwords, most reported plans to create these “new” passwords by reusing other passwords of theirs, leaving them vulnerable to similar attacks in the future. From our collected results, we establish best practices for maximizing the effectiveness of password-reuse notifications. However, because password-reuse notifications may not be sufficient on their own, we conclude with a discussion of additional steps for holistically addressing password reuse (Sec. 9).

2 BACKGROUND

We summarize prior work on password reuse and warning design.

2.1 Passwords and Password Reuse

Passwords are the dominant method of user authentication for online accounts due to their low cost, immediacy, convenience, and deployability [4, 27]. Although online account providers employ methods beyond passwords to improve security, such as 2FA [8, 24] and risk-based authentication [19, 23, 41], solutions such as password managers face steep adoption barriers [13]. Accounts therefore remain vulnerable to a number of password-related attacks [61].

Password reuse amplifies the severity of all password attacks. Once login credentials are compromised, all accounts with those same credentials become vulnerable. Various studies over the years have found that users reuse a majority of their passwords across sites [10]. Users have dozens of accounts, but only a few passwords that they cycle through [14, 16, 44, 54, 65]. Users reuse passwords to minimize the burden of memorization [17, 20], and they do so especially often for accounts they consider lower value [20, 54]. Even if users do not reuse passwords verbatim, they often modify existing passwords when creating new ones [34, 44, 52].

We consider password breaches to be cases where a hacker illegally obtains login credentials from a vulnerable system [30]. Once an account is breached, any other accounts sharing the same credentials become vulnerable [26]. Breaches are frequent, with over 4.5 billion credentials reported stolen in 2016 [30]. Leveraging stolen credentials enables attackers to perform online guessing with some success. Thomas et al. accumulated over 1.79 billion non-unique usernames and passwords from credential leaks, finding that 7–25 % of those credentials would enable attackers to log into a compromised account holder’s Google account [57]. Credential-stuffing, which automates logging into as many sites as possible with the stolen login credentials, generates more than 90 % of login traffic on many of the world’s largest websites and mobile apps [51]. Once accounts have been compromised, attackers may use them for spam, financial data, or distributing malware [43, 56].

2.2 Security Warnings and Notifications

A large body of prior work has researched security warnings and notifications broadly. Some examples include encouraging users to adopt 2FA [48] and detecting phishing [11, 53]. In their study of 25 million Google Chrome and Firefox users, Akhawe et al. found that user experience has a significant impact on behavior and that users often do look at warnings [2], contrary to other findings that users are susceptible to habituation and often ignore web warnings [5, 6].

Despite extensive prior work measuring password reuse, very few studies have examined password-reuse notifications. Jenkins et al. evaluated the efficacy of just-in-time fear appeals in warnings — “persuasive messages intended to better help someone be aware of a threat and to persuade them to engage in protective action” — at preventing users from reusing passwords, finding that such appeals resulted in a significant decrease in password reuse [33]. This suggests that notifications could encourage better password creation and management strategies. There is a need, however, to isolate what is effective or ineffective about these notifications. While Zou et al. studied reactions to notifications of the Equifax data breach, their work did not examine password reuse [66]. Huh et al. studied the notification LinkedIn sent users after their password database was breached, finding that less than half of participants changed their LinkedIn password upon receiving this notification [29]. While they asked respondents to self-report their actual reactions to receiving a single notification in the wild, we comparatively evaluate many different notifications, isolating the factors that contribute to particular reactions and understanding. Furthermore, while Huh et al. studied a notification that LinkedIn sent their users encouraging them to change their LinkedIn passwords due to a breach of LinkedIn, we focus on cases where cross-site password reuse substantially complicates the situation. To our knowledge, prior work has neither focused on notifications about cross-site password reuse nor compared such notifications.

3 STUDY 1

Study 1 explored current password-reuse notifications. It investigated user perceptions of, and reactions to, such notifications. Both Study 1 and Study 2 were approved by our IRB.

As both Study 1 and Study 2 rely on respondents’ self-reports of their feelings and actions they would intend to take, percentages reported below should not be taken as ground truth. Rather, we use our survey findings to inform the design of improved password-reuse notifications. While observing notification response in the wild may produce more accurate absolute reports of behavioral response, such observational studies fail to allow us to understand *why* people may react in certain ways and improve those reactions. Thus, similar to prior work on SSL warnings [15], we use Study 1 to identify potential areas of improvement for current password-reuse notifications, developing a model notice that we evaluate in Study 2.

3.1 Recruitment and Survey Structure

We recruited participants on Amazon’s Mechanical Turk, requiring that workers be 18 years or older, live in the US, and have a 95 %+ approval rate. We advertised our study as a survey about “online account notifications.” To avoid recruitment biases, we did not mention security or privacy. Study 1 was a scenario-based survey expected to take 15 minutes. Respondents were compensated \$2.50.

Respondents were first introduced to the survey scenario: “In the following survey, you will be asked to imagine that your name is Jo Doe. You have an online account with a major company called AcmeCo and can access your account through both a website and a mobile application. Imagine that this account is important to you and that it is like other accounts you may have, such as for email,

Table 1: Prominent characteristics of the six Study 1 notifications. The name of the condition identifies the provider that currently uses that text.

	Netflix	LinkedIn	Instagram	Google red bar	Google email	Facebook
Explicitly Mentioned						
Password reuse			✓			✓
Outside breach	✓		✓			
Outside security incident		✓				✓
Suspicious activity	✓			✓	✓	
Review activity				✓	✓	
Forced password reset	✓	✓				✓
Recommended password reset			✓			
Delivery Method						
Browser				✓		✓
Email	✓	✓			✓	
Mobile			✓			

banking, or social media.” Then, respondents were presented with one of six password-reuse notifications (Section 3.2).

Three sets of questions followed. The first set measured respondents’ overall understanding of the notification by asking what may have caused it to be sent through two open-ended questions: “In your own words, please describe what this notification is telling you” and “In your own words, please describe all factors that may have caused you to receive this notification.” The second set asked respondents to list three feelings they might have and three actions they might take upon receiving the notification, and why. The third set presented seven statements, in randomized order, about perceptions of the effectiveness of the notification’s explanation of the situation, its delivery method, and its apparent legitimacy. Respondents gave a Likert-scale response and free-text justification for each. Finally, respondents reported the following demographic information: gender, age, highest degree attained, and technical expertise. Appendix A.1 contains the full text of the survey.

3.2 Conditions

In Study 1 we evaluated six real notifications used by online account providers. To collect such notifications, four members of the research team searched for notifications sent by major online account providers after known data breaches that had been posted online or on social media. We deemed a notification in scope if the potential risk may have originated from password reuse. We verified all notifications as legitimate (not phishing) by cross-referencing Twitter accounts, company security blogs, and news articles.

We collected 24 real notifications about password reuse. To select a set of representative notifications, we used affinity diagramming [28] to categorize and group similar notifications. Three members of the research team created separate affinity diagrams for major types of variation in notifications. We uncovered stark differences in the degree to which a cause was explained, what actions were required or suggested, and how the notification was delivered.

From the 24 notifications, we selected six that captured the range of variation within and across these three dimensions. Table 1 summarizes the notifications, which we refer to with the name of the provider who originally sent that notification. To avoid priming

Figure 1: A notification we tested, rebranded from LinkedIn.

respondents with biases they might have about the companies that originally sent these notifications, as well as to minimize potential confounds from the visual layout of the notification, we visually rebranded all notifications to be from a hypothetical online account provider “AcmeCo.” Figure 1 depicts the rebranded LinkedIn notification. The five other notifications are in Appendix B.1.

Prior to launching the study, we conducted cognitive interviews to refine the survey wording iteratively and verify the intelligibility of questions. A limitation of survey studies is that responses can suffer from self-report and social desirability biases that may affect accuracy. Respondents’ reported reactions may differ from their reactions had they received the notification in real life. In line with survey best practices, we worked to minimize relevant biases through the aforementioned pre-testing and by using softening language to minimize social-desirability bias [36]. Despite potential biases, related work has shown that while survey responses to security messages may be biased, they correlate strongly with real-world reactions [50]. Our results should thus be interpreted as trends of user behavior rather than precise frequency estimates.

3.3 Analysis Methods and Metrics

We collected both quantitative and qualitative data. Our quantitative analysis centered on the seven statements to which participants responded on scales (four on Likert scales and three on other scales), which we treated as ordinal. To evaluate whether responses differed significantly across notifications while controlling for the effects of demographic factors, we built ordinal logistic regression models. In each model, the dependent variable was the set of Likert-scale responses to a given statement. We used the following independent variables: the notification the respondent saw; the respondent’s age; the respondent’s gender; the respondent’s level of education; and the respondent’s technical background. All independent variables

were treated as categorical; we selected the most prevalent categorical value as the baseline. We chose the LinkedIn notification as the baseline category for the notification term as it was most representative (as determined through affinity diagramming) of the 24 messages we originally collected.

In particular, we built parsimonious regression models using stepwise backward-elimination, minimizing AIC. All of these parsimonious final models contained the notification term. To determine whether this notification term was significant, we compared these final models to their analogous null models (removing the notification term) to calculate an omnibus p-value, which we report as the *regression* p-value. Furthermore, we report significant individual factors in the regression by providing that factor’s log-adjusted regression coefficient (e.g., odds ratio, denoted *OR*) and p-value. Our accompanying technical report¹ contains the full regression tables. If this omnibus test was significant, we performed pairwise comparisons between notifications using the Mann-Whitney U test, for which we report the test statistic (*U*) and the p-value. We set $\alpha = .05$ for all tests and used the Holm method to correct for multiple testing within a family of tests.

Finally, we analyzed responses to open-answer survey questions via qualitative coding. A member of the research team read the responses and performed a thematic analysis, iteratively updating the codebook as necessary. The researcher then used axial coding for consolidation and clarification, resulting in 11 themes for the causes of receiving the presented notification. To focus on recurring themes, we report codes that occurred for at least 10 % of responses. We also performed a thematic analysis of respondents’ free-text explanations in the third set of questions to more fully understand *why* respondents answered the questions the way they did. This process was largely the same as the one for the first section of questions, resulting in four or more codes for each question.

In addition, respondents provided in free text three feelings and three intended actions in response to the notification. We cleaned responses to condense tenses differences and misspellings. As the survey asked for these in any order, responses were not ranked during analysis. We used the NRC Word-Emotion Association Lexicon [42] to group feelings as positive, neutral, or negative.

4 STUDY 1 RESULTS

In Study 1, we found that the current password-reuse notifications we tested elicit worry and fear. While the notifications do motivate some respondents to report intending to change their passwords, respondents do not report intending to change their passwords in sufficiently security-enhancing ways. For example, many respondents report planning to make small adjustments to existing passwords, which will likely leave them susceptible to password-reuse attacks. This lack of sufficient action may be attributed in part to notification confusion. A majority of respondents report not understanding the notification, and their mental model may, therefore, be insufficient to elicit an appropriate response.

4.1 Respondents

180 people responded to our survey. Their ages ranged from 18 to 74 years, though most respondents were between 25 and 34 years

¹<https://super.cs.uchicago.edu/papers/ccs18-tr.pdf>

Figure 2: Sentiment analysis (using NRC EmoLex [42]) of respondents' reported feelings upon receiving a notification.

old. 44.4 % of our respondents were female, and a majority (62.8 %) of respondents had a two-year or higher degree. 70.6 % of respondents reported no experience (education or job) in a technical field.

4.2 Notification Response

Figure 3 highlights respondents' reactions to the notifications.

Notifications elicited negative responses. Of the 540 feelings reported by respondents, *worried*, *afraid*, and *anxious* were the main responses to receiving a password-reuse notification. Figure 2 displays feelings reported by ten or more respondents. Fortunately, some positive feelings, such as *safe* or *relieved*, were also common. As the notifications are communicating potential risks to accounts, it makes sense that an overall negative sentiment dominated. However, a password-reuse notification should induce more positive responses, as they are ultimately helping their users.

Notifications were concerning. Across notifications, most respondents (66.7 %) reported that they would feel extremely or moderately concerned upon receiving the notification. R56 explained, "The potential for losing an account and sensitive information is something to be concerned about. Anyone who wouldn't feel concerned is either ignorant or lying." Only 3.3 % reported no concern.

Respondents' reported concern differed significantly across notifications (regression $p = 0.003$). Respondents found the Facebook ($OR = 3.3$, $p = 0.011$) and the Google email notifications ($OR = 4.1$, $p = 0.003$) more concerning than the LinkedIn notification, the control in our regression. Respondents also reported a greater concern about receiving the Facebook notification ($U = 674.5$, $p = 0.019$) and the Google email notification ($U = 730.5$, $p = 0.011$) than the Instagram notification. 89.7 % reported the Facebook notification as concerning, 83.9 % reported the Google email notification as concerning, 54.9 % reported the LinkedIn notification as concerning, and 53.1 % reported the Instagram notification as concerning.

Ignoring the notifications would have consequences. Most respondents disagreed or strongly disagreed that ignoring the notification they received would not have consequences (77.1 %). Responses differed significantly across notifications (regression $p = .045$). Respondents noted potential consequences that included harm to their account "because hackers could steal my info" (R150), as well as "being locked out of my accounts" (R84). However, a

sizable minority was unsure (16.7 %). These "unsure" respondents wanted to get more information from AcmeCo, which shows the importance of clear communication of the situation at hand in a password-reuse notification. Finally, a few respondents were dismissive of any consequences: "Acme has so many accounts that the chances that my account is hacked are pretty slim" (R81).

Facebook, Google email notifications high-priority. Across notifications, responses about the priority of taking action differed significantly (regression $p = .012$). Compared to the LinkedIn notification, a significantly larger fraction of respondents reported that taking action in response to the Facebook ($OR = 4.3$, $p = 0.003$) and Google email ($OR = 3.0$, $p = 0.022$) notifications would be a high priority. Significantly more respondents reported the same for the Facebook notification relative to the Instagram notification ($U = 633.0$, $p = 0.044$). 100 % of respondents who received the Facebook notification, 93.5 % of those who received the Google email notification, 80.6 % of those who received the LinkedIn notification, and 71.0 % of those who received the Instagram notification reported that taking action in response to their respective notifications would be a high priority. We hypothesize respondents perceived the Facebook and Google email notifications to be a higher priority because the Facebook notification prohibited users from logging in, and Google's email included a prominent red color.

Nearly all respondents indicated that taking action in response to the notification was a priority. Across all notifications, 95.6 % of respondents indicated that taking action in response to receiving the notification would be either a very high, high, or a medium priority. In their free-response justifications, 76.6 % of respondents explained that they wanted to protect their personal information or prevent unauthorized account access. 29.4 % of responses specified that the high priority was due to a lack of time: "The quicker I act, the safer my account will be" (R54).

4.3 Understanding of the Notification

Few respondents recognized the notification's real cause. We asked respondents to describe all factors that may have caused them to receive that notification. Most respondents believed that the notification was sent because of circumstances beyond their control. R171 was typical in failing to account for password-reuse attacks as a cause, stating, "The chances of someone guessing that I use the same password are still incredibly low. Still, I would be worried that the password might be too common." 60 % of respondents attributed the notification to someone hacking their account or unsuccessfully attempting to log in. While this makes sense, as some notifications convey that someone may have tried to log in to their account, this is not the full truth: the login may have been attempted as part of a password-reuse attack. Further, 21.1 % of respondents believed that it may have been sent in error, as a false alarm due to the real user of the account using a new device, signing in from a new location, or entering the incorrect password too many times. A minority mentioned the potential real cause of the notification: either a data breach (20.6 %) or password reuse by the account holder (18.8 %).

4.4 Intended Response to Notification

Most respondents do not intend to change their password. While most respondents agreed that taking action was a priority,

Figure 3: Respondents reported their priority of taking *action* in response to the notification. Respondents also reported their agreement of whether the notification was sent via the appropriate *method*, could be ignored without *consequence*, would be sent by *real* companies, and explained how to *resolve* the situation. Finally, respondents reported the level of *concern* they would expect to have upon receiving the notification, and whether they had received such notifications *before*.

they disagreed on what to do and volunteered a wide variety of examples. Respondents wrote that they would take actions such as changing their password (29.3 % of respondents), investigating the situation (18.6 %), and logging into their account (15.4 %) in response to receiving the notification. While the self-reported intention to change their password was the most common, it is nevertheless extremely low as an absolute percentage. This is a cause for concern, as securing an account through password changes should be a priority for all users in situations of password reuse.

Overall, respondents found the notifications informative. A majority of respondents (62.8 %) either agreed or strongly agreed that the notification they received explained how to resolve the situation by giving specific, clear instructions (58.3 %). However, 26.7 % believed it did not do so, and 10.0 % of respondents indicated that resolving the situation would require more background information. As R137 explained, “I need more information as to what happened before I just blindly change my password.”

Notifications with prominent explanations perceived as most informative. We observed significant differences across notifications in respondents’ perceptions of whether the notification explained how to resolve the situation (regression $p < 0.001$). The agreement that the notification explained the situation differed starkly across notifications: LinkedIn (80.6 % of respondents agreed or strongly agreed), Facebook (75.9 %), Netflix (75.0 %), Instagram (68.7 %), Google email (54.9 %), and Google red bar (21.6 %). Agreement was higher for the LinkedIn notification than for the Google email ($OR = 0.2, p < 0.001$) and the Google red bar ($OR = .03, p < 0.001$) notifications. Compared to the Google red bar notification, agreement was also significantly higher for the Facebook ($U = 646.5, p < 0.001$), Google email ($U = 642, p = 0.011$), Instagram ($U = 177.5, p < 0.001$), and Netflix ($U = 131.0, p < 0.001$) notifications. The low reported percentages for the Google email and Google red bar notifications make sense because both notifications had a link that had to be clicked for more information and explanation. The other notifications had more detail and instructions in the notification itself.

4.5 Reactions to Structure and Delivery

Most respondents agreed that the notification they received used the appropriate method of contacting them (65.0 %), primarily because it was easy, convenient, or fast (58.3 % of respondents). However, some respondents would have preferred a more immediate method (17.8 %) or multiple methods (11.6 %). Agreement about the method’s appropriateness differed across notifications (regression $p < .001$).

Email perceived as the most legitimate delivery method. The Google email, LinkedIn, and Netflix notifications, all sent by email, were reported to be delivered with the most appropriate method and to seem the most legitimate. This is perhaps due to some respondents’ justification that email is official (10 %), and that they may have seen similar email notifications in the past. Respondents were more likely to report that the LinkedIn notification was appropriate than the Facebook ($OR = 0.2, p < 0.001$), Google red bar ($OR = 0.1, p < 0.001$), and Instagram ($OR = 0.3, p < 0.010$) notifications. For the LinkedIn, Instagram, Facebook, and Google red bar notifications, respectively, 98.6 %, 62.5 %, 51.7 %, and 48.2 % of respondents reported agreement that the notification they received was delivered with the appropriate method. Fewer respondents found the Facebook notification appropriate than the Google email notification ($U = 265, p = 0.049$).

Respondents’ expectations regarding real companies sending the notification also differed across conditions (regression $p = .012$). While 96.7 % of respondents who saw the LinkedIn notification reported expecting real companies would send it, only 67.7 % reported the same for the Instagram notification ($OR = 0.2, p = 0.003$).

Our notifications were relevant to real situations. Overall, most respondents agreed (86.7 %) that they would expect real companies to send notifications like these when necessary. Respondents reported receiving notifications similar to this one in the past: 52.1 % of respondents indicated receiving a similar notification a few times, and 9.5 % many times. Those who had received similar notifications explained that they were from sign-ins on other devices (20.0 %) or financial services (13.3 %).

Figure 4: Our model notification with the parts that varied highlighted in color and further specified in Table 2.

5 PASSWORD-REUSE NOTIFICATION GOALS

Password-reuse notifications take on a challenging task as the situation at hand is the cumulative result of multiple parties' actions. Further, the level of risk to convey and the appropriate actions to suggest or require are not always clear. While research has investigated best practices for other types of security notifications (cf. Section 2), we sought to create a framework for evaluating password-reuse notifications. Drawing on the Study 1 results, we identified five goals that effective notifications should achieve sufficiently: timeliness, legitimacy, action, background, and trust. We used these goals as a framework to evaluate notifications in Study 2.

First, notifications should reach their intended audience in a **timely** manner. A notification about a compromised password is only useful if the user sees the notification to create a new one.

Second, notifications should be perceived as **legitimate**. Some respondents in Study 1 were hesitant to trust our notifications, believing that they might be phishing. The presence of hyperlinks was cited as an indicator of phishing, and a few respondents were skeptical of any email that required password changes at all.

Third, a password-reuse notification should lead to **actions** that improve the security of the directly affected online account. Ideally, this would include taking productive actions for other accounts that may be at risk (i.e., those where similar passwords were used), as well as advising against unproductive or unrelated actions.

Fourth, the **background** information provided by a notification should be easily understood. In Study 1, 12.8 % of respondents were confused by how one service "got" their passwords for another service, which could potentially lead to confusion. Not all users will understand the mechanisms behind password databases or cryptographic hashes, but the root cause of the notification (password reuse) must be clearly conveyed.

Table 2: Notification dimensions varied in Study 2.

Delivery Medium	
<i>model</i>	Delivered by email
<i>inApp</i>	Mobile in-app
<i>mobile</i>	Mobile push notification and in-app
Incident	
<i>model</i>	This incident was likely a data breach of a service unrelated to AcmeCo, but because many people reuse similar passwords on multiple sites, your AcmeCo login information may have been affected.
<i>usBreach</i>	This incident was likely a data breach of one of our services.
<i>vagueCause</i>	—
Account Activity	
<i>model</i>	While we have not detected any suspicious activity on your AcmeCo account, ... as a precaution.
<i>suspicious</i>	Because we have detected suspicious activity on your AcmeCo account, ...
<i>omitActivity</i>	—
Remediation	
<i>model</i>	... you must create a new password ...
<i>recommend</i>	... we recommend that you create a new password.
Other Accounts	
<i>model</i>	Change all similar passwords on other accounts.
<i>noOthers</i>	—
Extra Actions	
<i>model</i>	To further improve your online security, we recommend: <ul style="list-style-type: none"> • Enabling AcmeCo's Two-Factor Authentication. • Using a password manager.
<i>noExtras</i>	—

Fifth and finally, notifications should improve **trust** between providers and users. Account providers send notifications to increase the security of users' accounts with that provider, as well as potentially with other providers, too. Therefore, notifications should aim to engender users' trust.

6 STUDY 2

In Study 1, we found that the content of a password-reuse notification impacted respondents' understanding of the situation at hand, as well as whether they would intend to take action in response. In Study 2, we sought to better isolate the factors of effective notifications by exploring the impact of making small changes to the content or delivery of these notifications. Our design of Study 2 focuses on key results from Study 1, along with the goals outlined in Section 5. We had six core research questions for Study 2.

First, we consider the delivery medium. The timeliness of a notification is largely determined by how it is sent to the recipient. Mobile push notifications interrupt the current workflow, whereas emails or in-app notifications require users to actively check those sources. The delivery medium of the notification may also change respondents' perception of the legitimacy of the notification.

RQ 1A: How does the delivery medium of a password-reuse notification affect its perceived effectiveness?

RQ 1B: If you, an online account provider, are breached, how important is the delivery medium in which you send your notification?

Next, we consider mentions of suspicious account activity and the nature of the data breach. These details address the goal of providing adequate background for users to understand the situation.

RQ 2: How does explicitly identifying the root causes of the incident influence the notification's effectiveness?

RQ 3: How does mentioning suspicious account activity influence the notification's effectiveness?

Depending on the importance of the account and the incident, notifications should force a password reset.

RQ 4: If a password change is only recommended, instead of required, will users report that they would change their passwords?

Finally, we consider various security suggestions beyond password changes. We hypothesize that these suggestions could improve the user's trust in the account provider by appearing to demonstrate proactive approaches to security.

RQ 5A: Is it important to explicitly recommend password changes on other sites in a notification?

RQ 5B: Is it important to explicitly recommend pro-security actions (e.g., 2FA, adopting a password manager) in a notification?

RQ 5C: If your service is breached, is it important to explicitly recommend password changes on other sites and pro-security behaviors beyond changing your password?

RQ 6: Will users report taking pro-security actions if they are not explicitly mentioned in a password-reuse notification?

6.1 Study 2 Conditions

We began developing our Study 2 conditions by creating a model notification (shown in Figure 4) that synthesized the individual aspects of notifications that were most successful in Study 1, filling in gaps relative to our aforementioned design goals. To disambiguate the impact of each aspect of the model notification's content and delivery, we created 14 additional variants of the model, each of which differed in a targeted way. These variants, as detailed in Table 2, reflect changes in the delivery method, description of the security incident, mention of account activity, suggested remediation, reference to other accounts, and additional pro-security actions mentioned. Each respondent was randomly assigned to see either the model notification or one of these fourteen variants. When presenting our results, we refer to these variants using multi-part names based on the nomenclature defined in Table 2. Special attention was given to increase the likelihood that our respondents perceived the notification as legitimate, rather than as phishing. Appendix B.2 contains additional images of the variants.

6.2 Study 2 Structure and Recruitment

We recruited respondents on Amazon's Mechanical Turk, again advertising a survey about online account notifications with no mention of security or privacy. Requirements for participation were the same as for Study 1, and participation in both Study 1 and Study 2 was prevented. The survey was again scenario-based, but this survey was structured into five sections and added additional questions to explore topics raised during the analysis of Study 1. Each respondent was compensated with \$2.50 for completing the 15-minute survey. Respondents were introduced to the survey scenario with the same text as Study 1 (cf. Section 3.1) and were then presented with their assigned notification.

The first section of survey questions measured respondents' overall reported conceptions of the notification with questions similar to Study 1, but with a key modification: respondents were given eleven closed-answer choices of the *causes* of receiving the notification. We chose to give closed-answer choices to measure

explicitly whether or not they expected some factor might have caused the situation, rather than relying just on what they thought to write. We based these choices on the responses to Study 1's open-ended version.

The second section asked whether respondents would intend to *change their passwords* for AcmeCo, as well as for other accounts. This section also contained follow-up questions about why they would or would not intend to change their passwords, as well as how they would create and memorize such passwords. The third section asked them to report their security perceptions of, and likelihood to take, ten *actions* beyond changing their password. These actions were again closed-answer and were selected based on free-text responses from Study 1.

The fourth section asked respondents about their *perceptions* of the notification with questions based on the corresponding section from Study 1 but modified to align with Study 2's research questions. The final section solicited the following *demographic* information: gender, age, highest educational degree attained, and technical expertise. We also asked respondents to report any previous experiences being notified about data breaches and history of having others gain unauthorized access to their online accounts. Appendix A.2 contains the survey text. As in Study 1, responses are reported behavioral intentions, rather than actual behavior. We again mitigated biases with softening language and pre-testing.

6.3 Analysis Method and Metrics

We again use regression models in our analysis. We had both binary (whether respondents selected each of the eleven potential causes, whether respondents reported intending to change any passwords or take the ten additional actions) and ordinal (responses on scales regarding perceptions of the notification, as well as a Likert-scale agreement with the security benefit of the ten actions) dependent variables. For binary dependent variables, we built logistic regression models. For ordinal dependent variables, we built ordinal logistic regression models. The independent variables were the notification, all covariates used in Study 1 (the respondent's age range, gender, education level, and technical background), whether the respondent had ever been notified that their information was exposed in a data breach and whether the respondent had experienced unauthorized access to an online account. These final two variables are proxies for prior experience with breaches [47, 49]. All independent variables were treated as categorical.

As in Study 1, we built parsimonious models through backward elimination. The full regression tables are again contained in our companion technical report. To determine whether the omnibus notification term was significant, we compared these final models to their analogous null models (removing the notification term) to calculate an omnibus p-value, which we report as the *regression* p-value. If the notification term was removed in backward elimination, we treated the notification as non-significant. For significant individual factors, we report the odds ratio and p-value.

When the omnibus notification term was significant, we made 18 comparisons between pairs of notifications to investigate our six research questions directly. For ordinal data, we used Mann-Whitney U tests (reporting *U* and the p-value). For categorical data more naturally expressed as a contingency table (e.g., whether and

how respondents intended to change their password), we performed χ^2 tests if all cell counts were greater than five, and Fisher's Exact Test (denoted *FET*) if they were not. We again set $\alpha = .05$ and used Holm correction within each family of tests.

Finally, in a process analogous to that for Study 1, we qualitatively coded free-response data.

7 STUDY 2 RESULTS

Across all variants of the model notification, respondents reported anticipating serious consequences to ignoring the notification and reported believing that changing their password would benefit their account security. While a majority of respondents indicated that they would intend to change their passwords, their intended password creation strategies would continue to expose them to password-reuse attacks. Unfortunately, many respondents did not perceive password reuse to be the root cause of the situation.

We found that adding extra security suggestions increases perceived risks, which may help the notification convey the seriousness of the situation and the need to take action. Omitting information about account activity or being vague about the origin of the security incident, however, warps perceptions of the situation.

7.1 Respondents

There were 588 respondents in Study 2. Most respondents were between the ages of 25 and 34 (44.6%), although 11.2% were younger and 44.1% were older. 48.4% of the respondents identified as female. Over half of our respondents had a two- or four-year degree, and 10.8% held higher degrees. A quarter of respondents reported experience (education or job) in technical fields.

53.2% of respondents in Study 2 indicated that they had been affected by a prior data breach. Most respondents were notified via email (55.9%), although receiving physical mail (17.3%) and reading the news or browsing social media (18.2%) were other common notification methods. The most common data breach mentioned by respondents was Equifax (12.1% of respondents). Less than one-third of respondents reported unauthorized access to an account. Of the 188 respondents that reported someone had gained unauthorized access to one of their online accounts, 23 personally knew the attacker, whereas 155 did not.

7.2 Perceived Causes of the Scenario

Many respondents did not perceive password reuse to be a cause of the situation. From among eleven potential causes of receiving the notification, we asked respondents to choose all they felt applied. Unfortunately, across all notifications, a minority of respondents chose "you reused the same or similar passwords for multiple online accounts" as a potential cause even though many variants of the notification mentioned password reuse. For example, the *model* notification (control condition) explained that their "AcmeCo account login and password may have been compromised" due to a data breach of a service unrelated to AcmeCo "because many people reuse similar passwords on multiple sites." Nonetheless, only 44.7% of respondents who saw the *model* notification chose password reuse as a cause of receiving the notification.

The rate of selecting password reuse as a cause varied by condition (regression $p < .001$). Among all variants, *model-suspicious*

(named using the keywords in Table 2) was most effective at conveying that password reuse was a potential cause. This variant augmented the model notification by noting *suspicious* activity had been detected on the account. Nonetheless, only 57.9% of respondents chose password reuse as a possible cause, which did not differ significantly from the control. Unsurprisingly, four variants that mentioned that AcmeCo itself suffered a breach had significantly lower rates of choosing password reuse as a cause: *model-usBreach-noMobile* (2.4%, $OR = 0.03$, $p < 0.001$); *model-usBreach-inApp* (2.4%, $OR = 0.03$, $p = 0.001$); *model-usBreach-noOthers* (10.0%, $OR = 0.13$, $p = 0.001$); and *model-usBreach-noOthers-noExtras* (2.6%, $OR = 0.03$, $p = 0.001$).

For *model-vagueCause*, 10.3% of respondents chose password reuse as a cause, which was also significantly lower than the control ($OR = 0.16$, $p = 0.003$). This is notable because that notification mentions a vaguely-worded "potential security incident" that may have led to a credential compromise, typical of many widely deployed notifications even when password reuse is the culprit.

Respondents also rarely chose "you have a weak password for your AcmeCo account" as a potential cause. Across conditions, only 15.0% of respondents selected this option. This did vary significantly by condition (regression $p = .011$), though we only observed a significant difference for the pair of conditions investigating the impact of mentioning suspicious activity (RQ3). While 38.9% of respondents indicated a weak password as a potential cause for *model-suspicious*, which mentioned suspicious activity, only 4.9% did so for *model-omitActivity*, which did not (*FET*, $p = .009$).

In contrast, across all notifications, respondents most commonly chose that "AcmeCo was hacked" (49.0% of respondents) or that a company unrelated to AcmeCo was hacked (41.7%). Note, however, that conditions varied in whether they reported that AcmeCo or some other company was breached, so these frequencies and the significant differences between conditions (both regressions $p < .001$) are unsurprising. More surprisingly, across conditions respondents selected three additional potential causes at higher rates than password reuse: "Someone hacked your AcmeCo account" (32.5% of respondents); "AcmeCo conducts regular security checks and this is just a standard security notification" (28.2%); "Someone is trying to gain unauthorized access to your account by sending this email" (27.4%). These did not vary significantly across conditions.

7.3 Creating New Passwords

Respondents rated whether fifteen potential actions would improve their account security. Six of these actions related to password changes. In addition, respondents selected whether or not "if [they] received this notification about an online account [they] had with a real company" they would change their password for that company. For brevity, we refer to this below as changing their AcmeCo password. We also asked them to report their likelihood to take five actions related to changing passwords, but for services other than the one that sent them the notification.

Most respondents perceived unique passwords as good for security. Overall, 86.0% of respondents agreed that changing their AcmeCo password to "a completely new password unrelated to the old one" would improve their account security. Most cited a "better safe than sorry" rationale for changing their password. For example,

Figure 5: Respondents' intentions for creating new passwords for their account on AcmeCo (who sent the notification) and on other providers. Respondents could change all passwords, only passwords that were the same or similar, only passwords for important accounts, or none at all.

Figure 6: Of respondents who intended to change passwords, their stated strategies for doing so for their account on AcmeCo and on other providers. They could generate a new password with a password manager or browser, make a completely new password, modify the old password, reuse a password they already use, or apply some other strategy.

R55 wrote, “It would bring me peace of mind to know I had done what I could to protect myself and my account.” Yet, 34.6 % also answered that changing their AcmeCo password to “a modification... of the old one” would improve their account security, while 26.0 % answered similarly about changing their AcmeCo password “to a password I use for another online account.”

To prevent password-reuse attacks, users should have a unique password for each account, and 84.1 % of respondents agreed that doing so would improve their account security. However, a concerning large fraction of respondents — 50.2 % — agreed that changing “all of my similar passwords on other online accounts to one new password” would improve their security. Unfortunately, doing so makes them susceptible to future password-reuse attacks. We did not observe significant differences in responses across conditions for any of these six actions related to password changes.

If they received our notification in real life, respondents would change their password, but ineffectively. The vast majority of respondents — 90.3 % — reported they would change their passwords if, in real life, they received the notification they saw (Figure 5). However, among these respondents, only 1.4 % of them said they would change their password to something completely unrelated. Additionally, only 9.7 % of them said they would use a password manager or their browser to generate the password.

The majority of respondents' new passwords would continue to expose their accounts to the same risks (Figure 6). Most respondents — 59.0 % — reported intending to create their new AcmeCo

password by changing a few characters in the old password, while 11.4 % reported intending to simply reuse another password they already used elsewhere. In reality, these strategies would not truly resolve their problems and would continue to facilitate password-reuse attacks. Attackers have adapted to users' tendency to modify passwords in small ways (e.g., common character substitutions, insertions, and capitalizations) and apply such common transformations in password-reuse attacks [10, 62]. Furthermore, self-reported intentions typically overreport actual behavior [58], suggesting that these results may already be overly optimistic.

Respondents' stated likelihood to “leave [their AcmeCo] password as-is” varied by condition (regression $p = .012$). Respondents who saw *model-{noOthers}-{noExtras}* were more likely to say they would keep their current password than those who saw *model* ($OR = 2.4, p = .042$) or *model-{noOthers}* ($W = 533, p = .040$). Respondents were also more likely to state the same if they had not previously received a data-breach notification ($OR = 1.4, p = .034$) or if they had a background in technology ($OR = 1.5, p = .038$). We hypothesize this last result may stem from overconfidence.

Some perceptions of security also varied across demographic factors. Female respondents were more likely to rate having unique passwords for all accounts as secure ($OR = 1.5, p = .012$) and less likely to rate keeping their current password as secure ($OR = 0.6, p = .008$). Respondents who had not previously received a data-breach notification were more likely to rate modifying their old password as secure ($OR = 1.4, p = .017$) and less likely to rate changing it to something unrelated as secure ($OR = 0.6, p = .002$). Surprisingly, respondents with a background in technology were also less likely to rate the latter as secure ($OR = 0.6, p = .007$).

Those who avoid password changes may do so due to suspicion of notifications or invincibility beliefs. Of the 52 respondents (9 % of the total) who said they would not change their password, 25 reported that it was because they would need to verify that the notification was legitimate, rather than a phishing attack. R534 elaborated that they would “wait and go to AcmeCo's website and see what was going on first.” Eight others said they would not change their password because of memorability concerns.

Seventeen respondents expressed various beliefs of invincibility: eleven said they use unique passwords on every account and thus would not worry about one password being compromised, while six believed their passwords were strong enough to eliminate the risk of compromise. As R2 wrote, “It is a very good password and I doubt someone would waste the time trying to crack it.” While non-experts have difficulties judging password strength [60], Pearman et al. observed in an in-situ study of 154 participants an average password strength that could resist up to 10^{12} guesses [44]. At the same time, real-world offline guessing attacks are on the order of 10^9 to 10^{12} guesses per day on a single GPU even against hash functions like scrypt [25]. Others consider 10^{14} guesses realistic in offline attacks [18]. While rate-limiting and risk-based authentication slow online guessing [21], password reuse remains a threat [38, 57].

Most respondents believed changing passwords for other accounts with similar passwords would improve security, yet they did not intend to do so. Even though, as previously mentioned, 84.1 % of respondents agreed that using unique passwords for each of their accounts would improve security, 35.2 % of respondents reported that they did not intend to change passwords for

Figure 7: Respondents' perceptions of whether actions would increase security, as well as their stated intention (across conditions) of taking those actions upon receiving the notification. We group actions by whether they relate to password reuse.

any accounts other than AcmeCo. As Figure 5 shows, an additional 15.6 % only intended to do so for accounts where they used exactly the same password, while 14.1 % only intended to change passwords for important accounts. The largest portion of respondents would not change their passwords on other accounts because they did not perceive connections between the account addressed by the notification and any other account. R69 explained, “Unless I heard from a company that was hacked, I’m not concerned.” Furthermore, respondents believed that if the account providers were unrelated, then the risks to account security also must be unrelated. A few respondents speculated that the threats were unrelated because “a potential hacker likely doesn’t know my additional accounts exist” (R23). Unfortunately, because reuse of both usernames and passwords across services is common [10, 62], attackers know to try the same or similar credentials across unrelated services.

In contrast, only 15.3 % of respondents reported intending to change their passwords on all other accounts, while 19.7 % reported intending to change all passwords that were similar to the one that was compromised. Unfortunately, even for respondents who said they would intend to change other passwords, their intended password-creation strategies would leave many at risk. The majority of respondents again intended to either modify (46.5 %) or directly reuse (9.6 %) passwords they already used elsewhere, as shown in Figure 6. On a more positive note, 38.8 % of respondents reported intending to use a password manager or browser to generate these other passwords, which balances usability and security for changing many passwords at once.

7.4 Taking Other Security-Related Actions

For nine additional actions unrelated to password changes, respondents again rated their expectation of how these actions impact

security, as well as their likelihood to take these actions upon receiving the notification. Notifications should encourage actions that are both *productive* and *relevant* for addressing password reuse. To account for these nuances, we included four actions that can potentially address password reuse, as well as five that are only tangentially related to the situation, as shown in Figure 7.

Notifications encourage 2FA adoption, yet are less effective at encouraging the use of password managers. The notifications had a divergent impact on two of the actions most relevant to mitigating threats from password reuse: enabling 2FA and using a password manager. While 83.3 % of respondents agreed that enabling 2FA would improve their security and 64.0 % rated it likely that they would do so, only 44.3 % agreed that using a password manager would improve their security, and only 37.3 % rated it likely they would adopt one after receiving the notification.

In contrast, 78.8 % of respondents agreed changing their password more frequently would improve security, and 51.9 % rated it likely they would do so. Furthermore, 80.9 % of respondents agreed that reviewing the recent activity on their account would improve security, and 89.5 % rated it likely they would do so.

Notification variants did not impact the likelihood of taking these actions. Which notification respondents saw did not significantly impact their stated likelihood of taking any of these nine actions. However, some demographic factors did. Respondents with a background in technology expressed a higher likelihood of using a password manager ($OR = 1.7, p = .002$), using an identity theft protection service ($OR = 1.6, p = .008$), and changing the way they lock their phone ($OR = 1.4, p = .033$) upon receiving the notification. Finally, female respondents expressed being more likely to review the activity on their account ($OR = 1.5, p = .022$).

Notification variants minimally impacted security perceptions. Respondents' agreement that updating their account's security questions would improve security varied across notifications (regression $p = .012$), though we did not observe the notification to significantly impact perceptions of any of the other eight actions. Compared to respondents who saw *model*, those who saw *model-vagueCause* ($OR = 2.7, p = .017$) or *model-suspicious* ($OR = 3.5, p = .004$) were more likely to agree that updating their security questions would improve security. We observed the same effect for three notifications that mentioned that AcmeCo itself had been breached: *model-usBreach-mobile* ($OR = 3.6, p = .002$), *model-usBreach-inApp* ($OR = 2.4, p = .035$), and *model-usBreach-noOthers* ($OR = 2.6, p = .026$). Female respondents were more likely to agree that it would improve security ($OR = 1.4, p = .047$), while those who had never received a data-breach notification were less likely to do so ($OR = 0.6, p = .009$).

Demographic factors were correlated with variations in respondents' perceptions of how these actions impacted security. Female respondents were more likely to agree that using an identity theft protection service ($OR = 1.6, p = .003$), changing their password more frequently ($OR = 1.7, p = .001$), and changing how they lock their computer ($OR = 1.5, p = .011$) would improve security. Respondents with a background in technology ($OR = 0.5, p < .001$) and those who had never received a data-breach notification ($OR = 0.7, p = .015$) were also less likely to agree with this statement. Respondents with a background in technology were also less likely to agree that changing their password in the future improves security ($OR = 0.7, p = .023$), while those who had never received a data-breach notification were less likely to agree that updating software improves security ($OR = 0.7, p = .043$).

7.5 Perceptions of the Notification

Most respondents would act in response within 24 hours. We found that most respondents would anticipate seeing and acting on the notification within a short period of time, despite our notifications varying in delivery method. 87.4 % reported that they would see the notification within 24 hours and 84.5 % would intend to take action within 24 hours; responses did not vary significantly across notifications. Respondents strongly preferred that account providers contact them via email (90.0 %), although SMS (43.9 %), mobile app (32.5 %), and mobile push notification (29.1 %) were also favorable options. Interestingly, this stated preference for email notifications conflicts with some respondents' hesitation to take action because of phishing concerns (Section 7.3).

Respondents' trust was lower when AcmeCo suffered a breach. We found that the level of reported trust varied significantly across notification conditions (regression $p = .004$). Compared to *model*, the reported trust of the provider was, perhaps unsurprisingly, lower for *model-usBreach-inApp*, which stated that AcmeCo itself was breached ($OR = 0.3, p = .003$). In their free-response justification, 13.8 % respondents overall reported decreased trust because they believed it to be AcmeCo's responsibility to prevent such breaches. On the contrary, 8.3 % of respondents' trust did not change, as "any company is bound to have security breaches" (R196). However, across conditions, 45.2 % of respondents increased trust in AcmeCo as a result of the notification, and 35.8 %

reported no change. This was because the notification conveyed a prioritization of their safety (29.3 %) or proactive and transparent policies (13.3 %). An additional 15.1 % of respondents believed that such a notification was simply expected of a company.

Experience with technology and data breaches impacted perceptions. In our models, we also compared the responses of respondents who had prior experience with data breaches to those who had no such experiences. Respondents who had never been notified about being in a breach reported that receiving the notification would lead to greater trust in AcmeCo compared to those who had previously received a data-breach notification ($OR = 1.7, p = .002$). Respondents who had never received such a notification were also more likely to agree that they would not know why they received such a notification ($OR = 1.6, p = .002$), more likely to perceive the notification as official ($OR = 1.6, p = .009$), and less likely to expect companies to send notifications like the one they saw ($OR = 0.7, p = .043$). This may be because prior experience gives respondents some expectations of provider behavior.

Respondents who reported a background in technology were more likely to agree that they would not know why they received such a notification ($OR = 1.6, p = .008$) and more likely to agree that ignoring the notification would have no consequences ($OR = 1.6, p = .006$). They were also less likely to agree that they expected companies to send such notifications ($OR = 0.6, p = .010$), less likely to agree that they would believe such a notification was official ($OR = 0.5, p < .001$), and less likely to prioritize taking action in response ($OR = 0.7, p = .031$). They were also less likely to agree that the notification explained how to resolve the situation ($OR = 0.6, p = .007$) and less likely to report that they would feel grateful about receiving the notification ($OR = 0.6, p < .001$).

8 LIMITATIONS

Like many survey studies, our results suffer from self-report biases. Respondents may have answered questions according to social desirability: selecting the answer they believe they *should* select, rather than their true answer [35]. To mitigate this bias, we did not explain that this was a study about security, and we included softening language in sensitive questions to remind respondents that people may have many different responses. That said, stated intentions are typically an upper bound on actual behavior [58]. As many respondents' intended actions would still leave them vulnerable to further attacks, reality may be even worse. This would be consistent with other researchers' finding that LinkedIn's actual breach notification was ineffective at prompting password resets [29].

Finally, we report on a convenience sample of MTurk workers receiving our hypothetical notifications. Such a design is inherently limited in its ecological validity. However, given that such notifications have rarely been studied, testing notifications for the first time in the field and potentially causing respondents to think they had been breached would create too high of a potential risk to human subjects. As in prior work on other types of notification messages [15], we chose to conduct a formative, controlled study to inform future research on password-reuse notifications.

Figure 8: The notification we found to be the most effective, relative to our notification goals, for respondents in Study 2.

9 DISCUSSION

We performed the first systematic study of how users understand and intend to respond to security notifications about situations related to password reuse. Through two complementary user studies, we identified best practices for the design of password-reuse notifications. Further, we identified where notifications are destined to fall short in helping users fully remediate password reuse issues. Our formative study lays the groundwork for future field studies. We recommend future work that does not rely on self-reporting, instead testing the best practices we developed for password-reuse notifications in more ecologically valid situations.

9.1 Best Practices

Our Study 1 results led us to identify five key design goals for password-reuse notifications (Section 5). Some goals (e.g., timeliness) are obvious from general guidelines about warning design, but the importance of providing an adequate background, as well as the subtle considerations around engendering trust, are more specific to the domain of password reuse. Our model notification in Study 2 performed the best according to these goals, suggesting these best practices for designing password-reuse notifications:

- The notification should be very explicit about the root causes of the situation, i.e., password reuse and a data breach.
- Providers should force a password reset on their service.
- The notification should strongly encourage changing similar passwords on other accounts and thoroughly explain why doing so staves off attacks.
- The notification should explicitly encourage enabling 2FA and using password managers.
- Notifications should be sent via both email and more immediate channels (e.g., a blocking notification upon login).

Table 3: How 24 real-world password-reuse notifications compare to our Study 2 model notification's best practices. Notifications are identified by their sender (and additional details if we collected multiple from the same provider).

Notification	Delivered by email	Mentions password reuse	Forces password change	Suggests changing similar passwords	Suggests 2FA and password manager
Adobe	✓		✓		
Amazon	✓	✓			✓
Carbonite	✓		✓	✓	✓
Digital Ocean	✓		✓	✓	
Edmodo	✓			✓	
Evernote	✓		✓	✓	
Facebook (Accessed)		✓	✓		
Facebook (Confirm Identity)					
Facebook (Logged In)			✓		
Freelancer	✓	✓	✓		
Google (2-Step)	✓				
Google (Someone Has...)	✓				
Google (Suspicious)					
Houzz	✓		✓	✓	
Instagram		✓			
LinkedIn	✓		✓		
Microsoft	✓				
Netflix	✓		✓	✓	✓
Pinterest (Read-Only)			✓		
Pinterest (Suspicious)	✓		✓		
Sony		✓	✓		
SoundCloud		✓	✓		
Spirit	✓		✓		
Spotify	✓	✓	✓	✓	

Therefore, we propose the wording of Figure 8 as a model notification for situations related to password reuse. Unfortunately, few real-world notifications currently follow these best practices. Table 3 compares the 24 real-world notifications we collected in Study 1 to the best practices we identified. None of these notifications met all of the best practices we established. In short, there is much room for improvement in widely deployed notifications.

9.2 Addressing Persistent Misunderstandings

While the real notifications we tested in Study 1 were successful in arousing concern, many respondents were unaware of the correct actions to take in response. The model notification we synthesized for Study 2, and its variants were successful in spurring the vast majority of respondents to report that they would change their password on the site that sent them the notification. This apparent success was tempered, however, by respondents reporting that their new password would often be a minor variation on their previous passwords, or even simply a password reused verbatim from another account. Furthermore, many respondents reported that they would leave their passwords unchanged on providers other than the one that sent them the notification. Collectively, these decisions would leave users vulnerable to future attacks leveraging password reuse.

The model notification had mixed success at encouraging respondents to take two other actions that could potentially mitigate password reuse. Users adopting 2FA erects another barrier for attackers in exploiting reused credentials, and nearly two-thirds of respondents reported being likely to do so after receiving the notifications we tested. Users adopting a password manager and using it to generate unique, strong passwords for each site is among the small number of usable solutions to combat password reuse, yet

under 40 % of respondents reported being likely to do so. These intentions did not vary significantly whether or not the notification explicitly encouraged respondents to take these actions. Future work could investigate whether describing the exact situation the given user is in even more explicitly, as well as why these particular actions are crucial in mitigation, might be more successful.

Our work thus underscores that it is unreasonable to expect users to maintain dozens of distinct and secure passwords simply by telling them to do so. Although notifications are a critical source of information to incite positive change in users' online security behaviors, they are only a band-aid on a gaping wound. In addition to improving notifications, we recommend devising ecosystem-level strategies to combat password reuse. Individual account providers cannot prevent password reuse across services without direct cooperation with others [64]. As our respondents already expressed much confusion about how providers "had this information in the first place, who they got it from and how they got it" (R159), other actors may be better positioned to make a difference.

Password managers and web browsers have a unique viewpoint on the full spectrum of a user's passwords that individual providers do not. Specifically, they have the opportunity to identify and prevent password reuse when users create, change, or import passwords. Unfortunately, current implementations of many password managers and browsers permit users to reuse passwords across accounts, often not even warning those users about why this is problematic. This behavior could be out of fear that users would not use those password managers or browsers if they felt burdened by onerous actions. Future work should thus investigate how password managers and browsers can be more explicit in preventing password reuse while maintaining a positive user experience. The current state of password reuse results from many actors' decisions. Remediation will require the contributions of many more.

REFERENCES

- [1] Parag Agrawal. 2018. Keeping Your Account Secure. https://blog.twitter.com/official/en_us/topics/company/2018/keeping-your-account-secure.html.
- [2] Devdatta Akhawe and Adrienne Porter Felt. 2013. Alice in Warningland: A Large-Scale Field Study of Browser Security Warning Effectiveness. In *Proc. USENIX Security Symposium*. 257–272.
- [3] Lujo Bauer, Cristian Bravo-Lillo, Elli Fragkaki, and William Melicher. 2013. A Comparison of Users' Perceptions of and Willingness to Use Google, Facebook, and Google+ Single-sign-on Functionality. In *Proc. DIM*. 25–36.
- [4] Joseph Bonneau, Cormac Herley, Paul C. van Oorschot, and Frank Stajano. 2012. The Quest to Replace Passwords: A Framework for Comparative Evaluation of Web Authentication Schemes. In *Proc. IEEE S&P*. 553–567.
- [5] Cristian Bravo-Lillo, Lorrie Cranor, Saranga Komanduri, Stuart Schechter, and Manya Sleeper. 2014. Harder to Ignore? Revisiting Pop-Up Fatigue and Approaches to Prevent It. In *Proc. SOUPS*. 105–111.
- [6] Christian Bravo-Lillo, Lorrie Faith Cranor, Julie S. Downs, and Saranga Komanduri. 2011. Bridging the Gap in Computer Security Warnings: A Mental Model Approach. *IEEE Security & Privacy Magazine* 9, 2 (March 2011), 18–26.
- [7] Cristian Bravo-Lillo, Saranga Komanduri, Lorrie Faith Cranor, Robert W. Reeder, Manya Sleeper, Julie Downs, and Stuart Schechter. 2013. Your Attention Please: Designing Security-decision UIs to Make Genuine Risks Harder to Ignore. In *Proc. SOUPS*. 6:1–6:12.
- [8] Jessica Colnago, Summer Devlin, Maggie Oates, Chelse Swoopes, Lujo Bauer, Lorrie Cranor, and Nicolas Christin. 2018. "It's Not Actually That Horrible": Exploring Adoption of Two-Factor Authentication at a University. In *Proc. CHI*. 456:1–456:11.
- [9] Sam Croley ("Chick3nman"). 2018. Abusing Password Reuse at Scale: Bcrypt and Beyond. <https://www.youtube.com/watch?v=5bYvTPVXC18&t=6h05m00s>.
- [10] Anupam Das, Joseph Bonneau, Matthew Caesar, Nikita Borisov, and XiaoFeng Wang. 2014. The Tangled Web of Password Reuse. In *Proc. NDSS*.
- [11] Serge Egelman, Lorrie Faith Cranor, and Jason Hong. 2008. You've Been Warned: An Empirical Study of the Effectiveness of Web Browser Phishing Warnings. In *Proc. CHI*. 1065–1074.
- [12] Serge Egelman, Andreas Sotirakopoulos, Ildar Muslukhov, Konstantin Beznosov, and Cormac Herley. 2013. Does My Password Go Up to Eleven?: The Impact of Password Meters on Password Selection. In *Proc. CHI*. 2379–2388.
- [13] Michael Fagan, Yusuf Albayram, Mohammad Maifi Hasan Khan, and Ross Buck. 2017. An Investigation Into Users' Considerations Towards Using Password Managers. *Human-centric Computing and Information Sciences* 7, 1 (2017), 12.
- [14] Sascha Fahl, Marian Harbach, Yasemin Acar, and Matthew Smith. 2013. On the Ecological Validity of a Password Study. In *Proc. SOUPS*. 13:1–13:13.
- [15] Adrienne Porter Felt, Alex Ainslie, Robert W. Reeder, Sunny Consolvo, Somas Thyagaraja, Alan Bettis, Helen Harris, and Jeff Grimes. 2015. Improving SSL warnings: Comprehension and Adherence. In *Proc. CHI*. 2893–2902.
- [16] Dinei Florencio and Cormac Herley. 2007. A Large-scale Study of Web Password Habits. In *Proc. WWW*. 657–666.
- [17] Dinei Florêncio, Cormac Herley, and Paul C. van Oorschot. 2014. Password Portfolios and the Finite-Effort User: Sustainably Managing Large Numbers of Accounts. In *Proc. USENIX Security Symposium*. 575–590.
- [18] Dinei Florêncio, Cormac Herley, and Paul C. van Oorschot. 2016. Pushing on String: The "Don't Care" Region of Password Strength. *Commun. ACM* 59, 11 (Oct. 2016), 66–74.
- [19] David Mandell Freeman, Sakshi Jain, Markus Dürmuth, Battista Biggio, and Giorgio Giacinto. 2016. Who Are You? A Statistical Approach to Measuring User Authenticity. In *Proc. NDSS*.
- [20] Shirley Gaw and Edward W. Felten. 2006. Password Management Strategies for Online Accounts. In *Proc. SOUPS*. 44–55.
- [21] Maximilian Golla, Theodor Schnitzler, and Markus Dürmuth. 2018. "Will Any Password Do?" Exploring Rate-Limiting on the Web. In *Proc. WAT*.
- [22] Dan Goodin. 2012. Why Passwords Have Never Been Weaker—and Crackers Have Never Been Stronger. <https://arstechnica.com/information-technology/2012/08/passwords-under-assault/>.
- [23] Google. 2010. Detecting Suspicious Account Activity. <https://security.googleblog.com/2010/03/detecting-suspicious-account-activity.html>.
- [24] Google, Inc. 2018. 2-Step Verification. <https://www.google.com/landing/2step/>.
- [25] Jeremi M. Gosney. 2017. Nvidia GTX 1080 Ti Hashcat Benchmarks. <https://gist.github.com/epixoi/ace60d09981be0954fdd35005051505>.
- [26] Weili Han, Zhigong Li, Minyue Ni, Guofei Gu, and Wenyuan Xu. 2018. Shadow Attacks Based on Password Reuses: A Quantitative Empirical Analysis. *IEEE Transactions on Dependable and Secure Computing* 15, 2 (April 2018), 309–320.
- [27] Cormac Herley and Paul C. van Oorschot. 2012. A Research Agenda Acknowledging the Persistence of Passwords. *IEEE Security & Privacy Magazine* 10, 1 (Jan. 2012), 28–36.
- [28] Karen Holtzblatt and Hugh Beyer. 2016. *Contextual Design* (second ed.). Elsevier.
- [29] Jun Ho Huh, Hyoungshick Kim, Swathi S.V.P. Rayala, Rakesh B. Bobba, and Konstantin Beznosov. 2017. I'm Too Busy to Reset My LinkedIn Password: On the Effectiveness of Password Reset Emails. In *Proc. CHI*. 387–391.
- [30] Troy Hunt. 2017. Password Reuse, Credential Stuffing and Another Billion Records in Have I Been Pwned? <https://www.troyhunt.com/password-reuse-credential-stuffing-and-another-1-billion-records-in-have-i-been-pwned/>.
- [31] Troy Hunt. 2018. Have I Been Pwned? Check If Your Email Has Been Compromised in a Data Breach. <https://haveibeenpwned.com>.
- [32] David Jaeger, Chris Pelchen, Hendrik Graupner, Feng Cheng, and Christoph Meinel. 2016. Analysis of Publicly Leaked Credentials and the Long Story of Password (Re-)use. In *Proc. PASSWORDS*.
- [33] Alexander Jenkins, Murugan Anandarajan, and Rob D'Ovidio. 2014. 'All that Glitters is not Gold': The Role of Impression Management in Data Breach Notification. *Western Journal of Communication* 78, 3 (Jan. 2014), 337–357.
- [34] Saranga Komanduri, Richard Shay, Patrick Gage Kelley, Michelle L. Mazurek, Lujo Bauer, Nicolas Christin, Lorrie Faith Cranor, and Serge Egelman. 2011. Of Passwords and People: Measuring the Effect of Password-composition Policies. In *Proc. CHI*. 2595–2604.
- [35] Frauke Kreuter, Stanley Presser, and Roger Tourangeau. 2008. Social Desirability Bias in CATI, IVR, and Web Surveys: The Effects of Mode and Question Sensitivity. *Public Opinion Quarterly* 72, 5 (2008), 847–865.
- [36] Jon A. Krosnick. 1999. Survey Research. *Annual Review of Psychology* 50, 1 (1999), 537–567.
- [37] Zhiwei Li, Warren He, Devdatta Akhawe, and Dawn Song. 2014. The Emperor's New Password Manager: Security Analysis of Web-based Password Managers. In *Proc. USENIX Security Symposium*. 465–479.
- [38] Deborah Logan. 2015. British Airways Among Latest Breaches. *Network Security* 2015, 4 (April 2015), 2–20.
- [39] Chris Long. 2014. Keeping Passwords Secure. <https://www.facebook.com/notes/protect-the-graph/keeping-passwords-secure/1519937431579736/>.
- [40] William Melicher, Blase Ur, Sean M. Segreti, Saranga Komanduri, Lujo Bauer, Nicolas Christin, and Lorrie Faith Cranor. 2016. Fast, Lean, and Accurate: Modeling Password Guessability Using Neural Networks. In *Proc. USENIX Security Symposium*. 175–191.
- [41] Grzegorz Milka. 2018. Anatomy of Account Takeover. In *Proc. Enigma*.
- [42] Saif M. Mohammad and Peter D. Turney. 2013. Crowdsourcing a Word-Emotion Association Lexicon. *Computational Intelligence* 29, 3 (2013), 436–465.

- [43] Jeremiah Onalapo, Enrico Mariconti, and Gianluca Stringhini. 2016. What Happens After You Are Pwned: Understanding the Use of Leaked Webmail Credentials in the Wild. In *Proc. IMC*. 65–79.
- [44] Sarah Pearman, Jeremy Thomas, Pardis Emami Naeini, Hana Habib, Lujo Bauer, Nicolas Christin, Lorrie Faith Cranor, Serge Egelman, and Alain Forget. 2017. Let's Go in for a Closer Look: Observing Passwords in Their Natural Habitat. In *Proc. CCS*. 295–310.
- [45] Nicole Perlroth. 2017. All 3 Billion Yahoo Accounts Were Affected by 2013 Attack. <https://www.nytimes.com/2017/10/03/technology/yahoo-hack-3-billion-users.html>.
- [46] Prabaharan Poornachandran, M. Nithun, Soumajit Pal, Aravind Ashok, and Aravid Ajayan. 2015. Password Reuse Behavior: How Massive Online Data Breaches Impacts Personal Data in Web. In *Proc. ICICSE*. 199–210.
- [47] Emilee Rader, Rick Wash, and Brandon Brooks. 2012. Stories As Informal Lessons About Security. In *Proc. SOUPS*. 6:1–6:17.
- [48] Elissa M. Redmiles, Everest Liu, and Michelle L. Mazurek. 2017. You Want Me To Do What? A Design Study of Two-Factor Authentication Messages. In *Proc. WAY*.
- [49] Elissa M. Redmiles, Amelia R. Malone, and Michelle L. Mazurek. 2016. I Think They're Trying to Tell Me Something: Advice Sources and Selection for Digital Security. In *Proc. IEEE S&P*. 272–288.
- [50] Elissa M. Redmiles, Ziyun Zhu, Sean Kross, Dhruv Kuchhal, Tudor Dumitras, and Michelle L. Mazurek. 2018. Asking for a Friend: Evaluating Response Biases in Security User Studies. In *Proc. CCS*.
- [51] Shape Security. 2017. 2017 Credential Spill Report. <http://info.shapesecurity.com/2017-Credential-Spill-Report-w.html>.
- [52] Richard Shay, Saranga Komanduri, Patrick Gage Kelley, Pedro Giovanni Leon, Michelle L. Mazurek, Lujo Bauer, Nicolas Christin, and Lorrie Faith Cranor. 2010. Encountering Stronger Password Requirements: User Attitudes and Behaviors. In *Proc. SOUPS*. 2:1–2:20.
- [53] Steve Sheng, Bryant Magnien, Ponnuram Kumaraguru, Alessandro Acquisti, Lorrie Faith Cranor, Jason Hong, and Elizabeth Nunge. 2007. Anti-Phishing Phil: The Design and Evaluation of a Game That Teaches People Not to Fall for Phish. In *Proc. SOUPS*. 88–99.
- [54] Elizabeth Stobert and Robert Biddle. 2014. The Password Life Cycle: User Behaviour in Managing Passwords. In *Proc. SOUPS*. 243–255.
- [55] San-Tsai Sun, Eric Pospisil, Ildar Muslukhov, Nuray Dindar, Kirstie Hawkey, and Konstantin Beznosov. 2011. What Makes Users Refuse Web Single Sign-on?: An Empirical Investigation of OpenID. In *Proc. SOUPS*. 4:1–4:20.
- [56] Kurt Thomas, Frank Li, Chris Grier, and Vern Paxson. 2014. Consequences of Connectivity: Characterizing Account Hijacking on Twitter. In *Proc. CCS*. 489–500.
- [57] Kurt Thomas, Frank Li, Ali Zand, Jacob Barrett, Juri Ranieri, Luca Invernizzi, Yarik Markov, Oxana Comanescu, Vijay Eranti, Angelika Moscicki, Daniel Margolis, Vern Paxson, and Elie Bursztein. 2017. Data Breaches, Phishing, or Malware? Understanding the Risks of Stolen Credentials. In *Proc. CCS*. 1421–1434.
- [58] Roger Tourangeau and Ting Yan. 2007. Sensitive Questions in Surveys. *Psychological Bulletin* 133, 5 (2007), 859.
- [59] Blase Ur, Felicia Alfieri, Maung Aung, Lujo Bauer, Nicolas Christin, Jessica Colnago, Lorrie Faith Cranor, Henry Dixon, Pardis Emami Naeini, Hana Habib, et al. 2017. Design and Evaluation of a Data-Driven Password Meter. In *Proc. CHI*. 3775–3786.
- [60] Blase Ur, Fumiko Noma, Jonathan Bees, Sean M. Segreti, Richard Shay, Lujo Bauer, Nicolas Christin, and Lorrie Faith Cranor. 2015. "I Added '!' at the End to Make It Secure": Observing Password Creation in the Lab. In *Proc. SOUPS*. 123–140.
- [61] Blase Ur, Sean M. Segreti, Lujo Bauer, Nicolas Christin, Lorrie Faith Cranor, Saranga Komanduri, Darya Kurilova, Michelle L. Mazurek, William Melicher, and Richard Shay. 2015. Measuring Real-World Accuracies and Biases in Modeling Password Guessability. In *Proc. USENIX Security Symposium*. 463–481.
- [62] Chun Wang, Steve T.K. Jan, Hang Hu, Douglas Bossart, and Gang Wang. 2018. The Next Domino to Fall: Empirical Analysis of User Passwords across Online Services. In *Proc. CODASPY*. 196–203.
- [63] Ding Wang, Zijian Zhang, Ping Wang, Jeff Yan, and Xinyi Huang. 2016. Targeted Online Password Guessing: An Underestimated Threat. In *Proc. CCS*. 1242–1254.
- [64] Ke Coby Wang and Michael K. Reiter. 2018. How to End Password Reuse on the Web. *CoRR abs/1805.00566* (May 2018), 1–16.
- [65] Rick Wash, Emilee Radar, Ruthie Berman, and Zac Wellmer. 2016. Understanding Password Choices: How Frequently Entered Passwords are Re-used Across Websites. In *Proc. SOUPS*. 175–188.
- [66] Yixin Zou, Abraham H Mhaidli, Austin McCall, and Florian Schaub. 2018. "I've Got Nothing to Lose": Consumers' Risk Perceptions and Protective Actions after the Equifax Data Breach. In *Proc. SOUPS*.

A SURVEY INSTRUMENTS

A.1 Study 1 Survey Instrument

Because the notifications were delivered through different channels, we inserted wording appropriate to the notification. For example, for LinkedIn, we used the following:

Prompt: Imagine that you receive, through email from AcmeCo,

VerbPrompt: receiving this notification through email from AcmeCo

NounPrompt: this notification through email from AcmeCo

PastTensePrompt: received this notification through email from AcmeCo

Introduction In the following survey, you will be asked to imagine that your name is Jo Doe. You have an online account with a major company called AcmeCo and can access your account through both a website and a mobile application. Imagine that this account is important to you, and that it is like other accounts you may have, such as for email, banking, or social media. This survey should take approximately 15 minutes to complete.

Prompt the following notification: (screenshot)

In your own words, please describe what this notification is telling you.

In your own words, please describe all of the factors that may have caused you to receive this notification.

Please list three feelings you might have after receiving this notification.

Please list three actions you might take after receiving this notification.

The first feeling you listed was *. Please explain why you might feel this way.

The second feeling you listed was *. Please explain why you might feel this way.

The third feeling you listed was *. Please explain why you might feel this way.

The first action you listed was *. Please explain why you might take this action.

The second action you listed was *. Please explain why you might take this action.

The first third you listed was *. Please explain why you might take this action.

I feel that *NounPrompt* explained to me how to resolve the situation.

☐ Strongly agree ☐ Agree ☐ Neither agree nor disagree ☐ Disagree ☐ Strongly disagree ☐ Don't know
Why?

Notifications can be received in many different ways, such as through email, on a webpage, or in a mobile app. Please select the answer choice that most closely matches how you feel about the following statement:

I feel that *NounPrompt* uses the appropriate method of contacting me.

☐ Strongly agree ☐ Agree ☐ Neither agree nor disagree ☐ Disagree ☐ Strongly disagree ☐ Don't know
Why?

I feel that ignoring *NounPrompt* would not have any consequences.

☐ Strongly agree ☐ Agree ☐ Neither agree nor disagree ☐ Disagree ☐ Strongly disagree ☐ Don't know
Why?

For me, taking action in response to *VerbPrompt* would be a

☐ Very high priority ☐ High priority ☐ Medium priority ☐ Low priority ☐ Not a priority ☐ Don't know
Why?

I would feel _____ about *VerbPrompt*.

☐ Extremely concerned ☐ Moderately concerned ☐ Somewhat concerned ☐ Slightly concerned ☐ Not at all concerned ☐ Don't know
Why?

I would expect real companies to send notifications like this one when necessary.

☐ Strongly agree ☐ Agree ☐ Neither agree nor disagree ☐ Disagree ☐ Strongly disagree ☐ Don't know
Why?

I have received notifications similar to this one in the past.

☐ Never ☐ A few times ☐ Many times ☐ Don't know

Briefly describe the notifications, if any, that you have received. Please include the context in which you received the notifications and who sent them.

A.2 Study 2 Survey Instrument

Introduction In the following survey, you will be asked to imagine that your name is Jo Doe. You have an online account with a major company called AcmeCo and can access your account through both a website and a mobile application. Imagine that this account is important to you, and that it is like other accounts you may have, such as for email, banking, or social media. This survey should take approximately 15 minutes to complete.

(Show notification and explain delivery method.)

Initial Questions

In your own words, please describe what this notification is telling you.

What may have caused you to receive this notification? Please check all that apply.

- ☐ Someone hacked your AcmeCo account. AcmeCo noticed suspicious activity, such as logins from an unexpected location, a new device being used, or multiple unsuccessful logins.
- ☐ Your AcmeCo account has not been hacked. Instead, you simply logged in from a new location or device, or accidentally entered the wrong password too many times.
- ☐ AcmeCo was hacked.
- ☐ A company unrelated to AcmeCo was hacked.
- ☐ You reused the same or similar passwords for multiple online accounts.
- ☐ Someone is trying to gain unauthorized access to your account by sending this email.
- ☐ AcmeCo conducts regular security checks and this is just a standard security notification.
- ☐ You have a weak password for your AcmeCo account.
- ☐ AcmeCo sent this by mistake.
- ☐ You went to a malicious website or downloaded malicious software.
- ☐ AcmeCo requires you to regularly change your password (e. g., every 90 days).
- ☐ Don't know

Password Change Actions

If you received this notification about an online account you had with a real company, which of the following best describes what you would do about passwords for that account?

- ☐ I would keep my password the same.
- ☐ I would change my password.
- ☐ Don't know

Why?

If "I would change my password" is selected.

What would you use for your new password on that account?

- ☐ Something related to the old password, but a few characters different.
- ☐ Something completely unrelated to the old password.
- ☐ A password that I already use for other accounts.
- ☐ A password generated by a password manager or browser.
- ☐ Other _____

If "I would change my password" is selected.

How would you try to remember your new password for that account? Select all that apply.

- ☐ Write it down (e. g., in a diary, on a sticky note).
- ☐ Use a password manager.
- ☐ Just try to remember it.
- ☐ Save it on my computer (e. g., in a document).
- ☐ Save it on my phone (e. g., in a note).
- ☐ Other _____

If you received this notification about an online account you had with a real company, which of the following best describes what you would do about passwords on other accounts? Please select all that apply.

- ☐ I would change all of my passwords I have on other accounts.
- ☐ I would change my passwords only for other accounts where I use the same password.
- ☐ I would change my passwords only for other accounts where I use similar passwords.
- ☐ I would change my passwords only for really important accounts (e. g., bank account).
- ☐ I would keep my passwords the same.
- ☐ Don't know.

Why? _____

If any of the first four from above were selected.

What would you use for your new password(s) on those other accounts?

- ☐ Something related to the old password, but a few characters different.
- ☐ Something completely unrelated to the old password.
- ☐ A password that I already use for other accounts.
- ☐ A password generated by a password manager or browser.
- ☐ Other _____

If any of the first four from above were selected.

How would you try to remember your new password(s) for those other accounts?

Select all that apply.

- ☐ Write it down on paper (e. g., in a diary, on a sticky note).
- ☐ Use a password manager.
- ☐ Just try to remember it.
- ☐ Save it on my computer (e. g., in a document).
- ☐ Save it on my phone (e. g., in a note).
- ☐ Other _____

People have different reactions and responses to notifications about their online accounts. If you received this notification about an online account you had with a real company, how likely would you be to take the following actions?

(Answer choices for each) ☐ Very Unlikely ☐ Unlikely ☐ Neither likely nor unlikely ☐ Likely ☐ Very Likely ☐ Don't Know

- ☐ Enable Two-Factor Authentication.
- ☐ Use a password manager.
- ☐ Update my security questions.
- ☐ Review my recent account activity.
- ☐ Leave my password as-is.

- ☐ Commit to change my password more frequently in the future.
- ☐ Sign up for an account with a company offering identity theft protection.
- ☐ Update the software my devices more regularly.
- ☐ Add a/Change my current password, PIN, pattern, fingerprint, etc. to lock my phone.
- ☐ Add a/Change my current password to lock my computer.

There are many different actions that people could take in response to notifications about their online accounts. Please select the answer choice that most closely matches how you feel about the following statements:

If I received this notification about an online account I had with a real company, it would improve my account security if I...

(Answer choices for each) ☐ Strongly agree ☐ Agree ☐ Neither agree nor disagree ☐ Disagree ☐ Strongly disagree ☐ Don't Know

- ☐ ... enabled Two-Factor Authentication.
- ☐ ... used a password manager.
- ☐ ... changed my password for this account to a new password that is a modification (changing a few characters) of the old one.
- ☐ ... changed my password for this account to a completely new password unrelated to the old one.
- ☐ ... changed my password for this account to a password I use for another online account.
- ☐ ... used unique passwords for each of my online accounts.
- ☐ ... changed all of my similar passwords on other online accounts to one new password.
- ☐ ... updated my security questions.
- ☐ ... reviewed my recent activity.
- ☐ ... left my password as-is.
- ☐ ... committed to change my password more frequently in the future.
- ☐ ... signed up for an account with a company offering identity theft protection.
- ☐ ... updated the software on my devices more regularly.
- ☐ ... added a/changed my current password, PIN, pattern, fingerprint, etc. to lock my phone.
- ☐ ... added a/changed my current password to lock my computer.

Notifications can be received in many different ways, such as through email, on a webpage, or in a mobile app. Please select the answer choice that most closely matches how you feel about the following statement: I feel that this notification uses the appropriate method of contacting me.

☐ Strongly Agree ☐ Agree ☐ Neither agree nor disagree ☐ Disagree ☐ Strongly disagree ☐ Don't know

If you were to receive a similar notification about an online account you had with a real company, how would you want to be contacted? Please select all that apply.

☐ Email ☐ Pop-up notification on mobile, such as if you received an SMS ☐ Text message ☐ Website on desktop or mobile browser ☐ In the mobile app ☐ Phone call ☐ Physical mail ☐ Other _____

Given that I received *NounPrompt*, I would probably see this notification:

☐ Within 3 hours ☐ Within 24 hours ☐ Within 3 days ☐ Within a week ☐ After a week ☐ Never ☐ Don't know

After receiving this notification, I would probably take action:

☐ Within 3 hours ☐ Within 24 hours ☐ Within 3 days ☐ Within a week ☐ After a week ☐ Never ☐ Don't know

I would expect real companies to send notifications like this one when necessary.

☐ Strongly Agree ☐ Agree ☐ Neither agree nor disagree ☐ Disagree ☐ Strongly disagree ☐ Don't know

If I received this notification about an online account I had with a real company, I would believe that this was an official notification sent by that company.

☐ Strongly Agree ☐ Agree ☐ Neither agree nor disagree ☐ Disagree ☐ Strongly disagree ☐ Don't know

I feel that ignoring this notification would not have any consequences.

☐ Strongly Agree ☐ Agree ☐ Neither agree nor disagree ☐ Disagree ☐ Strongly disagree ☐ Don't know

For me, taking action in response to *VerbPrompt* would be a:

☐ Very high priority ☐ High priority ☐ Medium priority ☐ Low priority ☐ Not a priority ☐ Don't know

If I received this notification about an online account I had with a real company, I would feel grateful.

☐ Strongly Agree ☐ Agree ☐ Neither agree nor disagree ☐ Disagree ☐ Strongly disagree ☐ Don't know

This notification adequately explains what is going on with my online account.

☐ Strongly Agree ☐ Agree ☐ Neither agree nor disagree ☐ Disagree ☐ Strongly disagree ☐ Don't know

If I received this notification about an online account I had with a real company, I wouldn't know why I received this notification.

☐ Strongly Agree ☐ Agree ☐ Neither agree nor disagree ☐ Disagree ☐ Strongly disagree ☐ Don't know

I feel that *NounPrompt* explained to me how to resolve the situation.

☐ Strongly Agree ☐ Agree ☐ Neither agree nor disagree ☐ Disagree ☐ Strongly disagree ☐ Don't know

People may have many different responses to receiving notifications about their online accounts. Please select the answer choice that most closely matches how you feel about the following statement: After receiving this notification, my trust in AcmeCo would:

☐ Significantly increase ☐ Increase ☐ Neither increase nor decrease ☐ Decrease ☐ Significantly decrease ☐ Don't know

Why?

To your knowledge, has anyone ever gained unauthorized access to one of your online accounts?
☐ Yes ☐ No ☐ Don't know

If yes selected. Who do you think accessed your online account? Please select all that apply.
☐ Someone you know personally ☐ Someone you don't know personally ☐ Don't know

If yes selected. Please describe what happened.

Do any of your accounts require you to change your password regularly (e.g., every 90 days)?
☐ Yes ☐ No ☐ Don't know

If yes selected. Please describe how you were informed of this regular password change policy

Have you ever been notified that your information was exposed in a data breach?
☐ Yes ☐ No ☐ Don't know

If yes selected. Please describe how you found out and what happened.

With what gender do you identify?

☐ Male ☐ Female ☐ Non-binary ☐ Other _____ ☐ Prefer not to say

What is your age?

☐ 18-24 ☐ 25-34 ☐ 35-44 ☐ 45-54 ☐ 55-64 ☐ 65-74 ☐ 75 or older ☐ Prefer not to say

What is the highest degree or level of school you have completed?

☐ Some high school ☐ High school ☐ Some college ☐ Trade, technical, or vocational training ☐ Associate's Degree ☐ Bachelor's Degree ☐ Master's Degree ☐ Professional degree ☐ Doctorate ☐ Prefer not to say

Which of the following best describes your educational background or job field?

☐ I have an education in, or work in, the field of computer science, computer engineering or IT.
☐ I do not have an education in, nor do I work in, the field of computer science, computer engineering or IT.
☐ Prefer not to say

(Optional) Do you have any final thoughts or questions about today's study?

B SCREENSHOTS OF NOTIFICATIONS

B.1 Study 1 Notifications

The six notifications used for Study 1 are shown below. Notifications are identified by their original sender, although all notifications were rebranded as AcmeCo for the purposes of the survey.

Study 1 Facebook notification:

Study 1 Google email notification:

Study 1 Google red bar notification:

Study 1 Instagram notification:

Study 1 **LinkedIn** notification: Study 1 **Netflix** notification:

B.2 Study 2 Notifications

Two of the three variants of the model notification's delivery medium — mobile and inApp — used for Study 2 are shown below. The third variant — email — was provided in the body of the paper. The text variations of the notifications are shown in Figure 4 with the corresponding changes in Table 2.

Study 2 *model-{mobile}* notification: Study 2 *model-{inApp}* notification: