The Effect of Repeated Login Prompts on Phishing Susceptibility

Peter Snyder and Chris Kanich

Department of Computer Science
University of Illinois at Chicago
{psnyde2,ckanich}@uic.edu

Michael K. Reiter

Department of Computer Science
University of North Carolina at Chapel Hill
reiter@cs.unc.edu

Abstract

Background. Understanding the human aspects of phishing susceptibility is an important component in building effective defenses. People type passwords so often that it is possible that this act makes each individual password less safe from phishing attacks.

Aim. This study investigated whether the act of re-authenticating to password-based login forms causes users to become less vigilant toward impostor sites, thus making them more susceptible to phishing attacks. Our goal was to determine whether users who type their passwords more often are more susceptible to phishing than users who type their passwords less often. If so, this result could lead to theoretically well-grounded best practices regarding login-session length limits and re-authentication practices.

Method. We built a custom browser extension which logs password entry events and has the capability of shortening session times for a treatment group of users. We recruited subjects from our local campus population, and had them run the extension for two months. After this time, we conducted a synthetic phishing attack on all research subjects, followed by a debriefing. Our research protocol was approved by the University's IRB.

Results. We failed to reject the null hypothesis. We found that login frequency has no noticeable effect on phishing susceptibility. Our high phishing success rate of 39.3% was likely a leading factor in this result.

Conclusions. This study confirmed prior research showing exceedingly high phishing success rates. We also observed that recruiting only in-person and campus-affiliated users greatly reduced our subject pool, and that the extension-based investigation method, while promising, faces significant challenges itself due to deployed extension-based malware defenses.

1 Introduction and Motivation

Of all cybersecurity attacks, phishing is perhaps most intimately tied to user decisions and behavior, rather than technical weaknesses of the platform on which it is perpetrated. Despite numerous studies both aimed at better understanding why and how people fall for phishing attacks [3, 4] and new systems to detect the sites themselves and protect users [9], the problem continues [6]. Thus, fully understanding the causes and effects of the phenomenon is a crucial component of successful defenses.

Fundamental to phishing as an attack is the user credential, which is most often a password. Passwords have been the subject of intense investigation along many dimensions; e.g., Bonneau et al. consider many of these dimensions and present a framework for comparing passwords with other authentication schemes [1].

Beyond the weaknesses of passwords discussed by Bonneau et al., the difficulties of password entry are exacerbated by the frequency with which users are asked to re-authenticate to sites. In the Internet's infancy, personal computers were commonly shared among different users, but software to enable efficient sharing via different profiles did not yet exist. Thus, automatically logging users out after some relatively short period of time became the de facto default security posture with respect to authenticated sessions. In the present day, however, browsing often takes place on single-user devices like smartphones, and even when not, streamlined browser profiles make separating different user accounts easy.

Several services like Facebook and Google have thus adopted a strategy where sessions remain logged in for an indeterminate amount of time.[1] However, to our knowledge there is no scholarship or best-practices document available which provides a breakdown of how successful this practice is, and it has not seen widespread adoption

---

[1]It remains to be seen whether such infrequent reauthentication might itself be harmful: if passwords are so rarely re-entered, the user runs the risk of forgetting it simply because they never use it.

on the rest of the web.

With respect to phishing defenses, this might indeed be a very good one: the longer a session remains in place on a given device, the less common the sight of the login screen (or even any login screen across all sites) is to the individual user. If login screens are less common, we hypothesize that this will cause the user to be more alert when logging in to services. While this may be the reasoning behind the long-lived sessions used by large Internet companies, being able to reproduce this effect in open scholarship would be an important step towards convincing more software authors and site owners to make these longer-lived sessions the default.

Specifically, financial services websites very often keep their session lengths limited to hours or even minutes. On one hand, this choice is an entirely reasonable: an online banking session left logged in on a shared or stolen device can cause immense damage for an individual. Conversely, being prompted for these passwords continually has the very real potential to make those users more susceptible to phishing attacks. When considering the threat models of phishing attacks and physical device takeover, the former can be perpetrated by anyone on the Internet rather than anyone with access to the device, potentially making the phishing concern far greater than the session length concern.

We claim that password re-entry frequency is not only a usability issue, but is also a security issue for all password-based login systems. Our hypothesis is that the more users are asked to authenticate with websites, the more they will experience security fatigue, and as a result become more susceptible to phishing attacks.

We further hypothesize that this effect is not local to the website. In other words, we expect that when users are prompted to frequently authenticate, they become security fatigued in general, on all websites they visit, not only on the specific sites that are prompting users for their credentials.

This study was built to test these hypothesises. We devised a methodology which includes a control group whose login frequency is unmodified, and a treatment group who have the length of their login session shortened, necessitating additional logins to simulate sites with short session timeouts. Unfortunately, we did not find statistical significance in our study, but we were able to replicate important results in phishing susceptibility and to find new but unremarkable results related to well crafted spear phishing attacks.

2 Methodology

This work was carried out in five stages, starting with the development of the software used to measure and manipulate the web browsing experience of the control group, and ending with informing the test participants about the experiment they participated in. This section describes each of these five sections in chronological order.

2.1 Web Browser Extension Development

Browser extensions are pieces of software that are installed in commodity web browsers to measure or modify the user's browsing experience. Though all recent versions of popular web browsers support the ability to write and install extensions, we limited our study to people using Firefox and Chrome.

Our study included two different extensions, one version for the control group that took measurements of the user's browsing activities, and another version for the experiment group that took the same measurements but also modified the browser to induce the user to (re)authenticate with popular websites more frequently than they normally would. Each of theses versions of the extension are described in greater detail in the following subsections.

When the user installed the extension in their browser, she was prompted to enter her email address. The extension then generated a random identifier for the user. The researchers never tied these two identifiers together; we were never able to associate a person's random identifier with her email address. These identifiers were used to identify users during different parts of the study. Finally, on installing the extension, users were randomly assigned to either the control or experiment group, with equal chance.

Throughout the experiment the extensions reported statistics to a central recording server. When the extension reported non-sensitive information to the recording server (such as how long a user has been participating in the study), the extension identified the user by her email address. Likewise, whenever the extension reported sensitive information to the server, the data was tied to the user's random identifier. This allowed us to track which users were still participating in the study, without tying any sensitive information to the participant.

2.1.1 Control Group Extension

The control group version of the extension did not modify the browsing experience at all; it only recorded information about the user's browsing activities. Each of these data points is described below.

User Participation: In order to determine which participants were staying active in the study (and to be able to remind participants if they were not staying active) the extension reported each active user's email address to the recording server once a day.

Pages Visited: In order to determine how regularly different study participants used the web, the extension recorded how many pages each participant visited each hour. The extension did not record *which* pages were visited, only a count of how many each hour. This information was periodically reported back to the recording server with the user's random identifier.

Passwords Entered: The extension also recorded how often users entered passwords on the web. Each time the user entered her password into a password field on the web, the extension recorded the URL of the page (used to determine what types of pages the user trusted with her password), a salted hash of the password (in order to determine if the user reused passwords, without revealing their passwords) and a NIST entropy measure of the password [5] (as a measure of password strength). These values were also periodically reported to the recording server, along with the user's random identifier.

2.1.2 Experiment Group Extension

The extension performed differently for users in the experiment group. In addition to recording the information discussed in Section 2.1.1, the experiment version of the extension also modified users' browsing sessions to cause them to need to authenticate with websites more frequently then they otherwise would.

- https://www.reddit.com

- https://www.facebook.com

- https://www.google.com

- https://mail.google.com

- https://www.tumblr.com

- https://twitter.com

- https://mail.yahoo.com

- https://www.pinterest.com

Figure 1: URLs of sites that participants in the experiment group were induced to reauthenticate with more frequently

We first selected eight popular sites where users needed to login to use the sites primary functionality, listed in Figure 1. When a user logs into one of these sites, the site sets a cookie in the user's browser. This cookie usually lasts for a long time: days, weeks or months. The site uses this information to identify the user to the site, so that the site knows who the user is and does not ask the user to re-login.

The experiment group version of the extension shortened the life span of these cookies to expire on average in 48 hours (some random variation, between plus-and-minus twelve-hours, was added into the cookie expiration times, to make it more difficult for experiment-group members to detect the manipulation). The net result of editing the expiration dates of the cookies is that users would need to re-login to these sites approximately every two days, instead of once a week or once a month.

2.2 Recruitment

Recruitment was conducted through university mailing lists of students, faculty and staff. A person was eligible to participate in the study if she 1) did most of her web browsing on a computer she could install software on; 2) was a student, faculty or staff member; 3) used Chrome or Firefox as her main browser; 4) spent some time on social media sites most days; 5) was at least 18 years old; and 6) was not currently incarcerated.

Participants were required to participate in the study for two months, during which they needed to use the computer with the extension installed at least once every three days. They were told that the study was about "measuring safe browsing practices online", but were not given any further detail about the purpose of the study. They were told that the extension would take anonymous measurements of their browsing habits, and that it would not harm their computer.

Participants were not told that some participants would have their cookies removed earlier than normal, and that they would need to log into sites more often. This deception was conducted with the review and permission of the IRB of the University of Illinois at Chicago.

2.3 Study Participation

Each participant who agreed to the above conditions arranged to meet with a member of the study to guide her through the process of installing the extension on her computer, confirm that she met the eligibility requirements, read and sign a consent form, and document her agreement to participate in the study.

Each participant was offered $30 in Amazon.com gift cards as compensation for her participation, receiving $15 on entering the study and the remaining $15 at the end of the study, if they followed all of the agreed-to terms.

During the study, participants operated their computers as normal, and carried out their typical browsing behaviors. We regularly checked to make sure that the extensions were functioning correctly and that the study participants were still using their browsers at least once every three days. In a few cases, we noticed that a study participant was not using her browser in accordance with

the study's terms. In such cases we contacted the participant by email to see if there was a technical problem, and in a few cases we removed participants from the study who were not able to meet the study's requirements.

2.4 Simulated Phishing Attack

At the end of the two-month study period, we sent an email out, telling the participants that the end of the study was approaching, and that they would need to make an appointment to have the software removed from their computers and be debriefed from the study.

We separately sent all email participants a fake phishing email. The email told all students that they should click on a link in the email to log into their university accounts, in order to complete a brief survey and qualify for the remaining $15 Amazon gift certificate. The email was constructed to only include content that an attacker would have access to and be able to forge.

Notably, the message was sent from a non-university account, which had never been used to interact with the study participants prior. Additionally, the link in the message that participants were asked to click on—and which claimed to be a link to the university's sign-on system—linked to a new domain that was not university owned and which had never been provided to participants before. Finally, when participants clicked on the link and were taken to the false, phishing, version of the university's sign-on system, they were asked to log into a domain that was also not university owned or affiliated with the school.

The fake, phishing version of the university's login page we constructed kept track of which study participants visited the page, how long they stayed on the page, if they entered a user name and password, and if they they submitted the form to attempt to log in.

Each participant who submitted values in the login form was asked to complete a survey. The survey attempted to assess whether she took standard precautions before submitting her university credentials by asking, among other questions, if she checked the URL before entering her user name and password, and if she noticed anything abnormal about the login page's URL.

To protect the participants, we did not record or transmit the entered user name or password over the network; we only recorded how many participants interacted with the page in the same manner one would interact with the true university login system, and if they trusted the false version of the page with their account credentials.

In order to avoid having study participants influence one another or inform each other about the deception in the study, we took care to not reveal the deception for one week, until all participants had a chance to receive and respond to the fake phishing email. If users did click on the

link in the sent email, and submitted their credentials to our fake-sign-on page, the messages and web pages they received appeared identical to those they would receive from the true university login system.

2.5 Debriefing

One week after the fake phishing email was sent to study participants, all study participants were sent another email, this time from the university email account that had been corresponding with them throughout the study. The email asked participants to schedule a debriefing meeting with the research assistant. At this meeting, each participant was told about the true purpose of the study, given the chance to ask about the purpose, methods, or outcomes of the research, and provided with the remaining $15 in Amazon credit.

3 Results

We were able to recruit 101 study participants, 89 of which completed the study[2]. Of those who completed the study, 43 were in the control group, and 46 in the experiment group.

3.1 Phishing Susceptibility

Of the 43 participants in the control group, 17 (39.5%) clicked on the link in the phishing email, or otherwise visited the phishing page. All 17 of these control-group members entered some value into the password field on the fake university-login page and submitted the form. In the experiment group, 19 of the 49 (38.8%) participants visited the phishing page, and 18 of them entered some value into the password field and submitted the form. We were not able to find a statistically significant difference between the control and treatment groups.

Participants who submitted the login form were taken to a survey that asked about their participation in the study, if they encountered any technical problems with the browser extension, and if they would like to be notified of the study results when available. Most relevant to the question of phishing susceptibility, participants were asked if they noticed that the URL for the fake university-login form was different from the URL where they normally logged in. Of the 17 members of the control group who completed the survey, 5 (29.4%) stated that they noticed the URL was different, versus 6 of 18 experiment-group members (33.3%) who noticed that the URL was different.

	Control	Experiment
Mean	32.88	35.56
Min	0	0
Max	407	546
St Dev	88.08	64.06

Table 1: Basic statistics on the number of passwords entered by users in the control and experiment groups.

3.2 Password Entry

Finally, the data gathered during this work allowed us to make some measurements about how many passwords participants entered during the two month study, and on how many different domains they submitted passwords. For the eight domains we affected in the study, participants entered on average 34.2 passwords, with our extension inducing users in the test group authenticating more often than users in the control group. Users in the control group entered on average 32.9 passwords on the eight watched sites during the study, while users in the experiment group entered 35.6 passwords to the eight relevant domains.

More broadly, users entered on average 185.74 passwords during the two month study, and submitted passwords to 28.69 domains.

4 Related Work

This work sits alongside other research establishing the effectiveness of phishing attacks as a means of stealing user credentials. Dhamija et al. [2] found that a well constructed phishing page fooled 20 out of 22 test subjects, and that this vulnerability seemed unrelated to demographic or personal characteristics, such as age, sex, or number of hours of computer use. Jagatic [8] investigated how social connections can affect the success rate of phishing attacks, and the success rate of a phishing attack went from 16% of 94 target students to 72% of 487 targeted students when the phishing message was forged to appear to be from a friend or other social contact.

Other research has established that common anti-phishing indicators in browsers do a poor job of alerting users to fraud. Schechter et al. [10] found that factors like the absence of https encryption on web pages and missing user-selected site images did not dissuade users from entering their credentials (all 27 users submitted their credentials to the phishing site in the former case, and 23 out of 25 users still did so in the latter case). Similarly, Wu et al. [12] found that even with additional anti-phishing tool-

bars and utilities installed, 10 out of the 30 participants in their study were still successfully phished. Whalen and Inkpen [11] used an eye-tracking system to determine what security indicators users viewed, and found that unless specifically prompted, users rarely looked at the browser's security indicators. Jackson et al. [7] found that web-browser users did not understand the browser's anti-phishing security warnings, and thus that they offered no protection, unless they received specific training in understanding the browser's indicators.

5 Lessons Learned

While the core experiment failed to reject the null hypothesis, several of our observations confirm previous studies and can otherwise be useful to the community performing further security-based user studies.

Confirmed very high phishing success rate. The success rate for phishing is high. 40.4% of participants who received the phishing email submitted a password to a untrusted domain, and 97.2% of participants who clicked through the email and visited the fraudulent university-login page submitted their credentials.

This result is possibly due to the priming effect of our phishing strategy (i.e., our subjects were expecting an email regarding payment). However, other factors were also likely at play, including a correctly functioning https url, a benign yet similar hostname, and an incredibly low-volume campaign such that typical defenses to prevent deliver of phishing messages would not have been triggered.

Note that we did not investigate whether any of our users had phishing defenses turned on, either directly through their browser or through additional software such as browser extensions or anti-virus programs.

Recruitment challenges. We chose to recruit through our university's email channel for mass advertisement to all faculty, staff, and students. While our university is far more diverse than average in terms of race and socio-economic status, this was still likely to be a less representative group than the general population.

Our reason for recruiting in this manner was that we wanted to ensure a standard phishing experience for the study. Everyone was required to use the same campus single-sign-on infrastructure, which we also required participants to use when selecting a time to meet with us to enter the study. We required this in-person meeting to minimize fraud and to ensure that the extension was installed correctly.

When we recruited in this manner, we had a far lower response rate than we expected, especially given the reward structure for our study. We believe that attempting to minimize participant fraud via in-person meetings was

---

[2]Seven members of the control group, and five members of the experiment group, exited the study mid-experiment.

likely not an effective use of time; other methods of filtering out fraudulent users would likely have been superior.

Extension installation challenges. The proliferation of extension-based malware made it particularly difficult to successfully install the extension. This included turning off various anti-malware features in the browser (temporarily) to successfully install the extension on users' machines.

Asking users to install an extension with such expansive permissions is a lot to ask, even in an IRB-approved, monetarily compensated study. We minimized the amount of data collected and ensured that all data was anonymized and encrypted during storage and transmission. We explained this process in layman's terms during the installation of the extension at the in-person enrollment events.

5.1 Advice for studying phishing susceptibility

For anyone wishing to attempt an experiment like this one, we believe that a few changes would raise the likelihood of observing a correlation—if one exists—between login frequency and phishing susceptibility. First, we believe that targeting more websites (or even doing so in a site-agnostic fashion) would be beneficial, as well as allowing for the collection of information about password entry more broadly.

Second, we recommend conducting further research in a way that can control for different amounts of natural (i.e., pre-test) password re-entry. The best way to control for this would likely be to recruit more research subjects. Allowing participants to sign up remotely would be a boon in this respect. However, it might filter for more technically savvy users who are able to install extensions on their own.

The effect of saved passwords or password managers would be an interesting angle to investigate. These tools associate the saved logins credentials with specific sites, and so phishing sites would not be auto-filled. Controlling for this effect would be important, as this process can drastically reduce the number of passwords typed by a user over a given amount of time.

References

[1] BONNEAU, J., HERLEY, C., VAN OORSCHOT, P. C., AND STAJANO, F. The quest to replace passwords: A framework for comparative evaluation of web authentication schemes. In *Security and Privacy (SP), 2012 IEEE Symposium on* (2012), IEEE, pp. 553–567.

[2] DHAMIJA, R., TYGAR, J. D., AND HEARST, M. Why phishing works. In *Proceedings of the SIGCHI conference on Human Factors in computing systems* (2006), ACM, pp. 581–590.

[3] DOWNS, J. S., HOLBROOK, M. B., AND CRANOR, L. F. Decision strategies and susceptibility to phishing. In *Proceedings*

of the second symposium on Usable privacy and security (2006), ACM, pp. 79–90.

[4] EGELMAN, S., CRANOR, L. F., AND HONG, J. You've been warned: an empirical study of the effectiveness of web browser phishing warnings. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (2008), ACM, pp. 1065–1074.

[5] GUIDELINE, N. E. A. Nist special publication 800-63 version 1.0. 2, 2006.

[6] HONG, J. The state of phishing attacks. *Communications of the ACM 55*, 1 (2012), 74–81.

[7] JACKSON, C., SIMON, D. R., TAN, D. S., AND BARTH, A. An evaluation of extended validation and picture-in-picture phishing attacks. In *Financial Cryptography and Data Security*. Springer, 2007, pp. 281–293.

[8] JAGATIC, T., JOHNSON, N., JAKOBSSON, M., AND MENCZER, F. Phishing attacks using social networks. *Indiana University Human Subject Study*, 05–9892.

[9] KHONJI, M., IRAQI, Y., AND JONES, A. Phishing detection: a literature survey. *Communications Surveys & Tutorials, IEEE 15*, 4, 2091–2121.

[10] SCHECHTER, S. E., DHAMIJA, R., OZMENT, A., AND FISCHER, I. The emperor's new security indicators. In *Security and Privacy, 2007. SP'07. IEEE Symposium on* (2007), IEEE, pp. 51–65.

[11] WHALEN, T., AND INKPEN, K. M. Gathering evidence: use of visual security cues in web browsers. In *Proceedings of Graphics Interface 2005* (2005), Canadian Human-Computer Communications Society, pp. 137–144.

[12] WU, M., MILLER, R. C., AND GARFINKEL, S. L. Do security toolbars actually prevent phishing attacks? In *Proceedings of the SIGCHI conference on Human Factors in computing systems* (2006), ACM, pp. 601–610.

A Appendix: Source Code

This section includes links and descriptions of the code used in the system. All code described here is managed in public git repositories.

- **Browser Extension**
  git@github.com:snyderp/uic-phishing-extension.git
  Javascript code used to build the Firefox and Chrome browser extensions.

- **Recording Server**
  git@github.com:snyderp/bits-phishing-server.git
  Python web server that records the information sent by each browser extension.

- **Phishing Server**
  git@github.com:snyderp/bits-phishing-survey.git
  Python web server that implements the simulated

phishing attack on the university's single sign-in system, along with the debriefing study.

- **Signup Server**
  git@github.com:snyderp/bits-phishing-signup-server.git
  Python web server that runs the system that publishes information about the study, allows users to signup for the study, and verifies that users meet the study participation requirements.