Alice in Warningland: A Large-Scale Field Study of Browser Security Warning Effectiveness

Devdatta Akhawe
*University of California, Berkeley**
devdatta@cs.berkeley.edu

Adrienne Porter Felt
Google, Inc.
felt@google.com

Abstract

We empirically assess whether browser security warnings are as ineffective as suggested by popular opinion and previous literature. We used Mozilla Firefox and Google Chrome's in-browser telemetry to observe over 25 million warning impressions *in situ*. During our field study, users continued through a tenth of Mozilla Firefox's malware and phishing warnings, a quarter of Google Chrome's malware and phishing warnings, and a third of Mozilla Firefox's SSL warnings. This demonstrates that security warnings can be effective in practice; security experts and system architects should not dismiss the goal of communicating security information to end users. We also find that user behavior varies across warnings. In contrast to the other warnings, users continued through 70.2% of Google Chrome's SSL warnings. This indicates that the user experience of a warning can have a significant impact on user behavior. Based on our findings, we make recommendations for warning designers and researchers.

1 Introduction

An oft-repeated maxim in the security community is the futility of relying on end users to make security decisions. Felten and McGraw famously wrote, "Given a choice between dancing pigs and security, the user will pick dancing pigs every time [21]." Herley elaborates [17],

> Not only do users take no precautions against elaborate attacks, they appear to neglect even basic ones. For example, a growing body of measurement studies make clear that ...[users] are oblivious to security cues [27], ignore certificate error warnings [31] and cannot tell legitimate web-sites from phishing imitations [11].[1]

The security community's perception of the "oblivious" user evolved from the results of a number of laboratory studies on browser security indicators [5, 11, 13, 15, 27, 31, 35]. However, these studies are not necessarily representative of the current state of browser warnings in 2013. Most of the studies evaluated warnings that have since been deprecated or significantly modified, often in response to criticisms in the aforementioned studies. Our goal is to investigate whether modern browser security warnings protect users in practice.

We performed a large-scale field study of user decisions after seeing browser security warnings. Our study encompassed 25,405,944 warning impressions in Google Chrome and Mozilla Firefox in May and June 2013. We collected the data using the browsers' telemetry frameworks, which are a mechanism for browser vendors to collect pseudonymous data from end users. Telemetry allowed us to unobtrusively measure user behavior during normal browsing activities. This design provides realism: our data reflects users' actual behavior when presented with security warnings.

In this paper, we present the rates at which users click through (i.e., bypass) malware, phishing, and SSL warnings. Low clickthrough rates are desirable because they indicate that users notice and heed the warnings. Clickthrough rates for the two browsers' malware and phishing warnings ranged from 9% to 23%, and users clicked through 33.0% of Mozilla Firefox's SSL warnings. This demonstrates that browser security warnings can effectively protect most users in practice.

Unfortunately, users clicked through Google Chrome's SSL warning 70.2% of the time. This implies that the user experience of a warning can have a significant impact on user behavior. We discuss several factors that might contribute to this warning's higher clickthrough rates. Our positive findings for the other five warnings suggest that the clickthrough rate for Google Chrome's SSL warning can be improved.

*The Mozilla Firefox experiments were implemented while the author was an intern at Mozilla Corporation.

[1] Citations updated to match our bibliography.

We also consider user behaviors that are indicative of attention to warnings. We find that Google Chrome's SSL clickthrough rates vary by the specific type of error. In Mozilla Firefox, a fifth of users who choose to click through an SSL warning remove a default option, showing they are making cognitive choices while bypassing the warning. Together, these results contradict the stereotype of the wholly oblivious user with no interest in security.

We conclude that users can demonstrate agency when confronted with browser security warnings. Users do not always ignore security warnings in favor of their desired content. Consequently, security experts and platform designers should not dismiss the role of the user. We find that the user experience of warnings can have an enormous impact on user behavior, justifying efforts to build usable warnings.

Contributions. We make the following contributions:

- To our knowledge, we present the first in-depth, large-scale field study of browser security warnings.
- We survey prior laboratory studies of browser security warnings and discuss why our field study data differs from prior research.
- We analyze how demographics (operating system and browser channel), warning frequency, and warning complexity affect users' decisions. Notably, we find evidence suggesting that technically skilled users ignore warnings more often, and warning frequency is inversely correlated with user attention.
- We provide suggestions for browser warning designers and make recommendations for future studies.

2 Background

Web browsers show warnings to users when an attack might be occurring. If the browser is certain that an attack is occurring, it will show an error page that the user cannot bypass. If there is a chance that the perceived attack is a false positive, the browser will show a bypassable warning that discourages the user from continuing. We study only bypassable warnings because we focus on user decisions.

A user *clicks through* a warning to dismiss it and proceed with her original task. A user *leaves* the warning when she navigates away and does not continue with her original task. A *clickthrough rate* describes the proportion of users who clicked through a warning type. When a user clicks through a warning, the user has (1) ignored the warning because she did not read or understand it or (2) made an informed decision to proceed because she believes that the warning is a false positive or her computer is safe against these attacks (e.g., due to an antivirus).

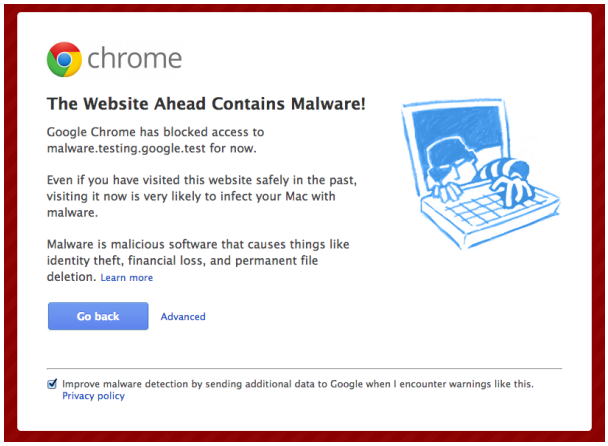

Figure 1: Malware warning for Google Chrome

Figure 2: Malware warning for Mozilla Firefox

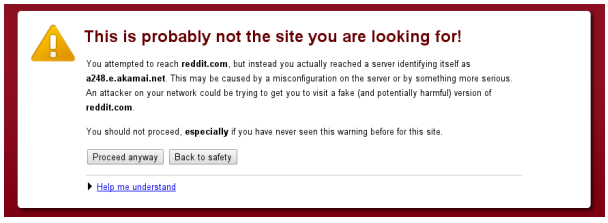

Figure 3: SSL warning for Google Chrome. The first paragraph changes depending on the specific SSL error.

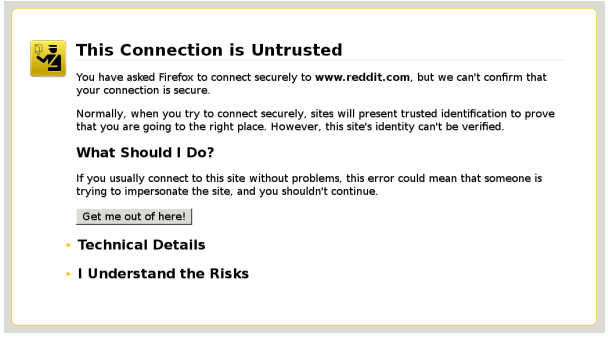

Figure 4: SSL warning for Mozilla Firefox

Figure 5: SSL Add Exception Dialog for Mozilla Firefox

We focus on three types of browser security warnings: malware, phishing, and SSL warnings. At present, all three types of warnings are full-page, interstitial warnings that discourage the user from proceeding.

2.1 Malware and Phishing Warnings

Malware and phishing warnings aim to prevent users from visiting websites that serve malicious executables or try to trick users. Google Chrome and Mozilla Firefox rely on the Google Safe Browsing list [26] to identify malware and phishing websites. The browsers warn users away from the sites instead of blocking them because the Safe Browsing service occasionally has false positives, although the false positive rate is very low [26].

Clickthrough Rate. If a malware or phishing warning is a true positive, clicking through exposes the user to a dangerous situation. Nearly all Safe Browsing warnings are true positives; the false positive rate is low enough to be negligible. The ideal clickthrough rate for malware and phishing warnings is therefore close to 0%.

Warning Mechanisms. The browsers routinely fetch a list of suspicious (i.e., malware or phishing) sites from Safe Browsing servers. If a user tries to visit a site that is on the locally cached list, the browser checks with the Safe Browsing service that the URL is still on the malware or phishing list. If the site is still on one of the lists, the browser presents a warning.

The two browsers behave differently if a page loads a third-party resource (e.g., a script) from a URL on the Safe Browsing list. Google Chrome stops the page load and replaces the page with a warning. Mozilla Firefox blocks the third-party resource with no warning. As a result, Mozilla Firefox users can see fewer warnings than Google Chrome users, despite both browsers using the same Safe Browsing list.

Warning Design. Figures 1 and 2 show the Google Chrome and Mozilla Firefox warnings. Their phishing warnings are similar to their respective malware warnings. When a browser presents the user with a malware or phishing warning, she has three options: leave the page via the warning's escape button, leave the page by closing the window or typing a new URL, or click through the warning and proceed to the page. The warnings also allow the user to seek more information about the error.

Click Count. Mozilla Firefox users who want to bypass the warning need to click one button: the "Ignore this warning" link at the bottom right. On the other hand, Chrome users who want to bypass the warning need to click twice: first on the "Advanced" link, and then on "Proceed at your own risk."

2.2 SSL Warnings

The Secure Sockets Layer (SSL/TLS) protocol provides secure channels between browsers and web servers, making it fundamental to user security and privacy on the web. As a critical step, the browser verifies a server's identity by validating its public-key certificate against a set of trusted root authorities. This validation will fail in the presence of a man-in-the-middle (MITM) attack.

Authentication failures can also occur in a wide variety of benign scenarios, such as server misconfigurations. Browsers usually cannot distinguish these benign scenarios from real MITM attacks. Instead, browsers present users with a warning; users have the option to bypass the warning, in case the warning is a false positive.

Clickthrough Rate. We hope for a 0% clickthrough rate for SSL warnings shown during MITM attacks. However, many SSL warnings may be false positives (e.g., server misconfigurations). There are two competing views regarding SSL false positives. In the first, warning text should discourage users from clicking through both true and false positives, in order to incentivize developers to get valid SSL certificates. In the other, warning text should provide users with enough information to correctly identify and dismiss false positives. The desired clickthrough rates for false-positive warnings would be 0% and 100%, respectively. In either case, false positives are undesirable for the user experience because we do not want to annoy users with invalid warnings. Our goal is therefore a 0% clickthrough rate for all SSL warnings: users should heed all valid warnings, and the browser should minimize the number of false positives.

Warning Design. Figures 3 and 4 present Google Chrome and Mozilla Firefox's SSL warnings. The user can leave via the warning's escape button, manually navigate away, or click through the warning. In Mozilla Firefox, the user must also click through a second dialog (Figure 5) to bypass the warning.

The browsers differ in their presentation of the technical details of the error. Google Chrome places information about the specific error in the main warning (Figure 3, first paragraph), whereas Firefox puts the error information in the hidden "Technical Details" section and the second "Add Exception" dialog (Figure 5).

Click Count. Mozilla Firefox's SSL warning requires more clicks to bypass. Google Chrome users click through a single warning button to proceed. On the other hand, Mozilla Firefox's warning requires three clicks: (1) click on "I Understand The Risks," (2) click on the "Add Exception" button, which raises a second dialog, (3) click on "Confirm Security Exception" in the second dialog. By default, Firefox permanently remembers the exception and will not show the warning again if the user reencounters the same certificate for that website. In contrast, Chrome presents the warning every time and does not remember the user's past choices.

2.3 Browser Release Channels

Mozilla and Google both follow rapid release cycles. They release official versions of their browsers every six or seven weeks, and both browsers update automatically. The official, default version of a browser is referred to as "stable" (Google Chrome) or "release" (Mozilla Firefox).

If users are interested in testing pre-release browser versions, they can switch to a different *channel*. The stable/release channel is the recommended channel for end users, but a minority of users choose to use earlier channels to test cutting-edge features. The "Beta" channel is several weeks ahead of the stable/release channel. The "developer" (Google Chrome) or "Aurora" (Mozilla Firefox) channel is delivered even earlier. Both browsers also offer a "nightly" (Mozilla Firefox) or "Canary" (Google Chrome) release channel, which updates every day and closely follows the development repository.

The pre-release channels are intended for advanced users who want to experience the latest-and-greatest features and improvements. They give website, extension, and add-on developers time to test their code on upcoming versions before they are deployed to end users. The early channels are not recommended for typical end users because they can have stability issues, due to being under active development. The rest of this paper assumes a positive correlation between pre-release channels use and technical ability. While this matches the intention of browser developers, we did not carry out any study to validate this assumption.

3 Prior Laboratory Studies

We survey prior laboratory studies of SSL and phishing warnings. The body of literature paints a grim picture of browser security warnings, but most of the warnings have since been deprecated or modified. In some cases, warnings were changed in response to these studies.

Only two studies evaluated warnings that are similar to the modern (June 2013) browser warnings that we study in this paper. Sunshine et al. and Sotirakopoulos et al. reported clickthrough rates of 55% to 80% for the Firefox 3 and 3.5 SSL warnings, which are similar but not identical to the current Firefox SSL warning [30, 31]. However, Sotirakopoulos et al. concluded that laboratory biases had inflated both studies' clickthrough rates [30].

3.1 SSL Warnings

SSL warnings are the most studied type of browser warning. Usability researchers have evaluated SSL warnings in both SSL-specific studies and phishing studies because SSL warnings and passive indicators were once viewed as a way to identify phishing attacks.[2]

Dhamija et al. performed the first laboratory study of SSL warnings in 2006. They challenged 22 study participants to differentiate between phishing and legitimate websites in Mozilla Firefox 1.0.1 [11]. In this version, the warning was a modal dialog that allowed the user to permanently accept, temporarily accept, or reject the certificate. When viewing the last test website, participants encountered an SSL warning. The researchers reported that 15 of their 22 subjects (68%) quickly clicked through the warning without reading it. Only one user was later able to tell the researchers what the warning had said. The authors considered the clickthrough rate of 68% a conservative lower bound because participants knew that they should be looking for security indicators.

In 2007, Schechter et al. studied user reactions to Internet Explorer 7's SSL warning, which is the same one-click interstitial that is present in all subsequent versions of Internet Explorer [27]. Participants encountered the warning while logging into a bank website to look up information. The researchers were aware of ecological validity concerns with laboratory studies and split their participants into three groups: participants who entered their own credentials, a role-playing group that entered fake passwords, and a security-primed role-playing group that entered fake passwords. Overall, 53% of the total 57 participants clicked through. However, only 36% of the non-role-playing group clicked through. The difference between the role-playing participants and non-role-playing partic-

[2] There is evidence that modern phishing sites can have valid SSL certificates [24].

ipants was statistically significant, illustrating one challenge of experiments in artificial environments.

Sunshine et al. performed multiple studies of SSL warnings in 2009 [31]. First, they conducted an online survey. They asked 409 people about Firefox 2, Firefox 3, and Internet Explorer 7 warnings. Firefox 2 had a modal dialog like Firefox 1.0.1, and Firefox 3's warning is similar but not identical to the current Firefox warning. Less than half of respondents said they would continue to the website after seeing the warning. As a follow-up, Sunshine et al. also conducted a laboratory study that exposed 100 participants to SSL warnings while completing information lookup tasks. The clickthrough rates were 90%, 55%, and 90% when participants tried to access their bank websites in Firefox 2, Firefox 3, and Internet Explorer 7, respectively. The clickthrough rates increased to 95%, 60%, and 100% when participants saw an SSL warning while trying to visit the university library website.

Sotirakopoulos et al. replicated Sunshine's laboratory SSL study with a more representative population sample [30]. In their study, 80% of participants using Firefox 3.5 and 72% of participants using Internet Explorer 7 clicked through an SSL warning on their bank website. More than 40% of their participants said that the laboratory environment had influenced them to click through the warnings, either because they felt safe in the study environment or were trying to complete the experimental task. Sotirakopoulos et al. concluded that the laboratory environment biased their results, and they suspect that these biases are also present in similar laboratory studies.

Bravo-Lillo et al. interviewed people about an SSL warning from an unspecified browser [5]. They asked 20 participants about the purpose of the warning, what would happen if a friend were to click through, and whether a friend should click through the warning. Participants were separated into "advanced" and "novice" browser users. "Advanced" participants said they would not click through an SSL warning on a bank website, but "novice" participants said they would.

Passive Indicators. Some studies focused on passive SSL indicators, which non-interruptively show the status of the HTTP(S) connection in the browser UI. Although browsers still have passive SSL indicators, interruptive SSL and phishing warnings are now the primary tool for communicating security information to users.

Friedman et al. asked participants whether screenshots of websites depicted secure connections; many participants could not reliably determine whether a connection was secure [15]. Whalen and Inkpen used eye-tracking software to determine that none of their 16 participants looked at the lock or key icon in the URL bar, HTTP(S) status in the URL bar, or the SSL certificate when asked to browse websites "normally" [34]. Some browsers modify the lock icon or color of the URL bar to tell the user when a website has an Extended Validation (EV) SSL certificate. Jackson et al. asked 27 study subjects to classify 12 websites as either phishing or legitimate sites, but the EV certificates did not help subjects identify the phishing sites [19]. In a follow-up study, Sobey et al. found that none of their 28 subjects clicked on the EV indicators, and the presence of EV indicators did not affect decision-making [29]. Similarly, Biddle et al. found that study participants did not understand Internet Explorer's certificate summaries [3].

In 2012, a Google Chrome engineer mentioned high clickthrough rates for SSL warnings on his blog [20]. We expand on this with a more accurate and detailed view of SSL clickthrough rates in Google Chrome.

3.2 Phishing Warnings

Phishing warnings in contemporary browsers are active, interstitial warnings; in the past, they have been passive indicators in toolbars. Researchers have studied whether they are effective at preventing people from entering their credentials into phishing websites.

Wu et al. studied both interstitial and passive phishing warnings [35]. Neither of the warnings that they evaluated are currently in use in browsers. First, they launched phishing attacks on 30 participants. The participants role-played during the experiment while using security toolbars that display passive phishing warnings. Despite the toolbars, at least one attack fooled 20 out of 30 participants. In their next experiment, they asked 10 study participants to perform tasks on PayPal and a shopping wish list website; they injected modal phishing warnings into the websites. None of the subjects entered the credentials into the PayPal site, but the attack on the wish list site fooled 4 subjects. The authors do not report the warning clickthrough rates.

Egelman et al. subjected 60 people to simulated phishing attacks in Internet Explorer 7 or Mozilla Firefox 2.0 [13]. Firefox 2.0 had a modal phishing dialog that is not comparable to the current Mozilla Firefox phishing dialog, and Internet Explorer had both passive and active warnings. Participants believed that they were taking part in a laboratory study about shopping. The researchers asked participants to check their e-mail, which contained both legitimate shopping confirmation e-mails and similar spear phishing e-mails sent by the researchers. Users who clicked on the links in the phishing e-mails saw a phishing warning. Participants who saw Mozilla Firefox's active warning, Internet Explorer's active warning, or Internet Explorer's passive warning were phished 0%, 45%, and 90% of the time, respectively. The clickthrough rates were an unspecified superset of the rates at which people fell for the phishing attacks.

3.3 Malware Download Warnings

Google Chrome and Microsoft Internet Explorer also display non-blocking warning dialogs when users attempt to download malicious executables. In a blog post, a Microsoft employee stated that the clickthrough rate for Internet Explorer's SmartScreen warning was under 5% [16]. We did not study this warning for Google Chrome, and Mozilla Firefox does not have this warning.

4 Methodology

We rely on the telemetry features implemented in Mozilla Firefox and Google Chrome to measure clickthrough rates *in situ*. Telemetry is a mechanism for browser vendors to collect pseudonymous data from end users who opt in to statistics reporting. Google Chrome and Mozilla Firefox use similar telemetry platforms.

4.1 Measuring Clickthrough Rates

We implemented metrics in both browsers to count the number of times that a user sees, clicks through, or leaves a malware, phishing, or SSL warning. Based on this data, we can calculate clickthrough rates for each warning type. As discussed in Section 2, we report only the clickthrough rates for warnings that the user can bypass. We measured the prevalence of non-bypassable warnings separately. To supplement the clickthrough rates, we recorded whether users clicked on links like "Help me understand," "View," or "Technical Details."

Bypassing some warnings takes multiple clicks, and our clickthrough rates for these warnings represent the number of users who completed all of the steps to proceed to the page. For Mozilla Firefox's SSL warning (which takes three clicks to proceed), we recorded how often users perform two intermediate clicks (on "Add Exception" or "Confirm Security Exception") as well as the overall clickthrough rate.

We also measured how often users encounter and click through specific SSL errors. In addition to the overall clickthrough rates for the warnings, we collected clickthrough data for each type of Mozilla Firefox SSL error and the three most common Google Chrome SSL errors.

Our Mozilla Firefox data set does not allow us to track specific telemetry participants. In Google Chrome, we can correlate warning impressions with psuedonymous browser client IDs; however, the sample size for most individual users is too small to draw conclusions. We therefore report the results of measurements aggregated across all users unless otherwise specified. The telemetry frameworks do not provide us with any personal or demographic information except for the operating system and browser version for each warning impression.

4.2 Measuring Time Spent on Warnings

We also used the Google Chrome telemetry framework to observe how much time Google Chrome users spent on SSL warnings. Timing began as soon as an SSL warning came to the foreground in a tab. In particular,

- We recorded the time spent on a warning and associated it with the outcome (click through or leave).
- We recorded the time spent on a warning and associated it with the error type, if it was one of the three most common error types (untrusted authority, name mismatch, and expired certificate).

Together, these correspond to five timing measurements (two for outcome and three for error type). For scalability, the telemetry mechanism in Google Chrome only allows timing measurements in discrete buckets. As a result, our analysis also treats time as a discrete, ordinal variable.

We used log-scaled bucket sizes (e.g., the first bucket size is 45ms but the last is 90,279ms) with 50 buckets, ranging from 0ms to 1,200,000ms, for the two outcome histograms. The three error type histograms had 75 buckets each, ranging from 0ms to 900,000ms. We used more buckets for the error histograms because we anticipated that they would be more similar to each other.

4.3 Ethics

We collected data from users who participate in their browsers' broad, unpaid user metrics programs. At first run of a browser, the browser asks the user to share usage data. If the user consents, the browser collects data on performance, features, and stability. In some pre-release developer channels, data collection is enabled by default. The browser periodically sends this pseudonymous data over SSL to the central Mozilla or Google servers for analysis. The servers see the IP addresses of clients by necessity, but they are not attached to telemetry data. All telemetry data is subject to strict privacy policies and participants can opt out by changing their settings [7, 23]. Multiple Google Chrome committers and Mozilla Firefox contributors reviewed the data collection code to ensure that the metrics did not collect any private data.

This work is not considered human subjects research by UC Berkeley because the student did not have access to database identifiers or personally identifying information.

4.4 Data Collection

Collection Period. Google Chrome's malware and phishing measurement code was in place in Chrome 24 prior to our work, and our SSL measurement code was added to Google Chrome 25. The Google Chrome data in this

paper was collected April 28 - May 31, 2013. Our Mozilla Firefox measurement code was added to Firefox 17, and a bug in the SSL measurement code was fixed in Firefox 23. The data on the Firefox malware warning, phishing warning, and SSL "Add Exception" dialog was collected May 1-31, 2013. The data on Firefox SSL warnings was collected June 1 - July 5, 2013, as the Firefox 23 fix progressed through the various release channels.

Sample Sizes. In Google Chrome, we recorded 6,040,082 malware warning impressions, 386,350 phishing warning impressions, and 16,704,666 SSL warning impressions. In Mozilla Firefox, we recorded 2,163,866 malware warning impressions, 100,004 phishing warning impressions, and 10,976 SSL warning impressions. Appendix A further breaks downs these sample sizes by OS and channel.

Number of Users. For Mozilla Firefox, we recorded warning impressions from the approximately 1% of Firefox users who opt in to share data with Mozilla via telemetry. In Google Chrome, we observed malware, phishing, and SSL warning impressions on 2,148,026; 204,462; and 4,491,767 clients (i.e., browser installs), respectively.

4.5 Method Limitations

Private Data. Due to privacy constraints, we could not collect information about users' personal demographics or browsing habits. Consequently, we cannot measure whether user behavior differs based on personal characteristics, the target site, or the source of the link to the site. We also cannot identify SSL false positives due to captive portals, network proxies, or server misconfigurations.

Sampling Bias. The participants in our field study are not a random population sample. Our study only represents users who opt in to browser telemetry programs. This might present a bias. The users who volunteered might be more likely to click through dialogs and less concerned about privacy. Thus, the clickthrough rates we measure could be higher than population-wide rates. Given that most of our observed rates are low, this bias augments our claim that clickthrough rates are lower than anticipated.

Overrepresentation. We present clickthrough rates across all warnings shown to all users. A subset of users could potentially be overrepresented in our analysis. Within the Google Chrome data set, we identified and removed a small number of overrepresented clients who we believe are either crawlers or malware researchers. We were unable to remove individual clients from the Mozilla Firefox set, but we do not believe this represents a bias because we know that the overrepresented clients in Chrome still contributed fewer than 1% of warning impressions. Some clients experienced multiple types of warning impressions; we investigated this in Chrome and found that the clickthrough rates do not differ if we remove non-independent clients. Our large sample sizes and small critical value ($\alpha = 0.001$) should further ameliorate these concerns.

Frames. Our original measurement for Mozilla Firefox did not differentiate between warnings shown in top-level frames (i.e., warnings that fill the whole tab) and warnings shown in iframes. In contrast, Google Chrome always shows malware and phishing warnings in the top-level frame and does not render any warning type in iframes. Since users might not notice warnings in iframes, the two metrics are not necessarily directly comparable.

Upon discovering this issue, we modified our Firefox measurement implementation to take frame level into account. Our new implementation is not available to all Firefox users yet, but we have data for recent pre-release channels. For malware and phishing warning impressions collected from the beta channel, the clickthrough rate for the top-level frame is within two percentage points of the overall clickthrough rate. This is due to the relative infrequency of malware and phishing warnings in iframes and the low overall clickthrough rate. Since the frame level does not make a notable difference for malware and phishing warnings, we present the overall rates (including both top-level frames and iframes) for the full sample sizes in Section 5.1. The difference is more important for SSL warnings: the clickthrough rate for top-level frames is 28.7 percentage points higher than the overall clickthrough rate of 4.3%. Consequently, Section 5.2 presents only the top-level frame rate for SSL warnings, although it limits our sample to pre-release users.

5 Clickthrough Rates

We present the clickthrough data from our measurement study. Section 5.1 discusses malware and phishing warnings together because they share a visual appearance. We then present rates for SSL warnings in Section 5.2.

5.1 Malware and Phishing Warnings

The clickthrough rates for malware warnings were 7.2% and 23.2% in stable versions of Mozilla Firefox and Google Chrome, respectively. For phishing warnings, we found clickthrough rates of 9.1% and 18.0%. In this section, we discuss the effects of warning type, demographics, and browser on the clickthrough rates.

5.1.1 Malware Rates by Date

The malware warning clickthrough rates for Google Chrome vary widely by date. We have observed clickthrough rates ranging from 11.2% to 24.9%, depending

Operating System	Malware		Phishing	
	Firefox	Chrome	Firefox	Chrome
Windows	7.1%	23.5%	8.9%	17.9%
MacOS	11.2%	16.6%	12.5%	17.0%
Linux	18.2%	13.9%	34.8%	31.0%

Table 1: User operating system vs. clickthrough rates for malware and phishing warnings. The data comes from stable (i.e., release) versions.

Channel	Malware		Phishing	
	Firefox	Chrome	Firefox	Chrome
Stable	7.2%	23.2%	9.1%	18.0%
Beta	8.7%	22.0%	11.2%	28.1%
Dev	9.4%	28.1%	11.6%	22.0%
Nightly	7.1%	54.8%	25.9%	20.4%

Table 2: Release channel vs. clickthrough rates for malware and phishing warnings, for all operating systems.

on the week, since the current version of the warning was released in August 2012. In contrast, the Mozilla Firefox malware warning clickthrough rate across weeks stays within one percentage point of the month-long average. We did not observe similar temporal variations for phishing or SSL warnings.

Recall from Section 2.1 that Google Chrome and Mozilla Firefox's malware warnings differ with respect to secondary resources: Google Chrome shows an interstitial malware warning if a website includes secondary resources from a domain on the Safe Browsing list, whereas Mozilla Firefox silently blocks the resource. We believe that this makes Google Chrome's malware clickthrough rates more sensitive to the contents of the Safe Browsing list. For example, consider the case where a well-known website accidentally loads an advertisement from a malicious domain. Google Chrome would show a warning, which users might not believe because they trust the website. Mozilla Firefox users would not see any warning. Furthermore, Chrome phishing warnings are less likely to be due to secondary resources, and that warning's clickthrough rates do not vary much by time.

5.1.2 Malware/Phishing Rates by Warning Type

In Mozilla Firefox, we find a significantly higher clickthrough rate for phishing warnings than malware warnings (χ^2 test: $p(1) < 0.0001$). This behavior is rational: a malware website can infect the user's computer without any action on the user's part, but a phishing website can only cause harm by tricking the user at a later point in time. Mozilla Firefox makes this priority ordering explicit by choosing to display the malware warning if a website is listed as both malware and phishing.[3] However, the practical difference is small: 7.2% vs. 9.1%.

In Google Chrome, the average malware clickthrough rate is higher than the phishing clickthrough rate. However, the malware clickthrough rate fluctuates widely (Section 5.1.1); the malware clickthrough rate is sometimes lower than the phishing clickthrough rate.

5.1.3 Malware/Phishing Rates by Demographics

We consider whether users of different operating systems and browser release channels react differently to warnings. As Table 1 depicts, Linux users have significantly higher clickthrough rates than Mac and Windows users combined for the Firefox malware warning, Firefox phishing warning, and Chrome phishing warning (χ^2 tests: $p(1) < 0.0001$). While the low prevalence of malware for Linux could explain the higher clickthrough rates for the Firefox malware warning, use of Linux does not provide any additional protection against phishing attacks. The Chrome malware warning does not follow the same pattern: Windows users have a significantly higher clickthrough rate (χ^2 tests: $p(1) < 0.0001$).

We also see differences between software release channels (Table 2). Nightly users click through Google Chrome malware and Firefox phishing warnings at much higher rates than stable users, although they click through Firefox malware and Google Chrome phishing warnings at approximately the same rates.

In several cases, Linux users and early adopters click through malware and phishing warnings at higher rates. One possible explanation is that a greater degree of technical skill – as indicated by use of Linux or early-adopter versions of browsers – corresponds to reduced risk aversion and an increased willingness to click through warnings. This does not hold true for all categories and warnings (e.g., nightly and stable users click through the Firefox malware warning at the same rate), suggesting the need for further study.

5.1.4 Malware/Phishing Rates by Browser

Google Chrome stable users click through phishing warnings more often than Mozilla Firefox stable users. This holds true even when we account for differences in how the browsers treat iframes (Section 4.5). Mozilla Firefox's beta channel users still click through warnings at a lower rate when we exclude iframes: 9.6% for malware warnings, and 10.8% for phishing warnings.

One possibility is that Mozilla Firefox's warnings are more frightening or more convincing. Another possi-

[3] Google Chrome will display both warnings. To preserve independence, our measurement does not include any warnings with both phishing and malware error messages. Dual messages are infrequent.

bility is that the browsers have different demographics with different levels of risk tolerance, which is reflected in their clickthrough rates. There might be differences in technical education, gender, socioeconomic status, or other factors that we cannot account for in this study. In support of this theory, we find that differences between the browsers do not hold steady across operating systems or channels. The gap between the browsers narrows or reverses for some categories of users, such as Linux users and nightly release users.

5.2 SSL Warnings

The clickthrough rates for SSL warnings were 33.0% and 70.2% for Mozilla Firefox (beta channel) and Google Chrome (stable channel), respectively.

5.2.1 SSL Rates by Demographic

In Section 5.1, we observed that malware and phishing clickthrough rates differed across operating systems and channels. For SSL, the differences are less pronounced.

As with the malware and phishing warnings, nightly users click through SSL warnings at a higher rate for both Firefox and Chrome (χ^2 tests: $p < 0.0001$).

The effect of users' operating systems on SSL clickthrough rates differs for the two browsers. In Firefox, Linux users are much more likely to click through SSL warnings than Windows and Mac users combined (χ^2 test: $p < 0.0001$), although it is worth noting that the Firefox Linux sample size is quite small (58). In Chrome, Windows users are very slightly more likely to click through SSL warnings than Linux and Mac users combined (χ^2 test: $p < 0.0001$).

5.2.2 SSL Rates by Browser

We find a large difference between the Mozilla Firefox and Google Chrome clickthrough rates: Google Chrome users are 2.1 times more likely to click through an SSL warning than Mozilla Firefox users. We explore five possible causes.

Number of Clicks. Google Chrome users click one button to dismiss an SSL warning, but Mozilla Firefox users need to click three buttons. It is possible that the additional clicks deter people from clicking through. However, we do not believe this is the cause of the rate gap.

First, the number of clicks does not appear to affect the clickthrough rates for malware and phishing warnings. Mozilla Firefox's malware and phishing warnings require one click to proceed, whereas Google Chrome's malware and phishing warnings require two. The Google Chrome malware and phishing warnings with two clicks do not have lower clickthrough rates than the Mozilla Firefox warnings with one click. Second, as we discuss in Section 5.2.3, 84% of users who perform the first two clicks in Mozilla Firefox also perform the third. This indicates that the extra click is not a determining decision point. Unfortunately, we do not have data on the difference between the first and second clicks.

Operating System	SSL Warnings	
	Firefox	Chrome
Windows	32.5%	71.1%
MacOS	39.3%	68.8%
Linux	58.7%	64.2%
Android	NC	64.6%

Table 3: User operating system vs. clickthrough rates for SSL warnings. The Google Chrome data is from the stable channel, and the Mozilla Firefox data is from the beta channel.

Channel	SSL Warnings	
	Firefox	Chrome
Release	NC	70.2%
Beta	32.2%	73.3%
Dev	35.0%	75.9%
Nightly	43.0%	74.0%

Table 4: Channel vs. clickthrough rates for SSL warnings.

Warning Appearance. The two warnings differ in several ways. Mozilla Firefox's warning includes an image of a policeman and uses the word "untrusted" in the title. These differences likely contribute to the rate gap. However, we do not think warning appearance is the sole or primary factor; the browsers' malware and phishing warnings also differ, but there is only about a 10% difference between browsers for these warnings.

Certificate Pinning. Google Chrome ships with a list of "pinned" certificates and preloaded HTTP Strict Transport Security (HSTS) sites. Users cannot click through SSL warnings on sites protected by these features. Certificate pinning and HSTS cover some websites with important private data such as Google, PayPal, and Twitter [8]. In contrast, Mozilla Firefox does not come with many preloaded "pinned" certificates or any pre-specified HSTS sites. As a result, Chrome shows more non-bypassable warnings: our field study found that 20% of all Google Chrome SSL warning impressions are non-bypassable, as compared to 1% for Mozilla Firefox.

Based on this, we know that Mozilla Firefox users see more warnings for several critical websites. If we assume that users are less likely to click through SSL warnings on these critical websites, then it follows that Mozilla Firefox's clickthrough rate will be lower. This potential bias could account for up to 15 points of the 37-point gap between the two clickthrough rates, if we were to assume that Google Chrome users would never click through SSL errors on critical websites if given the chance.

Remembering Exceptions. Due to the "permanently store this exception" feature in Mozilla Firefox, Mozilla Firefox users see SSL warnings only for websites without saved exceptions. This means that Mozilla Firefox users might ultimately interact with websites with SSL errors at the same rate as Google Chrome users despite having lower clickthrough rates. For example, imagine a user that encounters two websites with erroneous SSL configuration: she leaves the first after seeing a warning, but visits the second website nine times despite the warning. This user would have a 50% clickthrough rate in Mozilla Firefox but a 90% clickthrough rate in Google Chrome, despite visiting the second website at the same rate.

We did not measure how often people revisit websites with SSL errors. However, we suspect that people do repeatedly visit sites with warnings (e.g., a favorite site with a self-signed certificate). If future work were to confirm this, there could be two implications. First, if users are repeatedly visiting the same websites with errors, the errors are likely false positives; this would mean that the lack of an exception-storing mechanism noticably raises the false positive rate in Google Chrome. Second, warning fatigue could be a factor. If Google Chrome users are exposed to more SSL warnings because they cannot save exceptions, they might pay less attention to each warning that they encounter.

Demographics. It's possible that the browsers have different demographics with different levels of risk tolerance. However, this factor likely only accounts for a few percentage points because the same demographic effect applies to malware and phishing warnings, and the difference between browsers for malware and phishing warnings is much smaller.

5.2.3 SSL Rates by Certificate Error Type

To gain insight into the factors that drive clickthrough rates, we study whether the particular certificate error affects user behavior.

Google Chrome. Google Chrome's SSL warning includes a short explanation of the particular error, and clicking on "Help me understand" will open a more-detailed explanation. In case a certificate has multiple errors, Google Chrome only shows the first error out of untrusted issuer error, name mismatch error, and certificate expiration error, respectively.

Table 5 presents the clickthrough rates by error types for Google Chrome. If Google Chrome users are paying attention to and understanding the warnings, one would expect different clickthrough rates based on the warning types. We find a 24.4-point difference between the clickthrough rates for untrusted issuer errors and expired certificate errors. One explanation could be that untrusted issuer errors appear on unimportant sites, leading to higher clickthrough rates without user attention or comprehension; however, the Mozilla Firefox data suggests otherwise. An alternative explanation could be that expired certificates, which often occur for websites with previously valid certificates [1], surprise the user. In contrast, untrusted certificate errors always occur for a website and conform with expectations.

Certificate Error	Percentage of Total	Clickthrough Rate
Untrusted Issuer	56.0%	81.8%
Name Mismatch	25.0%	62.8%
Expired	17.6%	57.4%
Other Error	1.4%	–
All Error Types	100.0%	70.2%

Table 5: Prevalence and clickthrough rates of error types for the Google Chrome SSL warning. Google Chrome only displays the most critical warning; we list the error types in order, with untrusted issuer errors as the most critical. Data is for the stable channel across all operating systems.

Mozilla Firefox. Mozilla Firefox's SSL warning does not inform the user about the particular SSL error by default.[4] Instead, the secondary "Add Exception" dialog presents *all* errors in the SSL certificate. The user must confirm this dialog to proceed.

Table 6 presents the rates at which users confirm the "Add Exception" dialog in Mozilla Firefox. The error types do not greatly influence the exception confirmation rate. This indicates that the "Add Exception" dialog does not do an adequate job of explaining particular error categories and their meaning to the users. Thus, users ignore the categories and click through errors at the same rate. This finding also suggests that the differences in clickthrough rates across error types in Google Chrome cannot be attributed to untrusted issuer errors corresponding to unimportant websites; if that were the case, we would expect to see the same phenomenon in Firefox.

Error Prevalence. The frequency of error types encountered by users in our field study also indicates the base rate of SSL errors on the web. Our Google Chrome data contradicts a previous network telemetry study, which suggested that untrusted issuer errors correspond to 80% of certificate errors seen on the wire [18]. Also, Google Chrome users see fewer untrusted issuer errors than Mozilla Firefox users; this may be because Mozilla Firefox users are more likely to click on the "Add Exception" dialog for untrusted issuer errors. Recall that we collect the Mozilla Firefox error type statistics only after a user clicks on the "Add Exception" button.

[4] This information is available under the "Technical details" link, but our measurements indicate that it is rarely opened (Section 5.2.4).

Certificate Error	Percentage of Total	Confirmation Rate
Untrusted Issuer	38%	87.1%
Untrusted and Name Mismatch	26.4%	87.9%
Name Mismatch	15.7%	80.3%
Expired	10.2%	80.7%
Expired, Untrusted and Name Mismatch	4.7%	87.6%
Expired and Untrusted	4.1%	83.6%
Expired and Name Mismatch	0.7%	85.2%
None of the above	<0.1%	77.9%
All error types	100.0%	85.4%

Table 6: Prevalence and confirmation rates of error types for the Mozilla Firefox "Add Exception" dialog. The confirmation rate measures the percentage of users who click on "Confirm Security Exception" (Figure 5). The Mozilla Firefox dialog lists all the errors that occur for a certificate. Data is for the release channel across all operating systems; we did not need to limit it to the beta channel because frame level issues do not affect clickthrough rates inside the "Add Exception" dialog.

The high frequency of untrusted issuer errors highlights the usability benefits of "network view" SSL certificate verification systems like Perspectives and Convergence [10,33], which do not need certificates from trusted authorities. All of the untrusted certificate warnings—between 38% and 56% of the total—would disappear. Warnings with other errors in addition to an untrusted certificate error would remain. Nonetheless, our study also shows that these mechanisms are not a panacea: name mismatch errors constitute a large fraction of errors, and new systems like Perspectives and Convergence still perform this check.[5]

5.2.4 Additional SSL Metrics

We collected several additional metrics to complement the overall clickthrough rates.

More Information. Google Chrome and Mozilla Firefox both place additional information about the warning behind links. However, very few users took the opportunity to view this extra information. The "Help me understand" button was clicked during 1.6% of Google Chrome SSL warning impressions. For Mozilla Firefox warnings, 0 users clicked on "Technical Details," and 3% of viewers of the "Add Exception" dialog clicked on "View Certificate." This additional content therefore has no meaningful impact on the overall clickthrough rates.

Add Exception Cancellation. Not all Mozilla Firefox users proceed to the page after opening the "Add Exception" dialog: 14.6% of the time that a dialog is opened, the user cancels the exception. These occurrences indicate that at least a minority of users consider the text in the dialog before confirming the exception.

[5]Convergence does not check the certificate issuer, relying on network views instead. However, it performs name checks [10].

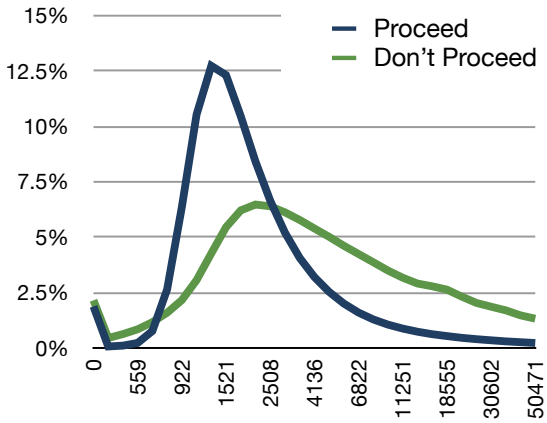

Figure 6: Google Chrome SSL clickthrough times (ms), by outcome. The graph shows the percent of warning impressions that fall in each timing bucket. The x-axis increases logarithmically, and we cut off the distribution at 90% due to the long tail.

Remember Exception. By default, the "Remember Exception" checkbox is checked in the Mozilla Firefox "Add Exception" dialog. Our measurements found that 21.3% of the time that the dialog is opened, the user un-ticks the checkbox. We hypothesize that these users are still wary of the website even if they choose to proceed.

6 Time Spent On SSL Warnings

In addition to MITM attacks, SSL warnings can occur due to server misconfigurations. Previous work found that 20% of the thousand most popular SSL sites triggered a false warning due to such misconfigurations [31]. Consequently, it may be safe and rational to click through such false warnings. The prevalence of a large number of such false warnings can potentially train users to consider *all* SSL warnings false alarms and click through them without considering the context.

In order to determine whether users examine SSL warnings before making a decision, we measured how much time people spent on SSL warning pages. In this section, we compare the click times by outcome (clickthrough or leave) and error type to gain insight into user attention. Our timing data is for all operating systems and channels.

6.1 Time by Outcome

Figure 6 presents the click times for different outcomes. Users who leave spend more time on the warning than

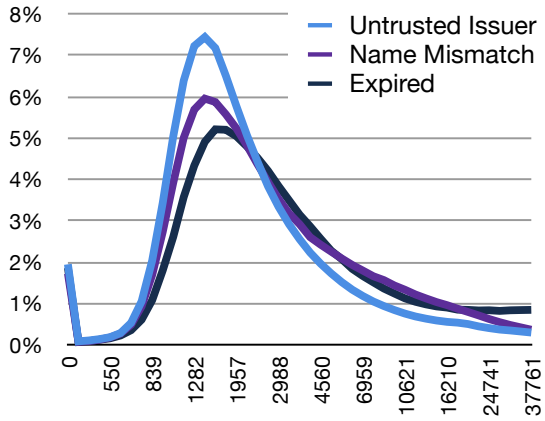

Figure 7: Google Chrome SSL clickthrough times (ms), by error type. The graph shows the percent of warning impressions that fall in each timing bucket. The x-axis increases logarithmically, and we cut off the distribution at 90% due to the long tail.

users who click through and proceed to the page. 47% of users who clicked through the warning made the decision within 1.5s, whereas 47% of users who left the page did so within 3.5s. We interpret this to mean that users who click through the warning often do so after less consideration.

6.2 Time by Error Type

Figure 7 depicts the click times for three error types (untrusted authority, name mismatch, and expired certificate errors). Users clicked through 49% of untrusted issuer warning impressions within 1.7s, but clicked through 50% of name and date errors within 2.2s and 2.7s, respectively. We believe that this data is indicative of warning fatigue: users click through more-frequent errors more quickly. The frequency and clickthrough rate of each error type (as reported in Section 5.2) are inversely correlated with that error type's timing variance and mode (Figure 7).

7 Implications

Our primary finding is that browser security warnings *can* be effective security mechanisms in practice, but their effectiveness varies widely. This should motivate more attention to improving security warnings. In this section, we summarize our findings and their implications, present suggestions for warning designers, and make recommendations for future warning studies.

7.1 Warning Effectiveness

7.1.1 Clickthrough Rates

Popular opinion holds that browser security warnings are ineffective. However, our study demonstrates that browser security warnings can be highly effective at preventing users from visiting websites: as few as a tenth of users click through Firefox's malware and phishing warnings. We consider these warnings very successful.

We found clickthrough rates of 18.0% and 23.2% for Google Chrome's phishing and malware warnings, respectively, and 31.6% for Firefox's SSL warning. These warnings prevent 70% (or more) of attempted visits to potentially dangerous websites. Although these warnings could be improved, we likewise consider these warnings successful at persuading and protecting users.

Google Chrome's SSL warning had a clickthrough rate of 70.2%. Such a high clickthrough rate is undesirable: either users are not heeding valid warnings, or the browser is annoying users with invalid warnings and possibly causing warning fatigue. Our positive findings for the other warnings demonstrate that this warning has the potential for improvement. We hope that this study motivates further studies to determine and address the cause of its higher clickthrough rate. We plan to test an exception-remembering feature to investigate the influence of repeat exposures to warnings. At Google, we have also begun a series of A/B tests in the field to measure the impact of a number of improvements.

7.1.2 User Attention

Although we did not directly study user attention, two results of our study suggest that at least a minority of users pay attention to browser security warnings.

- There is a 24.4-point difference between the clickthrough rates for untrusted issuer errors (81.8%) and expired certificate errors (57.4%) in Google Chrome.
- 21.3% of the time that Mozilla Firefox users viewed the "Add Exception" dialog, they un-checked the default "Permanently store this exception" option.

These results contradict the stereotype of wholly oblivious users with no interest in security.

7.2 Comparison with Prior Research

As Bravo-Lillo et al. wrote [5]:

> Evidence from experimental studies indicates that most people don't read computer warnings, don't understand them, or simply don't heed them, even when the situation is clearly hazardous.

In contrast, a majority of users heeded five of the six types of browser warnings that we studied. This section explores why our results differ from prior research.

Browser Changes. Most prior browser research was conducted between 2002 and 2009. Browsers were rapidly

changing during this time period; some changes were directly motivated by published user studies. Notably, passive indicators are no longer considered primary security tools, and phishing toolbars have been replaced with browser-provided, full-page interstitial warnings. As a result, studies of passive indicators and phishing toolbars no longer represent the state of modern browser technology.

Two studies tested an older version of the Mozilla Firefox SSL warning, in which the warning was a modal (instead of full-page) dialog. Dhamija et al. observed a 68% clickthrough rate, and Sunshine et al. recorded clickthrough rates of 90%-95% depending on the type of page [11, 31]. The change in warning design could be responsible for our lower observed clickthrough rates.

Ecological Invalidity. Sunshine et al. and Sotirakopoulos et al. recorded 55%-60% and 80% clickthrough rates, respectively, for a slightly outdated version of the Mozilla Firefox SSL warning [30, 31]. They evaluated the Firefox 3 and 3.5 warnings, which had the same layout and appearance as the current (Firefox 4+) warning but with different wording. It's possible that changes in wording caused clickthrough rates to drop from 55%-80% to 33.0%. However, during an exit survey, 46% of Sotirakopoulos's subjects said they clicked through the warning because they either felt safe in the laboratory environment or wanted to complete the task [30]. Since their study methodology was intentionally similar to the Sunshine study, Sotirakopoulos et al. concluded that both studies suffered from biases that raised their clickthrough rates [30]. We therefore attribute some of the discrepancy between our field study data and these two laboratory studies to the difficulty of establishing ecological validity in a laboratory environment.

In light of this, we recommend a renewed emphasis on field techniques for running and confirming user studies of warnings. Although we used in-browser telemetry, there are other ways of obtaining field data. For example, experience sampling is a field study methodology that asks participants to periodically answer questions about a topic [2, 6, 9, 28]. Researchers could install a browser extension on participants' computers to observe their responses to normally occurring warnings and display a survey after each warning. This technique allows researchers to collect data about participants' emotions, comprehension, and demographics. Participants may become more cautious or attentive to warnings if the purpose of the study is apparent, so researchers could obscure the purpose by surveying subjects about other browser topics. Network-based field measurements also provide an alternative methodology with high ecological validity. A network monitor could maintain its own copy of the Safe Browsing list and identify users who click through warnings. If the monitor can associate network flows with specific demographics (e.g., students), it can help understand the impact of these factors on user behavior. Similar studies could help understand SSL clickthrough rates; recent work addressed how to reproduce certificate validation at the network monitor [1].

7.3 Demographics

We found that clickthrough rates differ by operating system and browser channel. Our findings suggest that higher technical skill (as indicated by use of Linux and pre-release channels) may predispose users to click through some types of warnings. We recommend further investigation of user demographics and their impact on user behavior. Large-scale demographic studies might uncover additional demographic factors that we were unable to study with our methodology. If so, can warning design address and overcome those demographic differences?

Technically advanced users might feel more confident in the security of their computers, be more curious about blocked websites, or feel patronized by warnings. Studies of these users could help improve their warning responses.

7.4 Number of Clicks

Our data suggests that the amount of effort (i.e., number of clicks) required to bypass a warning does not always have a large impact on user behavior. To bypass Google Chrome's malware and phishing warnings, the user must click twice: once on a small "Advanced" link, and then again to "proceed." Despite the hidden button, users click through Google Chrome's malware/phishing warning at a higher rate than Mozilla Firefox's simpler warning. Furthermore, 84% of users who open Mozilla Firefox's "Add Exception" dialog proceed through it.

We find this result surprising. Common wisdom in e-commerce holds that extra clicks decrease clickthrough rates (hence, one-click shopping) [12, 32]. Google Chrome's warning designers introduced the extra step in the malware/phishing warning because they expected it to serve as a strong deterrent. One possible explanation is that users make a single cognitive decision when faced with a warning. The decision might be based on the URL, warning appearance, or warning message. Once the user has decided to proceed, additional clicks or information is unlikely to change his or her decision.

Our data suggests that browser-warning designers should not rely on extra clicks to deter users. However, we did not explicitly design our study to examine the effects of multiple clicks. Future studies on multi-click warnings could shed light on user decision models and impact security warning design. It is possible that extra clicks do not serve as a deterrent until they reach some threshold of difficulty.

7.5 Warning Fatigue

We observed behavior that is consistent with the theory of warning fatigue. In Google Chrome, users click through the most common SSL error faster and more frequently than other errors. Our findings support recent literature that has modeled user attention to security warnings as a finite resource [4] and proposed warning mechanisms based on this constraint [14].

Based on this finding, we echo the recommendation that security practitioners should limit the number of warnings that users encounter. Designers of new warning mechanisms should always perform an analysis of the number of times the system is projected to raise a warning, and security practitioners should consider the effects that warning architectures have on warning fatigue.

7.6 "More Information"

Users rarely click on the explanatory links such as "More Information" or "Learn More" (Section 5.2.4). Designers who utilize such links should ensure that they do not hide a detail that is important to the decision-making process.

Mozilla Firefox places information about SSL errors under "Technical Details" and in the "Add Exception" dialog instead of the primary warning. Thus, the error type has little impact on clickthrough rates. In contrast, Google Chrome places error details in the main text of its SSL warning, and the error has a large effect on user behavior. It is possible that moving this information into Mozilla Firefox's primary warning could reduce their clickthrough rates even further for some errors.

8 Conclusion

We performed a field study with Google Chrome and Mozilla Firefox's telemetry platforms, allowing us to collect data on 25,405,944 warning impressions. We find that browser security warnings can be successful: users clicked through fewer than a quarter of both browser's malware and phishing warnings and a third of Mozilla Firefox's SSL warnings. We also find clickthrough rates as high as 70.2% for Google Chrome SSL warnings, indicating that the user experience of a warning can have a tremendous impact on user behavior. However, warning effectiveness varies between demographic groups. Our findings motivate more work on browser security warnings, with particular attention paid to demographics. At Google, we have begun experimenting with new warning designs to further improve our warnings.

Acknowledgements

We thank the participants in Google and Mozilla's telemetry programs for providing us with valuable insight into our warnings. At Google, we would like to thank Matt Mueller for setting up the malware and phishing measurements, Adam Langley for making suggestions about how to implement SSL measurements, and many others for providing insightful feedback. At Mozilla, we would like to thank Sid Stamm for his mentorship and help collecting telemetry data, Dan Veditz for gathering data from Firefox 23, Brian Smith for providing information about the telemetry mechanisms, and the Mozilla contributors who reviewed our code and helped land this telemetry [22]. We also thank David Wagner, Vern Paxson, Serge Egelman, Stuart Schechter, and the anonymous reviewers for providing feedback on drafts of the paper.

References

[1] AKHAWE, D., AMANN, B., VALLENTIN, M., AND SOMMER, R. Here's My Cert, So Trust Me, Maybe? Understanding TLS Errors on the Web. In *Proceedings of the 2013 World Wide Web Conference* (2013).

[2] BEN ABDESSLEM, F., PARRIS, I., AND HENDERSON, T. Mobile Experience Sampling: Reaching the Parts of Facebook Other Methods Cannot Reach. In *Privacy and Usability Methods Powwow* (2010).

[3] BIDDLE, R., VAN OORSCHOT, P. C., PATRICK, A. S., SOBEY, J., AND WHALEN, T. Browser interfaces and extended validation SSL certificates: an empirical study. In *Proceedings of the ACM Workshop on Cloud Computing Security* (2009).

[4] BÖHME, R., AND GROSSKLAGS, J. The Security Cost of Cheap User Interaction. In *Proceedings of the New Security Paradigms Workshop (NSPW)* (2011).

[5] BRAVO-LILLO, C., CRANOR, L. F., DOWNS, J. S., AND KOMANDURI, S. Bridging the Gap in Computer Security Warnings: A Mental Model Approach. In *IEEE Security and Privacy* (March 2011), vol. 9.

[6] CHRISTENSEN, T., BARRETT, L., BLISS-MOREAU, E., LEBO, K., AND KASCHUB, C. A Practical Guide to Experience-Sampling Procedures. In *Journal of Happiness Studies* (2003), vol. 4.

[7] Google Chrome Privacy Notice. `http://www.google.com/chrome/intl/en/privacy.html`.

[8] CHROMIUM AUTHORS. HSTS Preload and Certificate Pinning List. `https://src.chromium.org/viewvc/chrome/trunk/src/net/base/transport_security_state_static.json`.

[9] CONSOLVO, S., AND WALKER, M. Using the Experience Sampling Method to Evaluate Ubicomp Applications. In *Pervasive Computing* (2003).

[10] Convergence. `http://www.convergence.io`.

[11] DHAMIJA, R., TYGAR, J. D., AND HEARST, M. Why phishing works. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (2006).

[12] DUTTA, R., JARVENPAA, S., AND TOMAK, K. Impact of Feedback and Usability of Online Payment Processes on Consumer Decision Making. In *Proceedings of the International Conference on Information Systems* (2003).

[13] 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 ACM CHI Conference on Human Factors in Computing Systems* (2008).

[14] FELT, A. P., EGELMAN, S., FINIFTER, M., AKHAWE, D., AND WAGNER, D. How to ask for permission. In *Proceedings of the USENIX Conference on Hot Topics in Security (HotSec)* (2012).

[15] FRIEDMAN, B., HURLEY, D., HOWE, D. C., FELTEN, E., AND NISSENBAUM, H. Users' Conceptions of Web Security: A Comparative Study. In *CHI Extended Abstracts on Human Factors in Computing Systems* (2002).

[16] HABER, J. Smartscreen application reputation in ie9, May 2011. `http://blogs.msdn.com/b/ie/archive/2011/05/17/smartscreen-174-application-reputation-in-ie9.aspx`.

[17] HERLEY, C. The plight of the targeted attacker in a world of scale. In *Proceedings of the Workshop on the Economics of Information Security (WEIS)* (2010).

[18] HOLZ, R., BRAUN, L., KAMMENHUBER, N., AND CARLE, G. The ssl landscape: a thorough analysis of the x.509 pki using active and passive measurements. In *Proceedings of the ACM SIGCOMM Internet Measurement Conference (IMC)* (2011).

[19] JACKSON, C., SIMON, D. R., TAN, D. S., AND BARTH, A. An evaluation of extended validation and picture-in-picture phishing attacks. In *Proceedings of the Workshop on Usable Security (USEC)* (2007).

[20] LANGLEY, A. SSL Interstitial Bypass Rates, February 2012. http://www.imperialviolet.org/2012/07/20/sslbypassrates.html.

[21] MCGRAW, G., FELTEN, E., AND MACMICHAEL, R. *Securing Java: getting down to business with mobile code*. Wiley Computer Pub., 1999.

[22] MOZILLA BUGZILLA. Bug 767676: Implement Security UI Telemetry. https://bugzil.la/767676.

[23] Mozilla firefox privacy policy. http://www.mozilla.org/en-US/legal/privacy/firefox.html#telemetry.

[24] NETCRAFT. Phishing on sites using ssl certificates, August 2012. http://news.netcraft.com/archives/2012/08/22/phishing-on-sites-using-ssl-certificates.html.

[25] PATERIYA, P. K., AND KUMAR, S. S. Analysis of Man in the Middle Attack on SSL. *International Journal of Computer Applications 45*, 23 (2012).

[26] PROVOS, N. Safe Browsing - Protecting Web Users for 5 Years and Counting. Google Online Security Blog. http://googleonlinesecurity.blogspot.com/2012/06/safe-browsing-protecting-web-users-for.html, June 2012.

[27] SCHECHTER, S. E., DHAMIJA, R., OZMENT, A., AND FISCHER, I. The Emperor's New Security Indicators. In *Proceedings of the IEEE Symposium on Security and Privacy* (2007).

[28] SCOLLON, C. N., KIM-PRIETO, C., AND DIENER, E. Experience Sampling: Promises and Pitfalls, Strengths and Weaknesses. In *Journal of Happiness Studies* (2003), vol. 4.

[29] SOBEY, J., BIDDLE, R., VAN OORSCHOT, P., AND PATRICK, A. S. Exploring user reactions to new browser cues for extended validation certificates. In *Proceedings of the European Symposium on Research in Computer Security* (2008).

[30] SOTIRAKOPOULOS, A., HAWKEY, K., AND BEZNOSOV, K. On the Challenges in Usable Security Lab Studies: Lessons Learned from Replicating a Study on SSL Warnings. In *Proceedings of the Symposium on Usable Privacy and Security* (2011).

[31] SUNSHINE, J., EGELMAN, S., ALMUHIMEDI, H., ATRI, N., AND CRANOR, L. F. Crying Wolf: An Empirical Study of SSL Warning Effectiveness. In *Proceedings of the USENIX Security Symposium* (2009).

[32] TILSON, R., DONG, J., MARTIN, S., AND KIEKE, E. Factors and Principles Affecting the Usability of Four E-commerce Sites. In *Our Global Community Conference Proceedings* (1998).

[33] WENDLANDT, D., ANDERSEN, D. G., AND PERRIG, A. Perspectives: Improving SSH-style Host Authentication with Multi-Path Probing. In *USENIX Annual Technical Conference* (2008).

[34] WHALEN, T., AND INKPEN, K. M. Gathering evidence: Use of visual security cues in web browsers. In *Proceedings of the Graphics Interface Conference* (2005).

[35] 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).

A Sample Sizes

	Malware	Phishing	SSL	Add Exception
Release	1,968,707	89,948	NC	1,805,928
Beta	74,782	3,058	10,976	66,694
Dev	61,588	2,759	15,560	53,001
Nightly	58,789	4,239	18,617	64,725

Table 7: Warning impression sample sizes for Mozilla Firefox warnings, by channel, for all operating systems.

	Malware	Phishing	SSL	Add Exception
Mac	71,371	3,951	534	154,129
Win	1,892,285	85,598	10384	1,634,193
Linux	1,750	112	58	17,606

Table 8: Warning impression sample sizes for Mozilla Firefox warnings, by operating system. The malware, phishing, and the "Add Exception" samples are from the release channel, whereas the SSL samples are from the beta channel. The frame issue does not affect statistics that pertain only to the "Add Exception" dialog.

	Malware	Phishing	SSL
Stable	5,946,057	381,027	16,363,048
Beta	44,742	3,525	232,676
Dev	14,022	1,186	66,922
Canary	35,261	612	42,020

Table 9: Warning impression sample sizes for Google Chrome warnings, by channel, for all operating systems.

	Malware	Phishing	SSL
Mac	598,680	20,623	947,971
Windows	9,775,104	333,522	13,399,820
Linux	15,456	577	515,319
Android	NC	NC	1,499,938

Table 10: Warning impression sample sizes for Google Chrome warnings, by operating system, for the stable channel.

In Google Chrome, we recorded 6,040,082 malware warning impressions, 386,350 phishing warning impressions, and 16,704,666 SSL warning impressions. In Mozilla Firefox, we recorded 2,163,866 malware warning impressions, 100,004 phishing warning impressions, and 45,153 SSL warning impressions. Tables 7, 8, 9, and 10 further separate the sample sizes based on OS and release channel.