Improving SSL Warnings: Comprehension and Adherence

Adrienne Porter Felt¹, Alex Ainslie¹, Robert W. Reeder¹, Sunny Consolvo¹,
Somas Thyagaraja¹, Alan Bettet¹, Helen Harris¹, Jeff Grimes²

¹Google, ²University of Pennsylvania

¹{felt, ainslie, rreeder, sconsolvo, somast, bettes, helenharris}@google.com, ²jeffgrimes@chromium.org

ABSTRACT

Browsers warn users when the privacy of an SSL/TLS connection might be at risk. An ideal SSL warning would empower users to make informed decisions and, failing that, guide confused users to safety. Unfortunately, users struggle to understand and often disregard real SSL warnings. We report on the task of designing a new SSL warning, with the goal of improving comprehension and adherence.

We designed a new SSL warning based on recommendations from warning literature and tested our proposal with micro-surveys and a field experiment. We ultimately failed at our goal of a well-understood warning. However, nearly 30% more total users chose to remain safe after seeing our warning. We attribute this success to opinionated design, which promotes safety with visual cues. Subsequently, our proposal was released as the new Google Chrome SSL warning. We raise questions about warning comprehension advice and recommend that other warning designers use opinionated design.

INTRODUCTION

Dissidents, drug dealers, and diplomats have one thing in common: they rely on SSL to help keep their online communication private. SSL protects their e-mails, tweets, and bank statements from eavesdropping or tampering in transit. When something goes wrong with a supposedly secure connection, most browsers alert the user with a warning. The user must then decide whether to adhere to or ignore the warning.

An ineffective SSL warning can cost users. For some, deciding to proceed under adverse conditions can lead to physical harm or imprisonment. For others, an attack in an urban coffee shop might lead to unauthorized credit card charges. We argue that an effective SSL warning ought to accomplish:

1. *Comprehension*. The user should be able to make an informed decision after seeing an SSL warning. He or she should understand the source of the threat, the data that is at risk, and the likelihood of a false positive warning.
2. *Adherence*. The warning should encourage users to act in a conservatively safe manner by not proceeding.

Comprehension is our preferred goal, but adherence is an alternative if and when we cannot achieve comprehension.

The literature cautions that SSL warnings may currently fail at one or both of these goals. Study participants mistake SSL warnings for messages about security updates, cookies, or malware [2, 10, 14, 16, 38]. This confusion leads to potentially harmful myths. For example, some people believe that their anti-virus software or operating system will protect them (it won't) [10, 38]. Furthermore, Google Chrome users adhered to only a third of SSL warnings [1].

We set out to design a new SSL warning that informs (or, failing that, convinces) users. We based our proposal on best practices drawn from warning literature. For comprehension, researchers advise that text should be simple, non-technical, brief, and specific (e.g., [6, 23, 24, 26, 38]). For adherence, the literature recommends promoting a clear course of action (e.g., [20, 38, 40]). Our proposal attempts to incorporate these suggestions. We hypothesized that the resulting proposal would improve comprehension and adherence.

We tested our proposal against three survey-based comprehension metrics, using 7,537 microsurvey responses. For comparison, we also tested SSL warning text from four browsers. Our proposed text slightly improved one aspect of comprehension (the threat source), but overall comprehension rates remained low for all warning texts. Why do all warnings — including ours — fail? Although we tried to follow best practices, we faced tradeoffs between contradictory advice. Our choices may not have been optimal. This suggests a need for more research into the relative importance of brevity, specificity, and non-technicality in security warnings.

Turning to our secondary goal of adherence, we ran a field experiment to measure how opinionated design affects adherence rates. *Opinionated design* is the use of visual design cues to promote a recommended course of action. Our proposal substantially increased adherence rates — by nearly 30 percentage points. This demonstrates that opinionated design can have a large impact on user safety and decision making. Following this experiment, our proposed SSL warning was released as the new Google Chrome 37 SSL warning. Adherence in the field subsequently increased from 37% to 62%, meaning that millions of additional users a month choose to act safely due to our warning design changes.

MOTIVATION: SSL WARNINGS

What Is SSL?

When a user visits a website over HTTPS, the browser tries to establish an SSL/TLS¹ connection to the website's server. SSL is supposed to ensure two properties: secrecy and authentication. *Secrecy* means that an eavesdropper should not

¹We henceforth say "SSL" for consistency with older literature.

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

CHI 2015, April 18–23, 2015, Seoul, Republic of Korea
ACM 978-1-4503-3145-6/15/04.

<http://dx.doi.org/10.1145/2702123.2702442>

be able to see or modify the e-mails and Tweets that a user sends over SSL. To this end, the browser encrypts the data that flows between the browser and a website's server. The browser also *authenticates* the server to make sure the server is not lying about its identity. Without authentication, a network attacker could pretend to be the server and thereby access the data. Authentication is needed to ensure secrecy.

Browsers display SSL warnings when the encryption is too weak or the server could not be authenticated. The connection is immediately halted, pending the user's decision about the warning. In some cases, the problem indicates a real attack. Syrian Internet users saw SSL warnings when the Syrian Telecom Ministry allegedly attacked Facebook users.² Similarly, SSL warnings alerted Chinese Internet users to attacks on Google³ and GitHub⁴. However, in other cases, misconfigured servers or firewalls cause spurious warnings.

Our Definition of Comprehension

To achieve our goal of comprehension, an ideal SSL warning would convey the following:

- *Threat source.* SSL warnings are about network attacks, not malicious or compromised servers. This means that the supposed attacker is at some point between the user's computer and the website's server. The threat does not stem from the website itself. An informed user might consider how much she trusts her local network connection and ISP. For example, a user might reasonably mistrust a WiFi connection at LaGuardia Airport — or any Internet connection in Syria. An informed user would not evaluate how benign or malicious the destination website is.
- *Data risk.* The user's data on the destination website is at risk of eavesdropping or tampering. Furthermore, this applies to all data already on the website, not just new data that the user enters after clicking through the warning. The user's local data and data on unrelated HTTPS websites are not at stake. For example, a user who faces a warning on a banking website should know that her banking statements might be read by someone else if she clicks through the warning. An informed user would therefore consider the sensitivity of her data on the destination website.
- *False positives.* Misconfigured websites and WiFi log-in screens cause spurious warnings in the absence of an attack. Websites with good security practices (such as e-mail providers and banks) are unlikely to be misconfigured. When weighing the likelihood of a false positive, an informed user would consider the website's reputation and whether the website normally works correctly.

Motivating Literature

Prior studies have shown that confusion about SSL warnings is widespread, and users overwhelmingly ignore some SSL warnings. However, experiments have shown that warning design changes can potentially help adherence.

² <https://www.eff.org/deeplinks/2011/05/syrian-man-middle-against-facebook>

³ <http://www.netresec.com/?page=Blog&month=2014-09&post=Analysis-of-Chinese-MITM-on-Google>

⁴ <http://www.netresec.com/?page=Blog&month=2013-02&post=Forensics-of-Chinese-MITM-on-GitHub>

Comprehension

SSL is a nuanced, technical topic that touches on other technical topics like network connections and authentication. Although people might understand that SSL relates to or provides security, they might not understand how.

Several studies have demonstrated that people conflate SSL warnings with other security topics. Interviewees told Bravo-Lillo et al. [10] that SSL warnings are about anti-virus software, security updates, or certifications for exemplary security practices. One interviewee described an SSL certificate as, "...like a credential or like a plug-in that allows you to use software, and it means your security is up to date on your computer." Dhamija et al. [14] asked participants to describe an SSL warning shortly after seeing it; participants said things like, "I accepted the use of cookies," and "It was a message from the website about spyware." Sunshine et al. [38] tested SSL warning comprehension with a survey, and some respondents said that SSL warnings are about viruses or worms. (Fewer than half of their respondents could accurately explain an SSL warning in their own words.) Researchers have also reported that their participants confused malware and phishing warnings with SSL warnings [2, 16].

This confusion leads some people to believe potentially harmful myths. As an example, one respondent told Sunshine et al. [38], "I use a Mac so nothing bad would happen" — perhaps true if the threat were malware, but incorrect and potentially dangerous for a network attack on SSL.

People who can reason about SSL false positives should be alarmed by SSL warnings on banking websites. However, Bravo-Lillo et al. [10] reported that six of their non-expert interviewees believe the opposite to be true: they said that SSL warnings could be ignored on banking websites because banks have good security practices. Some study participants similarly told Sunshine et al. [38] that they thought real attacks were less likely to occur on banking websites, although they were slightly more likely to heed warnings for banking websites than library websites in an experiment.

However, SSL is not impossible to understand. Biddle et al. [7] improved comprehension of Internet Explorer 7's certificate dialog boxes. They attributed their success to removing technical language and separating the concepts of secrecy and authentication. Although they did not study warnings, we hope the same success can be translated to warnings.

Adherence

Laboratory studies have consistently shown that participants ignore a majority of SSL warnings. In a laboratory phishing experiment, 68% of participants (15 of 22) clicked through an old Firefox SSL warning without stopping to read it [14]. (A click-through rate is the inverse of the adherence rate.) In another experiment, a third to half of participants clicked through Internet Explorer 7's SSL warning and entered credentials into a banking website [34]. Sunshine et al. exposed 100 study participants to SSL warnings during information lookup tasks; most participants (55% to 100%, depending on the warning and website) ignored the warnings [38]. Their study also showed that warning design could change click-

Figure 1. Our proposed warning. Default view on top, with the expanded “Advanced” section shown below.

through rates. Sotirakopoulos et al. replicated the Sunshine study with similar results, but 40% of their participants said during an exit survey that the laboratory environment had influenced their decision [37].

Field data is more optimistic about the ability of warnings to promote adherence. Large-scale field data from Mozilla Firefox showed that Firefox users adhered to nearly 70% of SSL warning impressions, which is a fairly good adherence rate [1]. Unfortunately, users adhered to only 30% of Google Chrome SSL warnings [1]. In a subsequent field experiment, Felt et al. found that differences in warning design accounted for a third to a half of the difference in behavior between browsers [18]. We briefly ran a version of the Firefox SSL warning in Google Chrome, and Google Chrome SSL warning adherence rates increased to 44%.

DESIGNING A NEW WARNING

Based on best practices from the warning literature, we proposed a new warning for Google Chrome. Our text is supposed to be simple, non-technical, brief, and specific, and the design is opinionated. Figure 1 shows our proposal. In this section, we describe how we arrived at the proposal.

Communicating the Threat

Technical Jargon

Warning researchers recommend using simple, non-technical language for warnings that are intended for a broad user base [6, 23]. People are more likely to read beyond the first sentence of a warning if it uses simple language [36]. Advertisements and warnings that contain technical language hold less interest and are less likely to be remembered or obeyed [3, 22]. For example, people are more likely to follow the simple instruction “open a window” than the more complex instruction “use in a well-ventilated room” [20].

SSL warnings have traditionally used technical terms such as “certificate” and “security credentials.” Both advanced and novice users find these terms confusing, and researchers recommend removing them [6, 7, 10, 38]. For example, a participant told Biddle et al. that “*I don’t know if my information is safe, because I don’t know what ‘encrypted’ means*” after seeing Internet Explorer 7’s certificate dialog box [7].

Warning	SMOG grade
Proposal	6.6
Google Chrome 36	11.0
Internet Explorer 11	10.5
Mozilla Firefox 31	8.7
Safari 7	8.0

Table 1. SMOG readability grades for different SSL warning texts.

Mozilla has removed more technical terms with each revision of the Firefox warning. In contrast, Google Chrome 36 and Internet Explorer 11 use many technical terms, respectively:

“...the server presented a certificate issued by an entity that is not trusted by your computer’s operating system.”

“The security certificate presented by this website was not issued by a trusted certificate authority.”

We removed all technical terms from the primary text of our proposal. Our decision to adhere to simple, non-technical language means that our proposed text is more general and less educational than other warnings. However, we hoped that the warning would be more accessible to a broad audience.

For curious or expert users, a more technical explanation can be read by clicking on “Advanced.” There, a secondary paragraph explains the cause of the error. Other SSL warnings similarly make use of secondary text. We do not expect many users to read the secondary text.

Reading Level

Following the recommendation to use simple language, we targeted a low reading level. Browsers are used by a broad audience. Newspaper articles — which have a similar audience — are typically written at a sixth grade reading level. While working on the text of our proposal, we tested it using the SMOG formula [30]. Table 1 shows the SMOG grades of our proposal and alternative warnings, using the title and primary text of each warning (see Table 2).

Brevity

Large quantities of text look like they will take effort to read, so people often read none of it [24]. Consequently, warnings should be as brief as possible [6, 7]. However, this presented a challenge: it is not possible to explain all aspects of the threat model in a single short paragraph. We therefore faced a tradeoff between brevity and comprehensiveness. We chose brevity, as we felt that conveying some of the topic was preferable to a user reading none of the text.

Specific Risk Description

People are more likely to comprehend and comply with a warning if it describes the risks explicitly and unambiguously [19, 26, 27, 40]. Warning researchers recommend providing specific, explicit, and comprehensive details about the consequences of ignoring a security warning [6, 10, 17].

This advice is at odds with the other recommendations for simplicity, non-technicality, and brevity. However, prior research demonstrates that all are important. We tried to accomplish both specificity and simplicity by providing a short list of example data types that could be stolen (“for example, passwords, messages, or credit cards”). This is more concrete than simply saying “information” or “data,” and we hoped it

Browser	Title	Primary text
Google Chrome 36	The site's security certificate is not trusted!	You attempted to reach example.com, but the server presented a certificate issued by an entity that is not trusted by your computer's operating system. This may mean that the server has generated its own security credentials, which Chrome cannot rely on for identity information, or an attacker may be trying to intercept your communications. You should not proceed, especially if you have never seen this warning before for this site.
Proposal (Chrome 37)	Your connection is not private	Attackers might be trying to steal your information from example.com (for example, passwords, messages, or credit cards).
Internet Explorer 11	There is a problem with this website's security certificate	The security certificate presented by this website was not issued by a trusted certificate authority. Security certificate problems may indicate an attempt to fool you or intercept any data you send to the server. We recommend that you close this webpage and do not continue to this website.
Mozilla Firefox 31	This Connection is Untrusted	You have asked to connect securely to example.com, but we can't confirm that your connection is secure. Normally, when you try to connect securely, sites will present trusted identification to prove that you are going to the right place. However, this site's identity can't be verified.
Safari 7	Safari can't verify the identity of the website "example.com"	The certificate for this website is invalid. You might be connecting to a website that is pretending to be "example.com", which could put your confidential information at risk. Would you like to connect to the website anyway?

Table 2. The text from different SSL warnings, for self-signed certificate errors.

would help users understand the data at risk. However, our text is less specific than warnings with technical terms.

Illustration

Illustrations can improve warning comprehension. Adding an illustrative symbol to a medical warning reduces the amount of time it takes to learn the information in the warning [31]. Furthermore, illustrations make warnings more attention-grabbing and memorable [13, 25, 28].

Existing SSL warnings use abstract illustrations that imply danger. Examples include a security guard (Mozilla Firefox 31) or a red shield (Internet Explorer 11). We followed suit and included an illustration of a red lock. Although this is abstract, it matches the security indicators used in the URL bar. Unfortunately, abstract symbols are more difficult to understand than concrete symbols [15, 21].

Risk Level

To gain user attention, warnings must stand out from the surrounding environment. Red warnings yield higher adherence rates and lower detection times than warnings printed in black or green [8, 28]. Sunshine et al. suggest that SSL warnings should be red to make them "scary" [38].

On the other hand, warnings need to communicate an appropriate level of risk. Several research studies have found that users confuse SSL warnings with phishing and malware warnings [2, 16, 38]. Since SSL warnings have much higher false positive rates, this confusion can harm the reputation of the more accurate phishing and malware warnings. ANSI recommends the use of signal words (e.g., "Danger," "Warning,") with associated colors (red, orange, yellow) for decreasing levels of risk [4]. Since we cannot determine the risk of physical harm for a particular user, none of the signal words are appropriate per ANSI standards.

We decided that color would still be appropriate for indicating risk levels without tying the warnings to physical risks. We therefore reserved a red background for malware and phishing warnings, which we consider higher-risk than SSL warnings. The SSL warning could have an orange or yellow background because it is lower risk but still a significant security threat. However, this presented a challenge: orange and yellow come close to failing our accessibility guidelines for contrast on a computer screen. At suitable levels of vibrancy,

yellow and orange lack sufficient contrast with either white or black text. We therefore prepared an SSL warning proposal with a gray background and red lock; although red is present, it is limited to an accent color.

Opinionated Design

We used visual design techniques to promote the safe choice as the preferred option. We call this *opinionated design*. Even without reading, the user should understand the instruction.

Clear instructions improve both comprehension and adherence rates [20, 40]. Simply providing information without a clear instruction does not necessarily influence behavior. For example, people don't always choose healthier products after reading nutritional labels [5]. Sunshine et al. recommend emphasizing a clear course of action in SSL warnings [38], and Microsoft's NEAT guidelines say that warnings should be clearly explained and actionable [32].

Choice Attractiveness

We wanted adherence to be a more visually attractive choice than non-adherence. The safe button is therefore a bright blue color that stands out against the background. Other Google properties use the same blue button style for primary actions, so people should associate the button style with a default choice. In contrast, the unsafe choice is a dark gray text link.

Choice Visibility

Several warnings hide the unsafe choice. The Google Chrome malware warning hides the "proceed" button behind an "Advanced" link, and users must click four times to proceed through the Mozilla Firefox SSL warning. These hurdles improve the adherence rate by 2–15%, depending on the complexity of each additional step [1, 18]. Similarly, Bravo-Lillo et al. found that adding small hurdles to installation dialogs improved participants' installation decisions [9, 11].

We hypothesize that the increase in adherence occurs due to a mix of factors. Finding the hidden choice requires effort, which can serve as a deterrent, and we believe that users view a hidden choice as not recommended by the browser. Furthermore, the additional click slightly increases the amount of time that users must spend looking at (and, ostensibly, thinking about) the warning [9, 11]. We view these mechanisms positively because they serve to convince undecided users.

However, hiding the link to proceed can have side effects. It increases the amount of effort to ignore a false positive, and some users might not realize that the hidden choice is available. For example, some participants asked Sunshine et al. to switch browsers because they found the Mozilla Firefox SSL warning too complex to proceed through [38].

We ultimately hid the link to proceed, but we tried to make its presence obvious. Figure 1 shows the default warning state; the user must click “Advanced” to reveal the link to proceed. Since “Advanced” is one of only two actions in the warning, we hope that users will try it and find the link to proceed before becoming frustrated.

COMPREHENSION

We tested whether the text of our proposed warning helps microsurvey respondents understand the threat model. We compared our proposed text to the SSL warning text from other major browsers (see Table 2).

Methodology

We created three survey-based comprehension metrics. Each survey consisted of a single comprehension question accompanied by a mock warning image. We collected a total of 7,537 Google Consumer Survey (GCS) responses.

Survey Questions

We tested each aspect of comprehension separately:

1. *Threat source.* The warning text should convey that a network attacker is the source of the threat. To measure this aspect of comprehension, we asked:

What might happen if you ignored this error while checking your email?
 - Your computer might get malware
 - A hacker might read your email

If respondents understood the warning text, they should have selected the e-mail answer choice. The answer order was randomized for display.

2. *Data risk.* The warning text should convey that a specific website’s data is at risk. We asked two versions of the following question:

If you ignored this error on {facebook.com, bankofamerica.com}, what information might a hacker be able to see?
 - Photos my friends have posted
 - Movies I have watched
 - Bank statements I have received
 - Documents on my computer
 - All of the above

Respondents should have been able to match the website to the data type: photos if asked about “facebook.com,” and bank statements for “bankofamerica.com.” The answer order was randomized, except for “All of the above.”

3. *False positives.* We also want users to believe that an SSL warning on a website with above-average website security practices is more — not less — likely to be a real attack. We asked two versions of the same question:

Figure 2. Mock warning images (our proposal and Google Chrome 36). We changed the URL appropriately for each question.

If you ignored this error while paying bills online, how likely is it that a hacker could see your bank account balance?

If you ignored this error while watching movies, how likely is it that a hacker could see what movies you have watched?

Each was accompanied by a five-point Likert-style scale, ranging from “Not at all likely” (left) to “Extremely likely” (right). Only the endpoints were labelled.

If respondents understood how false positives occur, they should have assigned a higher score to the question about banking than the question about movies.

Mock Warning Images

The survey questions were accompanied by mock warning images. We matched the URL in each image to the URL in the question it accompanied. Since we wanted to compare the text of multiple warnings, the images needed to clearly display the warning text without any distractions or confounds. To do this, we created a neutral, unbranded warning layout. We used the same layout with text from: our proposal, Google Chrome 36, Internet Explorer 11, Mozilla Firefox 31, and Safari 7. Figure 2 shows two examples of mock warning images.

The mock warning images contained the title and primary body text from each respective warning (Table 2). For all browsers, we chose the text that corresponds to a self-signed certificate error because it is the most common cause of SSL warnings [1]. Instead of varying the button text, each image had a blue “Ignore error” link to match the language used in our questions (“If you ignored this error...”). Wherever the warning texts included a specific browser name (e.g., “Safari”), we substituted the phrase “your browser.”

Incentives

Our GCS respondents answered these questions in lieu of a paywall. Each respondent had to answer a single question to proceed to the website. We chose to use GCS because they interrupt respondents on their way to websites, much like SSL warnings. GCSs are displayed on a variety of non-Google websites (e.g., news websites). We did not pay respondents, but they were able to gain free access to websites that normally cost money.

Sample Size

We requested a sample size of 300 for each of 5 variants and 5 conditions, although we occasionally received a few extra responses. We ended with a total of 7,537 survey responses.

Demographics

The GCS platform performs stratified sampling using inferred demographic and location information [29]. Their stratified sampling targets the most recent Current Population Survey (CPS) of Internet users [29]. Pew Research studied the GCS respondent population in 2012 and concluded that GCS respondents “conform closely to the demographic composition of the overall internet population” [33]. Sosik et al. found slight demographic differences but concluded that that technology usage and adoption was similar across GCS, Survey Sampling International (SSI), Knowledge Networks (KN), and Pew respondents [35].

We did not screen respondents for the threat source and data risk questions. However, we wanted to ensure that respondents for the false positive questions were familiar with online financial transactions. We screened respondents to the false positive questions with the same preliminary question: “Have you ever checked your bank account balance online?”

Limitations

Our respondents did not face any real danger. It is possible that people facing real warnings are more motivated to think about a warning, out of concern for his or her online safety. However, Google Chrome users decide within 2.1 seconds for half of SSL warning impressions in the field [1]. We therefore don’t believe that the lack of danger invalidates our findings.

Our samples per condition might not be completely independent. It is possible — albeit unlikely — that some respondents might have answered multiple questions. For example, a person might see multiple questions from the same set if (s)he answers surveys on multiple computers. Unfortunately, we can neither prevent nor identify this situation due to the platform’s anonymity. Given the low probability and our large sample size (at least 300 per condition), we treat our samples as independent when performing statistical testing.

We focused on comprehension amongst English-speaking U.S. Internet users. Our results may not translate to other cultures or languages. Once we achieve high comprehension rates in a single language, future work should investigate how to best translate that success to other countries.

Study Ethics

Our survey questions did not ask about sensitive topics, and we did not collect any personally identifiable information. The GCS platform anonymizes survey responses. The survey questions were reviewed by multiple researchers who have completed privacy and ethics training.

Results

Comprehension rates were low for the threat source and data risk metric for all warning versions. Our proposal improved the threat source metric, but we made no progress in conveying the data risk. However, respondents in aggregate were able to reason about false positives for all warning versions.

Threat Source

Our proposed warning text improved respondents’ understanding of the threat source. Table 3 shows how many respondents answered correctly per warning. Our proposal

Condition	“A hacker might read your email”	N
Proposal	49.2%	301
Chrome 36	37.7%	300
Safari	35.9%	304
Firefox	39.3%	300
Internet Explorer	39.3%	300

Table 3. Rate of correct responses to the threat source question. (“What might happen if you ignored this error while checking your email?”)

facebook.com						
Condition	Photos	Movies	Docs	Bank	All	N
Proposal	12%	6%	5%	2%	75%	302
Chrome 36	14%	7%	4%	2%	73%	303
IE 11	10%	7%	7%	3%	73%	302
Safari 7	13%	5%	5%	4%	73%	304
Firefox 31	11%	7%	6%	2%	74%	302

bankofamerica.com						
Condition	Photos	Movies	Docs	Bank	All	N
Proposal	3%	7%	8%	18%	65%	303
Chrome 36	6%	6%	9%	18%	62%	302
IE 11	4%	8%	8%	19%	51%	301
Safari 7	5%	8%	7%	14%	67%	301
Firefox 31	3%	5%	3%	20%	69%	302

Table 4. Responses to the data risk question. (“If you ignored this error on [website], what information might a hacker be able to see?”)

increased the chances of a correct response from 37.7% (Chrome 36) to 49.2% (proposal). Responses varied significantly across the warning versions ($\chi^2 = 13.43, p = 0.0093$).

Participants performed *worse than random chance* for all of the warnings except for our proposal. This suggests that the warning texts are counterproductive when describing the threat source. Although our proposal performed better than the alternatives, we still only achieved the same rate as random guessing. Substantial room for improvement remains.

Data Risk

Ideally, users should know specifically what data is at risk when they encounter an SSL warning. We provided a list of example data types (“passwords, messages, or credit cards”) in the hopes that it would help respondents draw this inference. However, respondents overwhelmingly overestimated the scope of the risk.

We asked respondents what type of data would be at risk on facebook.com or bankofamerica.com. As Table 4 shows, a majority of respondents selected “All of the above” instead of the more specific responses. The intended responses — photos for facebook.com, bank statements for bankofamerica.com — were a distant second in popularity. The warning text did not significantly affect participants’ ability to answer these questions precisely ($\chi^2 = 7.79, p = 0.955$; $\chi^2 = 22.72, p = 0.121$).

Some participants may have reasonably selected “All of the above” for the Facebook question because Facebook users can post movies in addition to photos. However, we do not believe that is why a majority of participants overestimated the risk. Participants could not reasonably make the same inference for Bank of America, yet participants responded similarly to that question.

Figure 3. Responses to the false positive questions. Banking (top) and watching movies (bottom). (“If you ignored this error..., how likely is it that a hacker could...?”)

False Positives

In prior work, interviewees told researchers that they thought SSL warnings were more likely to be false positives on banking websites [10, 38]. Although some individuals may hold this belief, our respondents in aggregate do not seem to. As Figure 3 shows, respondents rated an attack as more likely for a banking website than for a movie website. This relationship held true for all warning versions. We view this result with optimism, as it shows that participants have some ability to estimate the likelihood of a false positive.

Our proposal did not stand out from the alternatives. The warning version did not significantly affect responses to the banking question ($\chi^2 = 4.306, p = 0.9983$).

Discussion

None of the warning text that we tested — including our proposal — succeeded at our primary goal of comprehension for all three metrics. This suggests a need for future work on improving SSL warning comprehension.

We had hoped that following best practices would yield a dramatic improvement in comprehension. This did not happen, apart from a modest increase on the threat source metric. Why didn’t we see more of an improvement?

Not all best practices can be simultaneously satisfied, which meant that we had to make tradeoffs. In particular, it is challenging to design text that is brief, simple, and non-technical while also specific and comprehensive. We chose to emphasize brevity and simplicity, but the optimal warning text for comprehension might be longer or more technical. In future work, researchers could explore different points in this space (e.g., technical and brief vs. non-technical and long).

There are other recommendations that we did not pursue but might prove fruitful. One avenue is a more concrete illustration that explains an SSL warning. For example, one could illustrate an eavesdropper listening to a connection between a client and server. Another option is to make the warning highly specific to the situation, for example by displaying the user’s Facebook credentials in the warning if the error is for facebook.com.

ADHERENCE

Since we could not achieve high comprehension rates, we hoped to at least accomplish our secondary goal: guide users to act in a conservatively safe manner. We ran a large-scale, controlled field experiment in Google Chrome to measure how warning design affects adherence rates. Following this experiment, one of our proposals was released as the new SSL warning for Google Chrome.

Methodology

We deployed experimental SSL warnings on pre-release versions of Google Chrome and observed how they affected adherence. Figure 4 shows the experimental conditions.

Hypotheses

Our proposal is designed to promote the safe choice. In contrast, the Google Chrome 36 warning presents the safe and unsafe choices as visual equals. We hypothesized that the Google Chrome 36 warning ($Cond_A$) would have lower adherence rates than our proposal ($Cond_C$).

We wanted to test the influence of opinionated design, but we were faced with a confound: $Cond_C$ has both new text and a more opinionated design. To address this, we created $Cond_B$: our proposed new text, in the less opinionated Chrome 36 design. By comparing $Cond_A$ to $Cond_B$, we could see how much of the difference was due to the text.

We also experimented with the warning’s background color. Following ANSI standards, we initially planned to associate the SSL warning with a bright yellow color. However, this came close to failing accessibility guidelines for contrast on computer screens. We tested both yellow and gray backgrounds. If the yellow background ($Cond_D$) resulted in a higher adherence rate in the pre-release experiment, we would have adjusted the design to improve the contrast.

Experimental Platform

We pseudorandomly assigned a small fraction of pre-release Google Chrome users (Canary and Dev channels) to receive the experimental warnings. By default, pre-release users participate in a statistical reporting program to help identify regressions and test new features in Chrome. Their clients periodically send pseudonymous statistical reports to Google Chrome servers, including information such as whether the user clicked through any SSL warnings. Chrome tags the reports by experimental condition, allowing us to measure the adherence rate per condition. All of the conditions ran at the same time from June 24 – July 24, 2014. We limited the experiment to English.

Some release users also choose to opt in to statistical reporting. Once Google Chrome adopted $Cond_C$ as its new SSL warning, we used the statistical reports to measure how release users responded to the new warning. We measured adherence rates for the new SSL warning twice: from August 26 – September 20, 2014 (right after the release of Chrome 37), and from December 7, 2014 – January 4, 2015.

Study Ethics

A change to a warning carries a risk that the new warning might be less effective. However, users experience this same

Figure 4. Conditions for our field experiment. A is the Chrome 36 warning, and C is the Chrome 37 warning.

risk every time the Chrome team modifies a warning. We included only conditions that we thought would improve adherence, and we monitored the experiment closely to ensure it did not decrease adherence. If any of the conditions had decreased adherence, we would have halted the experiment.

Chrome's statistical reports are pseudonymous, and our experiment did not involve any personally identifiable information. The reports did *not* include any information about what websites triggered the SSL warnings. Users can opt out of statistical reporting at any time in Chrome settings.

Our experiment went through an internal experimental review process before it launched.

Limitations

We do not know whether the warning impressions were due to real attacks or false positives. Our overall goal is adherence, but we cannot calculate whether adherence is higher or lower during real attacks.

One possible experimental confound is the effect of novelty. Over time, users become accustomed to seeing — and disregarding — the same warning [12, 39]. Consequently, they may no longer pay much attention to it. A new warning can interrupt this rote behavior and gain user attention. It is therefore possible that our new warnings have a higher adherence rate simply because they are novel. However, SSL warnings are relatively rare events for most users; our statistical reports suggest that most users see one or fewer SSL warnings a month. Furthermore, we monitored the warning on release channels from August 2014 to January 2015. We did not see any change in adherence over time. This suggests that the improvement in adherence is not due solely to novelty.

Another consideration is the bias introduced by demographics. The release and pre-release versions of Google Chrome have different demographics. Early adopters seek out the pre-release versions of Chrome for early access to bleeding-edge features. As such, it is possible that they react differently to warnings than the general population. However, the adherence rates for $Cond_C$ were similar across pre-release and release versions. This suggests that the two populations might also react similarly to the other warning conditions, although we cannot be sure. We do not wish to test the other conditions on release users now that we believe the other conditions are more likely to result in users becoming the victims of attacks.

	Text	Design	Adherence	N
A	Original	Original	30.9%	4,551
B	Proposed	Original	32.1%	4,075
C	Proposed	Proposed (gray)	58.3%	4,633
D	Proposed	Proposed (yellow)	53.3%	4,528

Table 5. Adherence rates from the field experiment.

Results

Opinionated Design

Opinionated design substantially improved adherence. As Table 5 shows, our opinionated proposal ($Cond_C$) yielded a higher adherence rate than the older warning ($Cond_A$): 58.3% vs. 30.9%. This is a substantial increase of nearly 30 percentage points. On some days, that could amount to more than a million Google Chrome warning impressions. We attribute this improvement to the change in design and not to the change in text. The proposed text in the old design ($Cond_B$) increased the adherence rate by only 1.2%, which is a very small change.⁵

Following the experiment, $Cond_C$ was adopted as the new warning for Google Chrome. We monitored the adherence rate as the new warning rolled out, to ensure that release users also responded favorably. Indeed, the new warning improved adherence rates among release users who participate in statistical reporting. During the last month of Google Chrome 36, the adherence rate was 37% for 24,747,395 impressions. After the release of Google Chrome 37 in August 2014, the adherence rate increased to 62% for 20,214,251 impressions. Several months later, the adherence rate remained high at 61% for an additional 26,529,405 impressions.

Color Scheme

Contrary to our expectations, the yellow version of our proposal ($Cond_D$) performed worse than the gray version ($Cond_C$) by five percentage points (53.3% vs. 58.3%). We did not pursue the yellow background color further.

⁵A careful reader may note that the unsafe button is slightly larger in $Cond_B$ than $Cond_A$, due to differences in string lengths. This presents us with two possibilities: the text did not influence decision making, or did influence their decision making but the larger button size counteracted the effect. Either possibility supports the finding that opinionated design (using button appearance and visibility) increases adherence.

Discussion

We dramatically improved adherence rates with design cues that (a) promote the safe choice and (b) demote the unsafe choice. This demonstrates that opinionated design can substantially change how users react to browser security warnings. Our results therefore support prior researchers' recommendations [20, 38, 40]. Other warning designers may wish to adopt opinionated designs, if they have not already.

One remaining concern pertains to button visibility. Some users might miss the "Advanced" link and think it is impossible to click through the new warning. Some might switch browsers to proceed. We want to improve adherence rates by convincing users — not by confusing, frustrating, or tricking them. Unfortunately, we have no way to directly measure how many users grow frustrated or switch browsers. Instead, we have been monitoring the Google Chrome help center for complaints about the new design. Thus far, we have not seen complaints from users who could not find the link to proceed. However, warning designers should remain vigilant that excessively complex warnings might have this side effect. Other warning designers may want to consider less aggressive forms of opinionated design that do not completely hide the option to proceed.

To our surprise, our proposed text did not affect adherence rates even though it moderately improved threat source comprehension. We hope that future work will explore this topic further: how are comprehension and adherence related for SSL warnings? Within the same population, will improving comprehension increase (or decrease) adherence? Studying this question for the general population is different from separating users into expert and non-expert populations, as prior researchers have done [1, 10].

CONCLUSION

We proposed and evaluated a new SSL warning, which was released for Google Chrome 37. Our proposal dramatically improved adherence rates: from 31% to 58% in a controlled field experiment, and from 37% to 62% in the field following the release of the new warning. The increased adherence rate has held for more than four months after the release. Based on the results of our controlled field experiment, we attribute the improvement to our use of opinionated design. This demonstrates the potential power of opinionated design in helping users make safe security decisions.

Unfortunately, comprehension rates remain lower than desired for all of the SSL warning texts that we tested. This is disappointing, as we view comprehension as more important than adherence. We attempted to follow best practices for both adherence and comprehension, and we were surprised that this strategy yielded success for adherence but not comprehension. We attribute the low comprehension rates to the difficulty of creating an SSL warning that is simultaneously brief, non-technical, simple, and specific. Future work might improve comprehension by determining which pieces of advice are the most important. We urge readers to further pursue this line of research and develop understandable warnings.

GREETZ TO:

Chris Palmer, Glen Murphy, Ryan Sleevi, and Pariza Tabriz for their dedication to making security more usable; and to Ed Chi and Ulfar Erlingsson for their feedback on this project.

REFERENCES

1. Akhawe, D., and Felt, A. P. Alice in Warningland: A Large-Scale Field Study of Browser Security Warning Effectiveness. In *USENIX Security Symposium* (2013).
2. Almuhiemedi, H., Felt, A. P., Reeder, R. W., and Consolvo, S. Your Reputation Precedes You: History, Reputation, and the Chrome Malware Warning. In *SOUPS* (2014).
3. Anderson, R. E., and Jolson, M. A. Technical wording in advertising: implications for market segmentation. *Journal of Marketing* 44 (1980).
4. ANSI. Product safety signs and labels: Z535.4.
5. Arron, J., Egans, R., and Mela, D. Paradoxical Effect of a Nutrition Labeling Scheme in a Student Cafeteria. *Nutritional Research* 15 (September 1995).
6. Bauer, L., Bravo-Lillo, C., Cranor, L. F., and Fragkaki, E. Warning design guidelines (cmu-cylab-13-002).
7. Biddle, R., van Oorschot, P., Patrick, A. S., Sobey, J., and Whalen, T. Browser interfaces and extended validation ssl certificates: an empirical study. In *ACM Workshop on Cloud Computing Security* (2009).
8. Braun, C. C., Greeno, B., and Silter, N. C. Differences in behavioral compliance as a function of warning color. In *Human Factors and Ergonomics Society Annual Meeting* (1998).
9. Bravo-Lillo, C., Cranor, L., Komanduri, S., Schechter, S., and Sleeper, M. Harder to ignore?: Revisiting pop-up fatigue and approaches to prevent it. In *SOUPS* (2014).
10. Bravo-Lillo, C., Cranor, L. F., Downs, J., and Komanduri, S. Bridging the gap in computer security warnings: A mental model approach. *IEEE Security and Privacy* 9, 2 (2011).
11. Bravo-Lillo, C., Cranor, L. F., Downs, J., Komanduri, S., Reeder, R. W., Schechter, S., and Sleeper, M. Your attention please: Designing security-decision uis to make genuine risks harder to ignore. In *SOUPS* (2013).
12. Breznitz, S. Cry Wolf: The Psychology of False Alarms.
13. Bzostek, J. A., and Wogalter, M. S. Measuring visual search time for a product warning label as a function of icon, color, column, and vertical placement. In *Human Factors and Ergonomics Society Annual Meeting* (1999).
14. Dhamija, R., Tygar, J. D., and Hearst, M. A. Why phishing works. In *CHI* (2006).
15. Edworthy, J., and Adams, A. *Warning design: a research perspective*. Taylor and Francis, 1996.
16. Egelman, S., Cranor, L. F., and Hong, J. You've been warned: an empirical study of the effectiveness of web browser phishing warnings. In *CHI* (2008).

17. Egelman, S., and Schechter, S. The importance of being earnest [in security warnings]. In *Financial Cryptography and Data Security*, Lecture Notes in Computer Science. Springer Berlin Heidelberg, 2013.
18. Felt, A. P., Reeder, R. W., Almuhiemedi, H., and Consolvo, S. Experimenting At Scale With Google Chrome's SSL Warning. In *CHI* (2014).
19. Fischhoff, B., Riley, D., Kovacs, D. C., and Small, M. What information belongs in a warning? *Psychology Marketing* 15, 7 (1998).
20. Frantz, J. P. Effect of location and procedural explicitness on user process of and compliance with product warnings. *Human Factors* 36 (1994).
21. Frutiger, A. *Signs and symbols: their design and meaning*. Von Nostrand Reinhold, 1989.
22. Glover, B. L., and Wogalter, M. S. Using a computer simulated world to study behavior compliance with warnings: effect of salience and gender. In *Human Factors Society Annual Meeting* (1997).
23. Hancock, H., Bowles, C. T., Rogers, W. A., and Fisk, A. D. Comprehension and retention of warning information. *Handbook of warnings* (2006).
24. Hartley, J. *Designing instructional text*, 3 ed. Kagan Page and Nichols, 1994.
25. Kalsher, M. J., Wogalter, M. S., and Racicot, B. M. Pharmaceutical container labels and warnings: preference and perceived readability of alternative designs and pictorials. *International Journal of Industrial Ergonomics* 18 (1996).
26. Laughery, K. R., and Stanush, J. A. Effects of warning explicitness on product perceptions. In *Human Factors Society* (1989).
27. Laughery, K. R., and Vaubel, K. P. Explicitness in consequence information in warnings. *Safety Science* 16 (1993).
28. Laughery, K. R., Young, S. L., Vaubel, K. P., and Brelsford, J. W. The noticeability of warnings on alcoholic beverage containers. *Journal of Public Policy and Marketing* 12 (1993).
29. McDonald, P., Mohebbi, M., and Slatkin, B. Comparing Google Consumer Surveys to Existing Probability and Non-Probability Based Internet Survey. In *Google Whitepaper* (2012).
30. McLaughlin, G. H. Smog grading — a new readability formula. *Journal of Reading* (1969).
31. Morrow, D. G., Hier, C. M., Mendard, W. E., and Leirer, V. O. Icons improve older and younger adults' comprehension of medication information. *Journal of Gerontology: Psychological Sciences* 53B (1998).
32. Reeder, R., Kowalczyk, E. C., and Shostack, A. Poster: Helping engineers design NEAT security warnings. In *SOUPS* (2011).
33. Research, P. A Comparison of Results from Surveys by the Pew Research Center and Google Consumer Surveys.
34. Schechter, S. E., Dhamija, R., Ozment, A., and Fischer, I. The emperor's new security indicators. In *IEEE Symposium on Security and Privacy* (2007).
35. Schwanda-Sosik, V., Bursztein, E., , Consolvo, S., Huffaker, D. A., Kossinets, G., Liao, K., McDonald, P., and Sedley, A. Online Microsurveys for User Experience Research. In *CHI (Extended Abstracts)* (2014).
36. Silver, N., Leonard, D. C., Ponsi, K. A., and Wogalter, M. S. Warnings and purchase intentions for pest-control products. *Forensic Reports* 4 (1991).
37. 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 *SOUPS* (2011).
38. Sunshine, J., Egelman, S., Almuhiemedi, H., Atri, N., and Cranor, L. F. Crying Wolf: An Empirical Study of SSL Warning Effectiveness. In *USENIX Security Symposium* (2009).
39. Thorley, P., Hellier, E., and Edworthy, J. Habituation effects in visual warnings. *Contemporary ergonomics* (2001).
40. Trommelen, M. Effectiveness of explicit warnings. *Safety Science* 25 (1997).