Exploring User Reactions to Browser Cues for Extended Validation Certificates

Jennifer Sobey
School of Computer Science
Carleton University

Robert Biddle
Human-Oriented Tech. Lab
Carleton University

P.C. van Oorschot
School of Computer Science
Carleton University

Andrew S. Patrick
Inst. for Information Tech.
National Research Council

May 16, 2008

ABSTRACT

The ability of a user to reliably determine the true identity of a web site is important to online security. However, users often have trouble interpreting the browser security cues that are intended to help them with these decisions. With the introduction of Extended Validation SSL certificates in Internet Explorer 7.0, web browsers are introducing new indicators to convey status information about different types of certificates. We carried out a user study which compared a proposed new interface in the Mozilla Firefox browser with an alternative interface of our own design to investigate how users react to these new indicators. Our study included eye tracking data which provided empirical evidence with respect to which parts of the browser interface users tended to look at during the study and which areas went unnoticed. Our results show that, while the new interface features in the unmodified Firefox browser went unnoticed by all users in our study, the modified design was noticed by over half of the participants, and most users show a willingness to adopt these features once made aware of their functionality.

Keywords

Usable security, extended validation certificates, browser security, user study

1 Introduction

The ability of a user to reliably determine the true identity of a web site is important to online security. With the prevalence of phishing attacks, in which users are lured to fraudulent web sites, it is becoming increasingly important to provide users with effective tools to properly identify the true identity of a site. The use of certificates has traditionally been one way of providing identity information to the user, but studies have shown that many users have difficulty interpreting certificates or may not even be aware that they exist [2, 19].

With the introduction of Extended Validation (EV) SSL certificates [1], web browser software vendors are facing the design challenge of integrating support for these new certificates into their interfaces in a way that will be accepted and understood by users. Microsoft's Internet Explorer 7.0 was the first to introduce new interface features for Extended Validation which included a green background in the URL bar [6]. However, a preliminary study showed that these new visual cues did not provide a notable advantage for identifying a legitimate web site [8]. Other leading web browser vendors are currently working on plans to integrate support for Extended Validation in future releases [1].

After discussions with Mozilla developers [14], we decided to study the identity indicator being introduced in Mozilla's Firefox 3.0 browser. This interface includes a small clickable area to the left of the web site's address that produces a pop-up displaying information about the site certificate. The information displayed in the pop-up box indicates whether the site has an EV SSL certificate, a traditional SSL certificate, or no

certificate. We wanted to evaluate whether this interface would be effective in conveying identity information to the user and whether improvements could be made to make the indicator more effective.

We evaluated two different versions of the Firefox identity indicator – the version introduced in the Beta release of Firefox 3.0 and a modified version of this indicator that we designed, intended to better draw the user's attention. In a lab study, users interacted with both interfaces by performing tasks that required visiting an e-commerce web site and searching for several items they might purchase. Results were gathered by observation, questionnaire data, and by the use of an eye tracking device.

The remainder of the paper is divided as follows. *Section 2* provides a brief background on Extended Validation SSL certificates and summarizes related work in the area of web browser security. *Section 3* describes our user study methodology and the results we obtained. *Section 4* provides a further discussion of these results and the potential limitations of the study. *Section 5* contains our concluding remarks and ideas for future work in this area.

2 Background and Related Work

2.1 Extended Validation SSL Certificates

Extended Validation (EV) SSL Certificates are intended to provide improved authentication of entities who request server certificates for their web sites. These certificates build on the existing technology of the SSL certificate format but involve a more strictly defined certificate issuance process. A rigorous authentication process conducted by the EV Certification Authority (CA) is intended to allow visitors to a web site having one of these EV SSL certificates to have greater confidence in the site's identity. Whether this end-result turns out to be achievable remains an open question, relying on several factors including a suitable user interface for conveying trustworthy information to users. The guidelines for this certification process were established by the CA/Browser Forum, a voluntary organization consisting of CAs and Internet browser software vendors who support the new EV SSL standard [1].

Current support for EV SSL certificates relies on visual cues in the browser chrome – the frame of a web browser window that include menus, toolbars, scroll bars and status bars. As of March 2008, Microsoft's Internet Explorer 7.0 is the only browser to offer support for the EV SSL certificate in production software. When a user visits a web site having an EV certificate, the background of the browser's URL bar turns green and information regarding the web site owner and the issuing CA is displayed beside the padlock icon to the right of the address (see Figure 1) [11]. Mozilla Corporation, KDE, and Opera Software ASA are also members of the CA/Browser Forum and intend to provide EV certificate support in future releases of their software [1, 9, 13, 15].

2.2 Web Browser Security Indicators

One of the main challenges in the design of web browser security cues is the *unmotivated user property* noted by Whitten and Tygar [19]. Security is a secondary goal for most users; they are primarily focused on tasks such as checking email or browsing a web site. If security indicators are too subtle, many users will not be motivated to search for them or read manuals to learn their functionality. Conversely, if the user finds the security indicator too obtrusive there is a risk that the user will ignore security altogether, either because they become annoyed or they grow too accustomed to the indicator.

A lack of attention to security cues can result in users falling victim to phishing attacks. Dhamija, Tygar and Hearst [3] investigated why these attacks can be so effective and identified a number of factors that contributed to their success. Three groups of factors dealt directly with browser security indicators: (1) lack

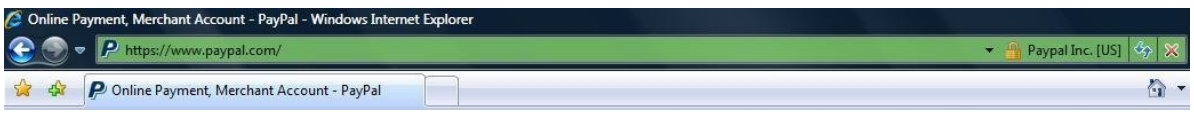

Figure 1: Internet Explorer 7.0's green URL bar for Extended Validation SSL certificates.

of knowledge of security and security indicators, (2) lack of attention to security indicators, and (3) lack of attention to the *absence* of security indicators. Even when these cues are actively being used, many users cannot reliably distinguish between a legitimate indicator and an attacker's image of one. Images placed in the content of a web page are often considered by users to be equally as trustworthy, since many users make no distinction between the page content and the chrome of a web browser [2].

2.2.1 The https Indicator

One indication of a secure connection to a web site is the placement of *https* in front of the address in the browser's URL bar. Several studies have shown that many users do not notice the presence or absence of the *https* indicator in a web site's address [3, 4, 16, 18]. One study by Schechter et al. [16] involved removing the *https* indicator and having users login to a banking web site. All 63 participants proceeded to enter their password and complete the task, despite the absence of the indicator.

2.2.2 The Lock Icon

In addition to *https*, secure connections are also indicated by the use of a lock icon located in the browser chrome. Its location varies depending on which browser is being used; the lock is often located either beside the address in the URL bar or in the bottom corner of the browser chrome. In several studies, this is the security indicator most often noticed [4, 18] but its absence often goes unnoticed [2]. Even when this indicator is used as a security cue by users, many do not fully understand its meaning [2, 3, 4].

Whalen and Inkpen [18] noted that while the lock metaphor alone may be a more powerful indicator of a secure connection than *https*, the icon is not being used to its full potential if there is no interaction with it. In browsers such as Internet Explorer and Firefox, the lock not only signifies a secure connection, but clicking on the lock icon results in the display of identity information based on the web site's certificate. The majority of users who do rely on this security indicator are not even aware of this identity feature [3, 4, 18] and do not reliably understand the concept of certificates at all [2, 3].

2.2.3 Extended Validation Indicators

Jackson et al. [8] performed an evaluation of the current EV certificate support in Internet Explorer 7.0 with respect to Picture-in-Picture phishing attacks. They found that the new security indicators had no significant effect on the users' ability to identify legitimate and fraudulent web sites, and reported that no one in the untrained group even noticed the new features. They do suggest that Extended Validation could become more useful in the future as users gain more awareness.

2.3 Browser Spoofing

When discussing the use of visual indicators to convey security and identity information, it is also necessary to consider how these indicators may be exploited by attackers. Felton et al. [5] describe a spoofing attack in which they were able to rewrite all of the URLs on a web page in order to direct users to an attacker site. They noted that their attack would be even more successful by overwriting the location and status bars using simple javascript so that the SSL indicators would appear as expected to the user. Ye et al. [20, 21] took this one step further by implementing an attack that removed the location and status bars provided by the browser and replaced them with their own. Since they had complete control over these new bars, they were able to spoof the traditional security indicators and even control the pop-up windows that displayed certificate information or security warnings.

Internet Explorer 7.0 has taken steps to help prevent these types of spoofing attacks. While the status bar can still be hidden, all windows (including pop-up windows) are required to display the location bar at the top. The developers of this browser have also placed all of the relevant security and identity indicators in the location bar, such as the lock icon and the green background for EV SSL certificates [12]. This makes it significantly more difficult for an attacker to overwrite the indicators in the location bar; they can no longer simply disable the default location bar and create their own. Restrictions such as this would be useful in all web browsers to decrease the likelihood of spoofed security indicators.

One attack that is no more difficult in this new IE 7.0 feature is the picture-in-picture attack, in which attackers make use of images, within the content of a web page, that mimic a browser window. Because of the similarity between the image and a legitimate browser window, the user can be fooled into thinking the site has simply opened a new window in front of the original [3]. Jackson et al. [8] acknowledge that without major changes to browser interface design, the only ways for users to identify these types of attacks are to notice which window has focus (two windows should not be in focus at once) or to try dragging or maximizing the window, and even these strategies are not fool-proof.

2.4 Use of Eye Tracking

Whalen and Inkpen [18] built upon the previous research on web browser security cues by incorporating eye tracking data into their evaluation. By tracking the user's gaze and fixation on the screen during the study tasks, they were able to obtain empirical results to cross-check what was reported by users with their actual behavior. There was very little variance between the visual cues that users reported using during the tasks and the data obtained from the eye tracker, but the tracking was also useful in identifying events that may otherwise have gone unreported, such as users looking for a padlock in the wrong location.

Another study by Kumar et al. [10] involved an eye tracker to implement a gaze-based password system that made use of the orientation of users' pupils to create passwords and authenticate to the system. The eye tracking data had a margin of error of 1° which resulted in some degree of inaccuracy, but despite this the error rates in their gazed-based password system were similar to those of passwords entered on a keyboard. These results support the use of eye tracking devices to reliably gather data on user gaze.

3 User Study

3.1 Implementation

3.1.1 Browser Interfaces

To evaluate the new identity indicators, we exposed participants to all three possible states of the indicator in both of the browser interfaces being studied (see Figure 2). The first browser used in the study was the Firefox 3.0 Beta 1 as proposed by Mozilla[1]. We refer to this browser interface as FF3 hereafter. The second browser was a modified Firefox 3.0 Beta 1, which we modified from the publicly available Beta code to insert our own identity indicator. We refer to this browser as FF3mod. In each of these browsers, the identity indicator had three possible states: (1) *identity unknown*, for web sites without SSL certificates or with self-signed certificates, (2) *location verified*, for web sites with traditional SSL certificates, and (3) *identity verified*, for web sites with EV SSL certificates[2].

A third browser was also included in the study as a control, giving a total of seven different interfaces. This consisted of the unmodified Firefox 2.0 browser (FF2) currently in circulation (circa March 2008), containing no support for EV SSL Certificates. Thus the user interface for this third browser in the study contained the traditional lock and *https* indicators but no additional identity indicators.

Because Mozilla had not yet implemented the functionality required to identify EV SSL certificates at the time of our build, we achieved the desired effect by building three separate versions for each of the two browsers – one for each state of the identity indicator. For FF3, the only distinguishing feature of the three versions of the browser was the information provided in the pop-up box of the identity indicator (see Appendix). In this browser, Mozilla developers buttonized the portion of the browser chrome to the left of the URL, which often contains a site's favicon, so that it would appear clickable to the user; clicking on this area would reveal the pop-up box for identity information.

We felt that this clickable indicator may be too subtle and go unnoticed by most users, so we designed FF3mod using a new identity indicator. Rather than buttonizing an existing feature in the browser chrome, we created an *identity confidence* button and displayed it in the same location to the left of the address bar.

[1]This was the current beta version in January, 2008

[2]The italicized names here are those assigned by Mozilla developers as identifiers for the three different SSL states. It is not our intention to explore, or tenure opinion, on the "true" level of security resulting from the use of the different types of certificates.

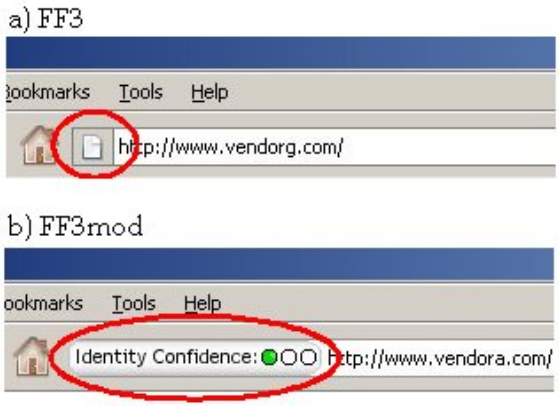

Figure 2: The image on the top shows the identity indicator used in FF3 (Firefox 3 Beta). The bottom image shows the identity indicator in our FF3mod (Modified Firefox 3 Beta).

The background of the button was colored white to provide a contrast against the dark gray chrome and contained an identity confidence meter consisting of three green lights. Web sites that had no certificate or had a self-signed certificate would have one green light lit up; two lights were lit on sites with traditional SSL certificates; and all three lights were lit for sites with EV SSL certificates (see Figure 3).

We chose to use one color for the lights rather than a traffic light metaphor for two reasons: (1) colorblind users may not otherwise be able to reliably distinguish between the different states; and (2) we did not feel it would be acceptable to produce red warning signals on a web site without a certificate, since many legitimate web sites simply do not offer secure connections. Similarly, we felt that a yellow signal for a web site with a traditional SSL certificate might falsely imply that the site may not be trustworthy. Other design considerations included catching a user's attention with the size and coloring of the button, and conveying some identity information on the button itself for users who chose not to click on it (or were not aware that it was clickable).

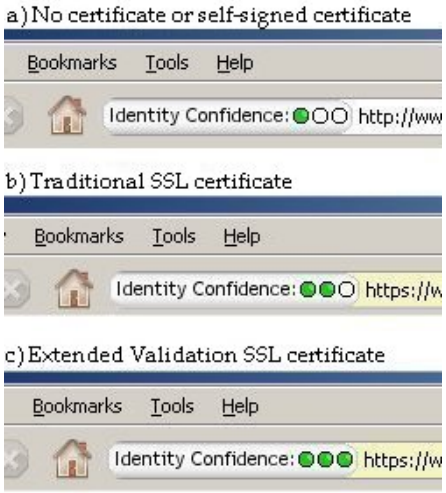

Figure 3: The *identity confidence* button in its three different states.

3.1.2 Other Technical Details

We were unable to use live web sites in our study as the EV functionality had not yet been fully implemented in Firefox 3.0 at the time of our evaluation. To provide the same experience as visiting live sites, we hosted the web sites used in our study on a Windows XP Professional machine using Apache 2.2.8. In order to emulate the correct behavior from the browsers, we created self-signed certificates for each web site to provide the information about the web site's identity to the browser interfaces. Despite the fact that self-signed certificates were used for all web sites in the study, each of the seven browser versions was hard-coded to display the appropriate SSL state no matter what type of certificate was used.

The web sites were based on a very simple design for an e-commerce web site selling computers, peripherals and accessories. All web sites were very similar in order to reduce biases introduced by the appearance of the web site. However, we created the illusion that the user was visiting seven different sites by changing the vendor name, the logo, and by interchanging product categories. This was intended to reduce the possibility that participants would dismiss the security cues once they believed they were interacting with the same web site for each task.

A Tobii 1750 eye tracker [17] set to a resolution of 1024 x 768 pixels at 96 dpi was used to capture and store data about each participant's gaze and fixation throughout the study. The stored data allowed playback of a recording of eye location on the screen and also captured the x and y co-ordinates of the each eye's location at intervals of 20 milliseconds. This device was located at the bottom of the monitor used by the participant for web browsing and captured eye movement as long as the user stayed within the range of the device. A second monitor was set up for the experimenter that displayed a real-time view of the eye tracking functionality. This allowed the experimenter to note any times where the user's gaze focused on the identity indicators and also provided a way to monitor that the eye tracker was functioning properly throughout the study. A calibration done at the beginning of the tasks ensured that the eye tracker device was configured correctly for each user.

3.2 Participants

A total of 28 participants took part in the user study. They were recruited through the use of an online campus recruiting system as well as posters displayed on campus. Sixteen were male and twelve were female, with ages ranging from 18 to 29. Twenty-four participants were undergraduate students and had a variety of majors and years of university education. Two participants were Computer Science and Engineering students who had high technical knowledge of computer security. With the exception of one other participant, the remaining users had relatively little computer security knowledge. Despite this, 21 out of 28 participants rated their concern for using their credit cards online as 8 or higher on a Likert scale from 1 (low) to 10 (high). All participants had made a purchase online in the past; 50% of participants reported making online purchases at least once per month. All participants browsed the Internet at least 5-10 hours per week, and consequently were very familiar with the use of a web browser. Microsoft's Internet Explorer was the web browser customarily used by 15 of the participants, 8 used Mozilla Firefox, 4 used Apple's Safari, and one participant reported using Netscape.

3.3 Tasks

Each participant in the study was asked to complete a 60 minute lab session. The participants were randomly assigned to one of two groups, with each group having the same distribution of gender, age, and education. Before proceeding to the tasks, participants in Group 1 were informed that the study's purpose was "to evaluate different web browsers and web sites that could be used for Internet shopping". This was not intended to deceive the participants in any way, but to ensure that there was no specific focus on security so that they would not be influenced to act any differently than they normally would. Participants in Group 2 were provided with the same purpose statement but were also told that "we are interested in such things as visual appearance, item pricing, amount of contact details, trust in the site's authenticity, and ease of use". The purpose of the additional information was to evaluate whether the subtle reference to trust in the site's authenticity would influence the participant to focus more on the identity indicators. All other aspects of the study were identical for both groups.

After the introduction, the participant performed a sequence of seven tasks. Each task involved the following steps:

1. Read a brief description of three items to be located on a web site.
2. Double-click on a desktop icon corresponding to the task number to open one of the web sites within one of the seven browser interfaces.
3. Locate the three requested items on the web site and record the price of each on a sheet provided.
4. Answer a series of two questions: (1) on a 10-point Likert-scale, "How willing would you be to make purchases on this web site with your own credit card?" and (2)"What factors did you use in making your decision?"

The order of presentation of the browser interfaces, web sites and tasks were counterbalanced using spatially balanced 7x7 latin squares [7] to avoid bias created by the order in which the independent variables are presented. Once all seven tasks were completed, a follow-up interview was conducted in which participants were asked for their opinions regarding the web browsers used in the study and whether or not they noticed the various identity indicators being evaluated. Finally at the end of the lab session, each participant filled out a questionnaire used to collect demographic information.

3.4 Results

Each participant completed all seven tasks, giving us a total of 196 tasks from which to draw data. The results were analyzed based on both qualitative data (observation of the participant's behavior during the study, post-task questionnaires and interviews at the end of the session) and quantitative data (gathered by the eye tracker and the post-task questionnaire).

3.4.1 Self-Reported Attention to the Identity Indicators

We were able to determine which identity indicators were noticed by observing the participants during the study and by reviewing their responses to the follow-up interview. Our results showed that the identity indicator introduced in the FF3 web browser went unnoticed by all of the participants in our study, regardless of the group condition. Because the indicator was not even noticed, no one attempted to click on this indicator and therefore no one saw the pop-up information box that distinguished between the three certificate levels (see Appendix).

Of the fourteen participants in Group 1 (those given minimal instructions), six reported noticing the FF3mod *identity confidence* indicator while performing the tasks. This same indicator was reported to be noticed by nine participants in Group 2 (the group given enhanced instructions). Of these fifteen participants who reported noticing the FF3mod *identity confidence* indicator, seven reported seeing it on at least two different interfaces. Five participants were unsure of how many times they had seen this indicator, while the other three said they only noticed it once near the end of their tasks as they became more observant of the browser features.

All seven participants who reported noticing the FF3mod *identity confidence* indicator at least twice while performing their tasks also reported noticing the different states of the indicator. Three of these participants immediately caught on to the meaning of the indicator and actively used this indicator when making decisions about their willingness to transact with the web sites. The participants who did not use the indicator in their decision-making dismissed it, stating reasons such as "I don't understand what it means" or "I just assumed all of the web sites were the same". None of these participants made any attempt to interact with (click on) the *identity confidence* button and therefore did not see the pop-up information box at any point.

During the follow-up interview, participants were explicitly shown the two different browsers (FF3 and FF3mod) and the identity indicators that were evaluated in the study and were asked which they would prefer to use at home if given the option. The FF3mod browser with the *identity confidence* button was chosen by twenty-two of the twenty-eight participants (78.6%). When asked why they would choose this option, participants gave reasons such as the indicator being more eye-catching and easier to notice, and

Figure 4: A screenshot of the eye tracker replay function. The large circle near the *identity confidence* indicator shows the participant's fixation on that region of the screen.

the fact that it provides some identity information without having to click on the button. Most felt the unmodified FF3 version was too subtle. The four participants who preferred FF3 stated that they liked the fact that it took up less space in the chrome but commented that they would need to somehow be made aware that it existed. One participant had no preference for either indicator, and one other clearly stated they preferred the traditional lock icon to either of these identity indicators.

3.4.2 Objective Measures of Attention to Identity Indicators

The results obtained with respect to participants' self-reported attention to identity indicators were verified with the eye tracker data. The eye tracker allowed us to replay each session in order to visually analyze times at which the user may have looked at the indicators. The replay screen portrays a moving blue dot that signifies the user's gaze; the larger the dot becomes, the longer the user has fixated on that region of the screen (see Figure 4). We were also able to analyze data files that recorded the x and y co-ordinates of the gaze at intervals of 20 milliseconds to determine times at which the participant's gaze was fixated on the indicator's co-ordinates.

The eye tracker data confirmed that the fifteen participants who reported noticing the FF3mod *identity confidence* indicator throughout the tasks did in fact fixate on the co-ordinates where the button was displayed for an average of 1.1 seconds at a time. Data from the participants who did not report noticing the *identity confidence* indicator showed that if their gaze did fall on the co-ordinates of interest, it was only for approximately 0.26 seconds at most.

In addition to the identity indicators being studied, seven participants also reported using the traditional indicators (the lock icon or *https*) to help make decisions about identity and trust. The eye tracker data confirmed that these users did in fact fixate their gaze on the appropriate co-ordinates throughout the seven tasks. There were also four participants who did not report using the traditional indicators in their decision-making but whose gaze fixated on their co-ordinates during most tasks.

One of the more interesting findings in the eye tracking data was how long users spent gazing at the content of the web pages as opposed to gazing at the browser chrome. On average, the fifteen participants who gazed at traditional indicators, new identity indicators, or both, spent only about 9.5% of time gazing at any part of the browser chrome. The remaining thirteen participants who did not gaze at indicators spent only 4.3% of their time focusing on browser chrome as opposed to content. While other studies have found that many users are unable to distinguish between web page content and chrome, ours suggests they do distinguish between the two and that they rarely glance at the chrome at all. This finding was also supported by the participants' comments during the follow-up interview. When the identity indicators were pointed out to participants, many made comments such as "I didn't even think to look up there" or "I was

only focusing on the web page itself."

3.4.3 Willingness to Transact

There was a wide range of answers to our Likert-scale question, "How willing would you be to make purchases on this web site with your own credit card?" Nine participants assigned the same rating across all seven browser interfaces, basing their decisions solely on visual appearance and professionalism (which was kept relatively constant across all seven web sites).

We took the mean of all ratings assigned to each browser interface and found the numbers to be consistent with what we would expect. A significant overall effect of interface on the ratings was found by performing an Analysis of Variance (ANOVA)[3] on the data ($F(6,156)= 4.09$, $p<.001$). There was no significant difference in ratings between the two groups ($F(1,26)=.52$,$p<.48$). There was also no interaction between the group condition and the interface. Post hoc tests were conducted to determine whether there were any significant pairwise differences among the means using a Tukey HSD test[4]. There was a significant difference in the means between the FF3 non-SSL interface and the FF3mod EV-SSL interface, as well as between the FF3mod non-SSL interface and both the FF3 EV-SSL and FF3mod EV-SSL interfaces. There were no significant differences between non-SSL and SSL interfaces or SSL and EV-SSL interfaces. Since the FF2 control interface was not rated differently than any other interface, we chose to remove this condition from further analysis.

A second ANOVA was performed to compare the two browser conditions with the three different states of each browser. There was no interaction found between the factors of browser and state, and no significant difference between the browsers ($F(1,27)=0.40$, $p<.53$). There was however a significant difference in SSL state ($F(2,54)=6.03$,$p<.005$). We followed up these results with a Tukey HSD test and found the significant difference to be between the non-SSL and EV-SSL states. There were no significant differences between non-SSL and SSL states, nor between SSL and EV-SSL states.

With the eye tracking data, we were able to classify participants as "gazers" or "non-gazers." Participants who were considered to be gazers looked at either the traditional security indicators (lock icon, *https*), the FF3mod *identity confidence* indicators, or both, during each task. There were eleven participants classified as gazers in the study. All other participants were classified as non-gazers, regardless of what they reported looking at during the study. By making this distinction, we were able to identify that participants who look at SSL indicators (either traditional lock and *https* or the new identity indicators) de-value non-SSL connections and assign higher ratings to web sites with SSL or EV SSL certificates (see Figure 5); but in our study, less than 40% of participants were gazers.

To verify the difference in ratings assigned to non-SSL and SSL connections, we performed ANOVAs on the data. Among non-gazers, as expected, there was no significant difference in ratings across SSL state ($F(2,32) = 1.61$, $p < .22$). However, there was a very significant difference in ratings across SSL state among the gazers ($F(2,20) = 6.32$, $p < .008$). A Tukey HSD test was used to verify the differences among the various SSL states among gazers; there was a significant increase in mean ratings from non-SSL(3.41) to SSL(5.50) interfaces, as well as from non-SSL(3.41) to EV SSL(5.95) interfaces. The increase from SSL to EV SSL interfaces was not significant. Figure 6 gives an overall picture of the ratings between gazers and non-gazers for all browser versions and SSL states.

In addition to analyzing these ratings with respect to gazers vs. non-gazers, we also compared the three users who reported using the FF3mod *identity confidence* indicator in their decision-making with other gazers who did not. Participants who used the FF3mod *identity confidence* indicator in their decision-making assigned a mean rating of 8.33 to the EV SSL interfaces, 6.50 to interfaces with SSL, and 3.83 to interfaces with non-SSL connections. Although these differences appear large, the small number of participants who made use of this new indicator in their decision-making prevented a meaningful statistical comparison.

[3]As is well known, an ANOVA is a statistical method used to make simultaneous comparisons between two or more means; the values can be tested to determine whether a significant relation exists between variables.

[4]The Tukey Honestly Significant Difference test is a method of multiple comparisons that test for a significant difference between a pair of means based on rankings from smallest to largest

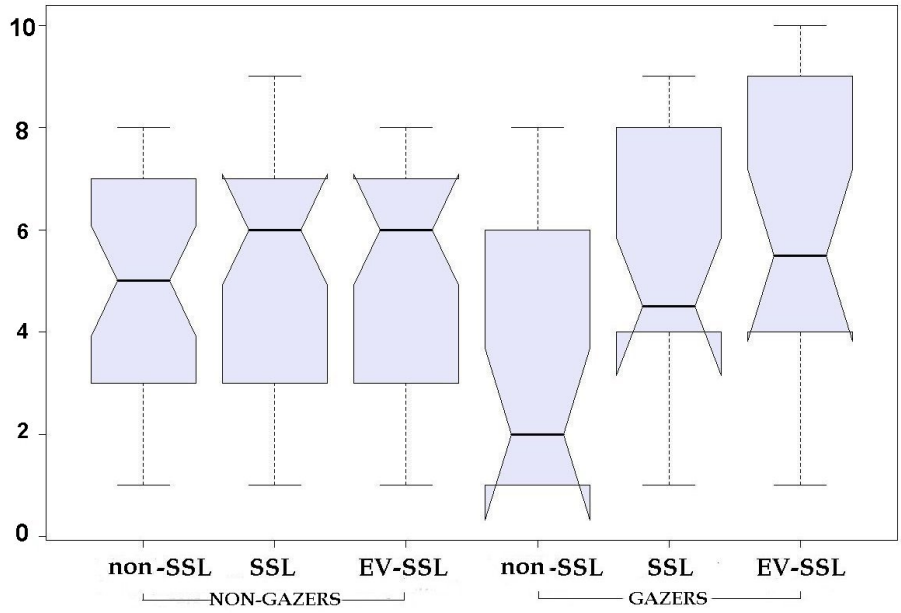

Figure 5: Boxplot of participants' mean *willingness to transact* ratings based on SSL state, grouped by gazer or non-gazer. Horizontal lines of each box represent 25th, 50th and 75th percentile ratings, with the actual distribution shown by the lines that extend from the box. Note the small variance in ratings among non-gazers compared to a larger variance by gazers.

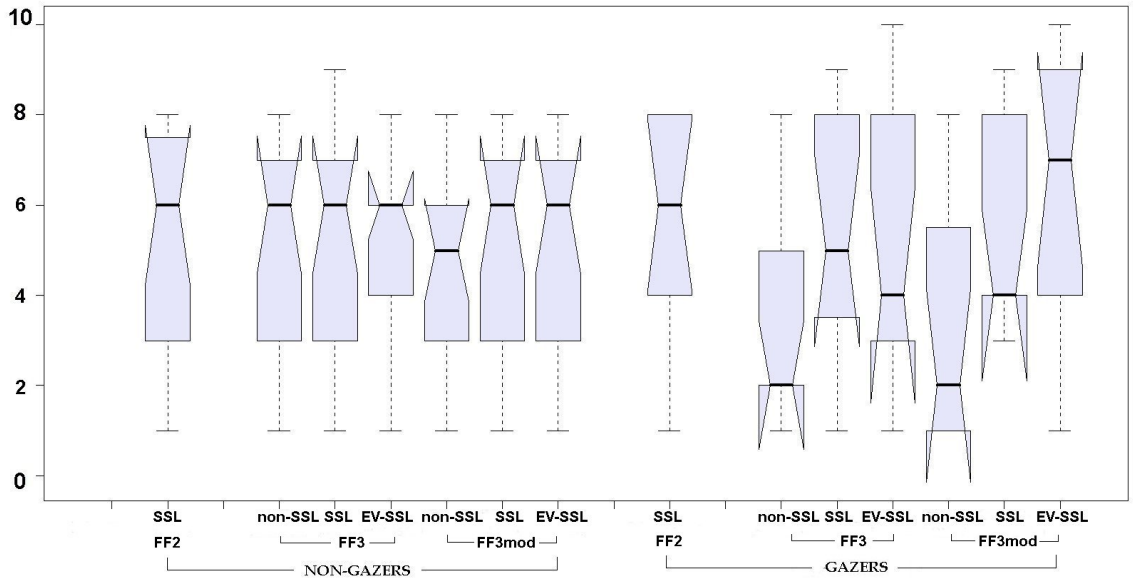

Figure 6: Boxplot of participants' mean *willingness to transact* ratings based on Browser and SSL state, grouped by gazer or non-gazer.

3.4.4 Decision Strategies

In the post-task questionnaires and in the participant information questionnaire at the end of the session, participants were asked what factors they use in making decisions about whether or not to trust a web

site. Twenty of the participants in the study reported the visual appearance and professionalism of the web site as a factor in their decision-making. Other common factors included the amount of detail on the website, the contact information or location of the company, and their familiarity with the website. Nine participants reported using the lock icon and/or *https* indicator to make their decisions (which we were able to reliably verify using the eye tracker) and, as mentioned earlier, three participants used the FF3mod *identity confidence* indicator in their decision-making. Several participants mentioned looking for security logos or security and privacy statements in the content of the web site itself.

The participant information questionnaire at the end of the study identified several other factors that participants tend to use during their normal web browsing but which were irrelevant for our particular study. These included strategies such as relying on the company's reputation, noticing regular updates to the web site, searching the web for user reviews of web sites, asking friends and family for recommendations, contacting the company by telephone to ensure it exists, or contacting the Better Business Bureau to inquire about the organization. It is possible the absence of these features from the web sites used in our study may have uniformly lowered participants' willingness to transact.

4 Discussion

4.1 Extended Validation Indicators

Our results showed that the identity indicators used in the unmodified FF3 browser did not influence decision-making for the participants in our study in terms of user trust in a web site. These new identity indicators were ineffective because none of the participants even noticed their existence. Had they known that this clickable area existed beside the browser's URL bar, they would have been able to distinguish between the three SSL states by clicking on that area and seeing the pop-up information box. Since this functionality was not discovered, the indicators were of no value to the user. The differences in ratings based on state for this browser can only be explained by the use of the traditional lock icon and *https* indicators.

While many participants also disregarded the new FF3mod *identity confidence* indicator or did not notice it at all, it was promising to note that three participants did make use of it in their decision-making and seemed to understand its meaning immediately. This supports the idea that users may be able to reliably make use of such indicators to evaluate web site identity. In addition to these three users, twelve others reported noticing the *identity confidence* indicator but did not report using it in decision-making, possibly because they did not fully understand its purpose. Many participants reported noticing the indicator late in the study after most of the tasks were completed; this suggests that as users are given more exposure to the new indicators, they may be more likely to take notice of them.

It is also interesting to note that, while there were significant differences in ratings given to non-SSL interfaces vs. SSL or EV SSL interfaces, there was no significant difference in ratings given to SSL vs. EV SSL interfaces. Even among the three participants who reported using the FF3mod *identity confidence* indicator in their decision-making, only one participant gave notably different ratings to the FF3mod SSL and the FF3mod EV SSL interfaces. The other two participants also used the traditional lock icon or *https* to aid in their decisions and thus assigned high levels of trust to all SSL interfaces. Since we did nothing to educate participants on the differences between SSL and EV SSL, and they had no background knowledge in the area, it is not surprising that ratings given to these interfaces did not differ greatly. If the goal of EV SSL certificates is to give users a higher level of confidence in a web site's identity than traditional SSL certificates, we believe that users will need to be better educated on the different levels of identity indicators.

4.2 User Attention to Browser Chrome

We have seen in previous studies that most users may not be able to distinguish between web page content and browser chrome [2]. With the use of eye tracking data, our study also showed that many users spend very little time looking at any parts of the browser chrome. This presents an important challenge when it comes to incorporating security cues into web browsers; any content provider can trivially modify the content of a website to include security information. This problem is amplified by the fact that many users actually look for security information in the page content of a website. During our study, several users mentioned they had looked for security logos within the websites or looked for statements on the payment pages regarding

the security of their credit card information. These types of security cues could be effortlessly incorporated into an attacker's web site and many users would evidently be fooled by this type of technique.

While elements of the browser chrome can also be spoofed [5, 20, 21], it is more work for the attacker. It becomes even more difficult when browsers such as Internet Explorer 7.0 place restrictions on which parts of the window can be hidden, but it is still not impossible. However, in order to provide identity indicators that can best aid users in identifying web sites, designers need to place these identity cues in the chrome. Two main open questions remain: (1) how can users be persuaded that the elements of the chrome are worth looking at; and (2) how can it be ensured that users can distinguish a legitimate indicator from a spoofed indicator?

4.3 Design Implications

The fact that most users tend to ignore the browser chrome suggests that designers need to somehow find a way to draw visual attention to any security cues provided by the browser. We attempted to do this with our FF3mod *identity confidence* indicator by making it larger than the original FF3 indicator and using a color contrast to the browser chrome surrounding it. However, this was still not enough to get half of our study participants to take notice of it. We feel that better techniques for drawing user attention to important security indicators are needed, especially if these indicators are meant to be intuitive for the user. (Of course, parties responsible for other buttons in the chrome likely feel similarly about the importance of their buttons unrelated to security). However, this is also dangerous advice since attackers can counter this by finding ways of spoofing these parts of the chrome. The design of these indicators should be done in a way that makes it much more difficult for attackers to replicate. Mozilla developers [14] attempted to do this by having the identity indicator's pop-up window overlap slightly with the location bar (see Appendix), but we believe this is unlikely to be noticed by most users.

Another important design issue to note was the "clickable" feature of both the FF3 and FF3mod indicators. Not one participant in our study clicked on any of the indicators, even those who did notice and use the FF3mod *identity confidence* indicator. We designed our FF3mod indicator to have rounded edges and shading in an attempt to make it appear button-like, however this failed to cause users to click on this button. Perhaps more shading or a different shape would have been more effective. It is also possible that including action words, such as "click here" might have had more of an effect, but clearly it seems unreasonable for every clickable button to be so annotated.

4.4 Limitations of the Study

One of the major limitations of our study was the fact that it was conducted in a laboratory setting rather than in the field. This may have led to participants acting differently than they normally would in their own environments. Some participants may have felt more secure than during their normal web browsing because it was a university setting, while others may have paid more attention to security because of the more formal setting. The eye tracker may have also influenced people to behave differently since they were aware that their eye movements were being recorded. However, the eye tracker provided us with valuable data for our analysis and this was the main reason for the use of a laboratory setting; it would not be realistic to expect participants to install eye trackers in their home environments for the purposes of the study.

The fact that the tasks involved recording prices rather than following through with financial transactions may have also influenced participants to be less concerned with security. This effect was balanced by asking them questions after each task regarding their willingness to transact with the web site; these questions were intended to draw their attention to security issues.

Another potential limitation of our study was participants' lack of familiarity with the various components of the study. Twenty of the twenty-eight participants did not use Mozilla Firefox as their usual web browser. The novelty of an unfamiliar browser may have distracted participants because not only were the identity indicators new to them, but so was the overall look and feel of the browser window. The concept of EV SSL certificates is also relatively new, so we expect many users are not even aware that they should be looking for cues relating to the certificate types. As users gain more knowledge of EV SSL certificates, they may become more likely to use the types of identity indicators used in this study to make decisions about online security.

5 Conclusion and Future Work

While the introduction of Extended Validation SSL certificates was intended to help users make informed decisions regarding the identity and authenticity of a web site, our study shows that the unmodified Firefox 3.0 browser cues fail to effectively convey this information, at least in the absence of additional user training or awareness. By introducing a modified design of the Firefox 3.0 browser, we were able to increase the number of users who reported noticing an identity indicator to fifteen (over 50% of the study participants) and observed three users who showed immediate understanding of the indicator. However, to have users take notice of this new *identity confidence* button, we were forced to use more valuable space in the browser chrome. Regardless of the size of the indicator, many users tend to look to the content of the web site for security information rather than the browser chrome. This presents a challenge for browser interface designers who wish to provide to the user intuitive identity cues that will not go unnoticed.

A natural extension of our study is to evaluate user reactions to the indicators as a function of users being given increasingly more information before the study tasks. The hope would be to have more participants notice the identity indicators so as to better evaluate their understanding and interpretations of the various states and their preferences for each indicator. A future field study would also be interesting to measure behavior over time as users become more aware of the EV SSL features to see whether these indicators would continue to aid them in their decision-making or whether they would eventually be dismissed.

One feature of the indicator that we were unable to study was the information in the pop-up box triggered by clicking on the indicator (none of the participants attempted to interact with the indicators), and users' interpretation of the information presented by these boxes. Another natural aspect to study is the effect of the particular wording of this pop-up box as well as its behavior in the browser. Having the browser display a message pointing out the new features of this box might successfully draw the user's attention to the identity indicator. Until users are aware that identity indicators exist in the browsers and are able to effectively interpret their meaning, we believe that Extended Validation SSL certificates will have little effect on online security.

6 Acknowledgments

We thank Johnathan Nightingale from Mozilla for discussions, insight and technical advice regarding the Firefox 3.0 beta. We also thank Tim Moses for his background and insight on Extended Validation certificates, Sonia Chiasson and Alain Forget for their help with various aspects of the user study, and Rachna Dhamija.

References

[1] CA/Browser Forum. `http://www.cabforum.org/`.

[2] R. Dhamija and J. Tygar. The battle against phishing: Dynamic security skins. In *Proceedings of the Symposium on Usable Privacy and Security (SOUPS '05)*, 2005.

[3] R. Dhamija, J. Tygar, and M. Hearst. Why phishing works. In *Human Factors in Computing Systems (CHI 2006)*, April 22-27 2006.

[4] J. S. Downs, M. Holbrook, and L. Cranor. Decision strategies and susceptibility to phishing. In *Proceedings of the 2006 Symposium on Usable Privacy and Security*, July 2006.

[5] E. Felton, D. Balfanz, D. Dean, and D. Wallach. Web spoofing: An internet con game. In *Proceedings of the 20th National Information Systems Security Conference*, 1996.

[6] R. Franco. Better website identification and extended validation certificates in IE7 and other browsers. `http://blogs.msdn.com/ie/archive/2005/11/21/495507.aspx`.

[7] C. Gomes, M. Sellmann, C. Van Es, and H. Van Es. Computational methods for the generation of spatially balanced latin squares. `http://www.cs.cornell.edu/gomes/SBLS.htm`.

[8] C. Jackson, D. R. Simon, D. S. Tan, and A. Barth. An evaluation of extended validation and picture-in-picture phishing attacks. In *Proceedings of Usable Security 2007 (USEC '07)*, 2007.

[9] K Desktop Environment. `http://www.kde.org`.

[10] M. Kumar, T. Garfinkel, D. Boneh, and T. Winograd. Gaze-based password entry. In *Proceedings of the 2007 Symposium on Usable Privacy and Security*, 2007.

[11] Microsoft. Extended validation SSL certificates. `http://www.microsoft.com/windows/products/winfamily/ie/ev/default.mspx`.

[12] Microsoft. Internet Explorer 7.0 features. `http://www.microsoft.com/windows/products/winfamily/ie/features.mspx`.

[13] Mozilla. EV-Certs for Firefox. `http://mozillalinks.org/wp/2007/05/ev-certs-for-firefox/`.

[14] J. Nightingale. Personal Communication, September 19th, 2007.

[15] Opera Software. `http://www.opera.com`.

[16] S. E. Schechter, R. Dhamija, A. Ozment, and I. Fischer. The emperor's new security indicators. In *Proceedings of the 2007 IEEE Symposium on Security and Privacy*, May 2007.

[17] Tobii Technology AB. `http://www.tobii.com`.

[18] T. Whalen and K. Inkpen. Gathering evidence: Use of visual security cues in web browsing. In *Proceedings of Graphics Interface 2005*, pages 137–145, May 2005.

[19] A. Whitten and J. D. Tygar. Why Johnny can't encrypt: A usability case study of PGP 5.0. In *Proceedings of the 8th USENIX Security Symposium*, August 1999.

[20] E. Z. Ye, Y. Yuan, and S. Smith. Web spoofing revisited: SSL and beyond. Tech. Rep. Department of Computer Science, Dartmouth College, TR2002-417.

[21] Z. Ye, S. Smith, and D. Anthony. Trusted paths for browsers. *ACM Transactions on Information and System Security*, pages 153–186, May 2005.

APPENDIX

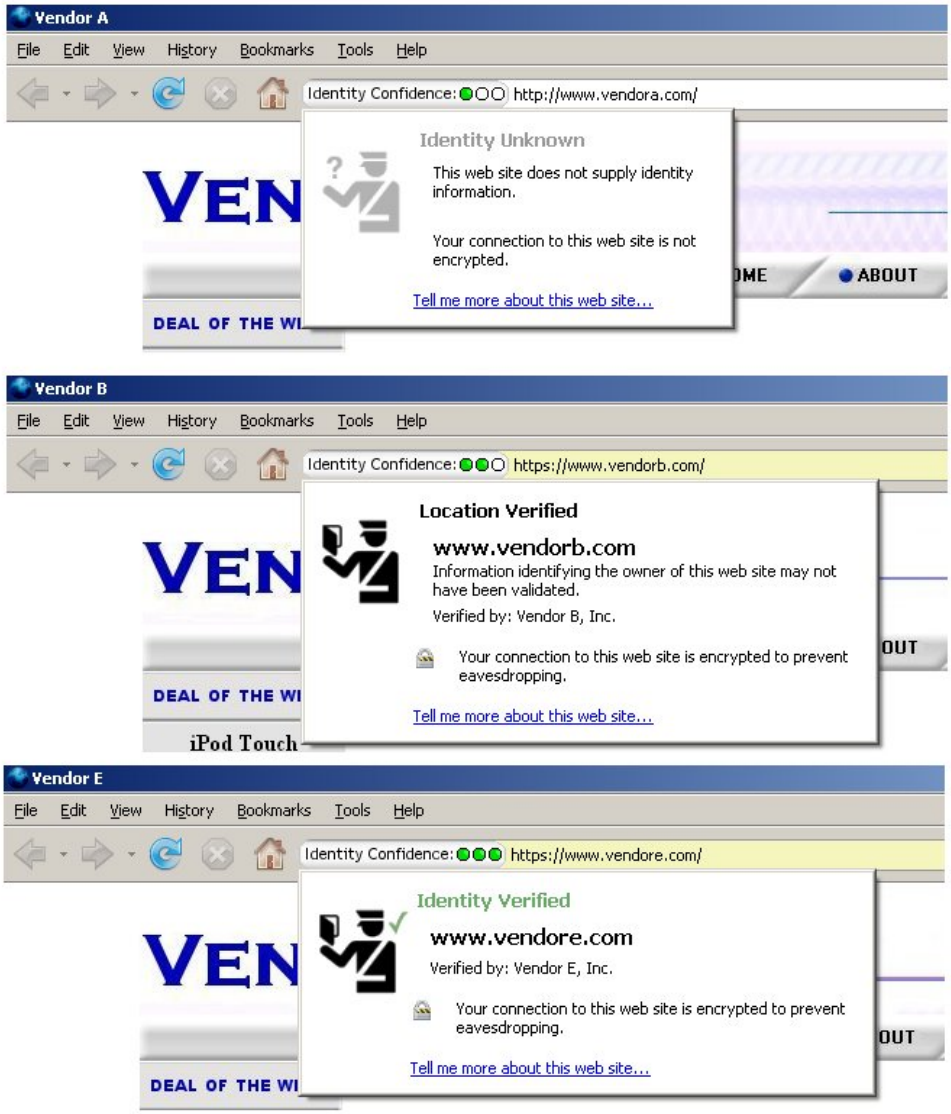

Figure 7: Text boxes corresponding to different states of Firefox 3.0 Beta 1. The information drop down box appears when users click on the identity indicator. The box on the top is for web sites with self-signed certificates or no certificate, the middle box is for web sites with traditional SSL certificates, and the bottom box is for web sites with an EV SSL certificate. The images shown here display the text box with the FF3mod *identity confidence* indicator, but the boxes are identical to those in Firefox 3.0 Beta 1.