Insights into User Behavior in Dealing with Internet Attacks

Kaan Onarlioglu
Northeastern University
Boston, MA
onarliog@ccs.neu.edu

Utku Ozan Yilmaz
Bilkent University
Ankara, Turkey
uyilmaz@cs.bilkent.edu.tr

Engin Kirda
Northeastern University
Boston, MA
ek@ccs.neu.edu

Davide Balzarotti
Institute Eurecom
Sophia Antipolis, France
balzarotti@eurecom.fr

Abstract

The Internet is a lucrative medium for criminals targeting Internet users. Most common Internet attacks require some form of user interaction such as clicking on an exploit link. Hence, the problem at hand is not only a technical one, but it also has a strong human aspect. Although the security community has proposed many technical solutions to common attacks, the behavior of users when they face current threats, and the way they evaluate the security implications of their actions remain largely unexplored.

In this paper we describe an online experiment platform we built for testing the behavior of users when they are confronted with prevalent, concrete attack scenarios such as reflected cross-site scripting, session fixation, and file sharing scams. We conducted experiments with 164 Internet users with diverse backgrounds. Our findings suggest that many non-technical users can exhibit performance comparable to security experts at averting relatively simple threats that they are frequently exposed to in everyday life. They can do so solely by following their intuition, without actually perceiving the severity of the threat. However, when facing more sophisticated attacks, these non-technical users often rely on misleading cues such as the "size" and "length" of artifacts (e.g., URLs), and hence, fail to protect themselves. We also show that trick banners that are common in file sharing websites and shortened URLs have high success rates of deceiving non-technical users, thus posing a severe security risk.

1. Introduction

The Internet has become a critical infrastructure, and any disruption in services adversely affects our lives and causes significant damage (e.g., the recent Amazon EC2 cloud outage affected several Fortune 500 companies and millions of users [21]). Clearly, as the importance of an information medium increases, so does its attractiveness for criminal activity with the aim of making quick, illegal financial gains. In fact, because of their high popularity and a user base that consists of millions of Internet users, web applications have become primary targets for attackers. According to SANS [36], attacks against web applications constitute more than 60% of the total attack attempts observed on the Internet. Many web applications are exploited every day to convert trusted websites into malicious servers hosting client-side browser exploits. Once the victim's machine has been infected with malware, the attackers then start collecting sensitive information such as credit card numbers and passwords. According to SANS, most website owners fail to scan their applications for common flaws. In contrast, from an attacker's point of view, automated tools designed to target specific web application vulnerabilities simplify the discovery and mass infection of websites.

Although some types of attacks are technically difficult for users to detect and prevent (e.g., stored cross-site scripting attacks [27] on a popular social networking website), most Internet attacks actually require user interaction (e.g., clicking on an exploit link, installing risky software, failing to recognize a phishing website, ignoring an SSL certificate warning, etc.). Hence, the user often becomes the weakest link in the chain, and the attackers often rely on social engineering techniques to trick victims into engaging in risky behavior, thus compromising their security.

To date, the security community has proposed many technical solutions to mitigate current Internet threats such as botnets (e.g., [15, 31, 32]), malware (e.g., [16, 39, 62]), cross-site scripting (e.g., [61]), cross-site request forgery (e.g., [17]), and drive-by download exploits (e.g., [47]).

However, it is clear that the problems at hand are not only technical, but they also involve a strong human aspect as some form of user interaction is typically required for many of these attacks to be successful.

In [24], Dhamija et al. attempted to understand which phishing attack strategies work better in practice and why. The paper provided the first empirical evidence on which malicious strategies are successful at deceiving general users by conducting experiments with 22 users. Recently, Sunshine et al. [52] presented an empirical study of SSL warning effectiveness, which showed that users do not react to the warnings as expected and they often exhibit dangerous behavior. Based on the lessons that they were able to learn, the authors conducted experiments with 100 users and designed new warnings that performed significantly better than the SSL certificate warnings used in browsers today.

Note that while the literature is rich in studies on general security usability and the human decision making process, besides several papers that focus on phishing attacks and warning effectiveness, there have not been any previous work on how users are able to cope with current threats such as cross-site scripting, session fixation and file sharing scams. This paper presents the first empirical findings that shed light on how different user groups deal with, and react to, the aforementioned attacks. In this study, we do not try to identify the techniques used by attackers to trick their victims. Instead, we investigate the problem from the users' perspective, and determine how users evaluate the security implications of their own actions.

Our results suggest that, even though non-technical users sometimes have a wrong understanding of security risks and of what constitutes an attack, when confronted with threats that they are frequently exposed to, they can still successfully mitigate them based on previous experience. Our findings also identify critical artifacts (e.g., the length of a URL) that users pay particular attention to when making security judgments.

We empirically confirm the general intuition that security education has a significant effect in preventing the more complex Internet attacks, and that a general "security awareness" is critical for user protection. Finally, we believe that online test systems such as the one we have constructed are useful in educating users about popular attacks on the Internet.

This paper makes the following contributions:

- We introduce an online security test system that presents to the users 44 typical benign and malicious scenarios, and records their behavior. We have conducted empirical experiments with a diverse set of 164 users. To the best of our knowledge, the study we present is the largest that has been conducted to date on common Internet attacks such as cross-site scripting, session fixation and file sharing scams.
- We show that many non-technical users can successfully mitigate the common attacks (such as email scams) that they are frequently exposed to, even when they cannot assess the severity of the threat. On the contrary, more advanced attacks (e.g., session fixation) are still only detected through security expertise. We also show that users with a security education are better at assessing the consequences of a possible threat.
- We provide empirical evidence that many users treat "length" and "size" as a sign of maliciousness (e.g., length of URLs and size of files).
- We show that non-technical users are frequently tricked by shortened URLs, and are largely not aware of simple web-security tools and services available for expanding shortened URLs.
- We provide empirical evidence that trick banners that are common in file sharing websites have a high success rate of deceiving users, and, therefore, pose an important security threat.

The rest of the paper is structured as follows: In Section 2, we summarize the related work. In Section 3, we present our experiment platform and give details of each test we performed on the participants. In Section 4, we show the results we obtained from the tests and list our observations. In Section 5, we discuss these results, and summarize the insights we distilled. Finally, in Section 6, we briefly conclude the paper.

2. Related Work

Although attacks such as cross-site scripting, session fixation and social engineering-based malware are common on the Internet, there have not been empirical studies specifically designed to determine the awareness level of Internet users about these attack vectors, and how they are able to deal with such attacks in practice.

As mentioned before, one well-known work that attempts to understand why phishing strategies work was conducted by Dhamija et al. [24]. Jackson et al. [37] investigated the effectiveness of extended validation certificates, and in a recent work, Lin et al. [42] conducted a user study to determine whether domain highlighting techniques help users identify malicious websites.

Another set of related efforts investigate user reaction to security warnings. Sunshine et al. [52] presented an empirical study of SSL warning effectiveness. Egelman et al. [26] studied the effectiveness of active and passive phishing warnings. In [44], Maurer et al. presented the first results of using data type-based alert dialogs for increasing security awareness of users. Raja et al. [46] designed and evaluated firewall warnings based on a physical security mental model.

Note that these studies focused on single specific attacks. In comparison, our study covers general web attacks, as well as email and file sharing scams.

Friedman et al. [28] conducted a number of general interviews about web security and concluded that many participants could not reliably determine whether a connection is secure. The participants were shown screenshots of a browser connecting to a site and they had to decide if the connection was secure or not. In another study [29], Friedman et al. interviewed participants about their concerns on risks and potential harms of web use. In [48], Schechter et al. evaluated some of the common website authentication mechanisms. In comparison, our work focuses on concrete technical attacks such as session fixation, cross-site scripting, and malicious links provided by URL shortening services. We report our findings on how users behave when confronted with such realistic attacks.

A recent work by Conti et al. [22] suggested a taxonomy of malicious interface techniques. The authors conducted a survey on a group of users to measure their frustration and tolerance when they encounter such interfaces. However, this study does not discuss the effectiveness of such techniques at deceiving users.

Research has also focused on understanding the decision-making process in security-critical contexts. Cranor [23] proposed a framework for reasoning about the human-factors involved in a security-critical system and identifying possible failure points. Herley [33] examined why users reject security advice and suggested that most security advice have a poor cost-benefit trade-off.

Other researchers have attempted to measure the effectiveness of social engineering attacks in social networks. For example, in [38], Jagatic et al. performed realistic phishing attacks on undergraduate students based on the information they were able to harvest from social networking websites. In another work, Bilge et al. [18] were able to show that users tend to have a higher level of trust in messages they receive from their social networks.

Finally, there exist several studies on the usability of security solutions. For example, in [19], Chiasson et al. described a usability study they had conducted on 26 users which shows that some previously proposed security solutions have serious usability problems. In [20], Clark et al. presented another study on the usability of anonymous web browsing. Wu et al. [60] evaluated the effectiveness of anti-phishing toolbars. Motiee et al. [45] conducted a study to determine whether users apply the User Account Control implemented in the Windows Vista and Windows 7 operating systems correctly. Ho et al. [34] surveyed users of home wireless networks to determine whether they are aware of the security features available to them.

Previous attempts at understanding the human factor in security focuses on understanding how and why certain attack techniques trick users, or they study the usability of proposed security solutions. In contrast, in this work, our focus is on determining and understanding how users react to threats, how they evaluate the security implications of their decisions, what cues they use to this end, and how they assess the risks involved. Moreover, our test platform has a broad scope, incorporating many popular and concrete attack scenarios that have not been studied in this context before.

3. Design of the Experiments

In this section, we describe the setup of the system we developed to test and simulate the typical security threats that users may encounter in their everyday Internet usage. In order to be able to reach a large number of users with diverse backgrounds, we conducted online experiments using an interactive test platform.

We designed the experiments as a within-subject study in which all participants responded to a series of tests in various security contexts. After studying a wide range of common Internet attacks that require some sort of user interaction or decision, we created 44 security-related scenarios, grouped in three test suites: web-based attacks, email-based attacks, and file sharing-related attacks. The tests included both malicious and benign scenarios, distributed in a random fashion. Each of the malicious scenarios exemplified different attack techniques in order to prevent the participants from building knowledge along the way and performing better in subsequent tests.

After reaching the homepage of our online test platform, we informed the participants that they were going to take part in "an experiment to determine the security-awareness of Internet users" and we asked them to provide an email address. We also informed the participants that the tests would take about an hour.

In order to mitigate the negative effects of motivational confounds and prevent inaccurate results due to loss of concentration, we gave participants the option to leave after completing any number of tests, and come back again later to continue from where they left off. As we explain in the following sections, most of the tests asked the participants to briefly explain the reasoning behind their decisions. By manually checking these responses and logging the times spent on each test, we ensured that the participants did not rush to the end of the tests by answering the questions arbitrarily but gave sufficient thought to their decisions. Apart from their email addresses, which we used to uniquely identify the participants when they wanted to continue the tests later, we did not ask for any personally identifiable information.

We recruited the participants through announcements on Twitter and Facebook, and by directly asking people in non-technical disciplines (e.g., medicine, and geology) to publicize our test platform URL. We did not offer a financial incentive to the participants; instead, we aimed at maximizing the task performance and obtaining accurate results by promoting our experiments as an opportunity for the participants to test their security knowledge and get feedback on their performance. In our recruitment strategy, we focused on having a balanced number of participants with technical and non-technical backgrounds in order for us to make meaningful comparisons between the performances of these groups. After eliminating the data from 5 respondents who had a poor command of English, or who completed the tests in an unreasonably short amount of time by giving arbitrary responses to the questions, we completed our study with 164 participants. However, the file sharing test suite was available only to the participants that reported previous experience with BitTorrent or with one-click-hosting services (91 and 97 participants, respectively). In the following sections, we describe in more detail the security tests we conducted.

3.1. Demographic Information

Before starting the tests, we collected standard demographic information, as well as inquiring the participants about their computer and Internet usage habits. In this way, we ensured that all the participants were reasonably familiar with basic computer terminology and daily tasks such as surfing the Internet and reading emails, in order for them to understand and respond to our tests correctly.

In addition, since we were interested in observing the effects of technical background on the results of the experiments, we asked the participants questions to estimate their security proficiency. Specifically, we asked them if they are comfortable with doing everyday tasks using their computers, if they have previous programming experience, if they are professionally involved in software/hardware development, if they have a degree in computer science or a related technical field, and if they have specialized computer security expertise.

3.2. Web-Based Attacks

The first test suite presented the participants with various URLs, and asked them to rate the "risk they perceived" for each link. That is, the participants were asked to rate how dangerous or safe they believed it would be to click on the links. The risk perception ratings were expressed in the 5-point Likert scale [43], ranging from "*Definitely safe*" to "*I cannot decide*", to "*Definitely dangerous*". After assessing the risk for the link, the participants were also asked if, in the context of the presented scenario, they would click on the link and were prompted to briefly explain the rationale behind their decisions. The test suite included both malicious and benign URLs. In particular, we tested the following attacks:

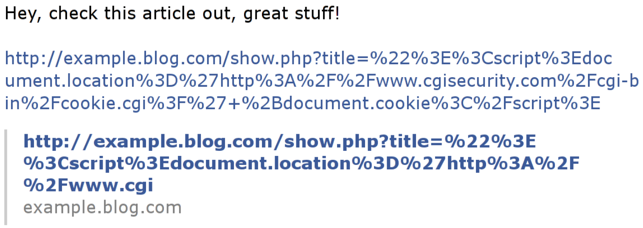

Hey, check this article out, great stuff!

http://example.blog.com/show.php?title=%22%3E%3Cscript%3Edoc
ument.location%3D%27http%3A%2F%2Fwww.cgisecurity.com%2Fcgi-b
in%2Fcookie.cgi%3F%27+%2Bdocument.cookie%3C%2Fscript%3E

**http://example.blog.com/show.php?title=%22%3E
%3Cscript%3Edocument.location%3D%27http%3A%2F
%2Fwww.cgi**
example.blog.com

Figure 1. A Facebook wall post containing a link with a reflected cross-site scripting attack. The parameter "title" contains a malicious script.

- *Cross-site Scripting:* This is a prevalent Internet threat in which an attacker injects client-side scripts into the browsers of users when they visit a vulnerable web page. The most common type of this vulnerability, called reflected cross-site scripting, typically occurs when a web application fails to properly sanitize the data it receives from a web client (e.g., HTTP parameters) and directly uses it in a generated web page. An attacker can then create a link to the vulnerable web page, and include a script as one of the parameters. When a user, unaware of the threat, clicks on this link, the malicious script is served to her browser along with the requested web page, and is executed. An attacker can distribute these malicious links via spam emails, or post them on the Internet. An example of this attack that we included in the web-based attacks suite is shown in Figure 1.

- *Session Fixation:* When a web application that authenticates its users using session identifiers fails to properly invalidate the previous sessions, an attacker can exploit this behavior to steal authenticated sessions. In a simple attack scenario, the attacker creates a new session on a web application, records the corresponding session identifier, and crafts a link to the web application using the recorded session identifier (e.g., *https://www.mybank.com/online/signup?sessionid=395hd74mcue7nb2h1j09*). After a user clicks on this link to access the web application and authenticate herself using the recorded session identifier, the attacker can use the same identifier to hijack the active session.

- *Link Manipulation Tricks:* Link manipulation deceives users by making them believe that they are following a legitimate and safe link, while in reality, the link typically leads to a malicious page that immediately infects the user's computer (e.g., by a drive-by download attack) or to a phishing page. Examples of link manipulation tricks in the web attacks suite included misspelled URLs, the use of subdomains (e.g., *http://www.paypal.hostding.com*, which appears to be PayPal, but in fact is a subdomain of *hostding*), including usernames in the URLs (e.g., *http://www.twitter.com@twiter.com*, which appears to be the official Twitter website, but actually links to the suspicious *twiter.com*), and using anchor texts (i.e., the text enclosed by the `<a>` HTML tag) that does not match the real link destinations.

The rest of the tests included perfectly benign, but less conventional links with, for instance, a very long parameter list, a non-HTTP protocol, and mixed-case characters.

We also showed the participants two additional URLs, the destinations of which could not be determined without further analysis: a raw IP address, and a shortened TinyURL [12] link. In these two tests, apart from following the link or ignoring it, the participants were given a third option: to verify the destination and then decide if they would follow the link. The participants who chose this option were asked to briefly explain how they would attempt to determine the destination.

Each URL and its related questions were presented on a separate page. We also provided a short scenario to describe to the participants the context in which they encountered the URL. The URLs were created using actual HTML anchor tags to allow the browsers to render them authentically, and to allow the participants to hover their mouse over the links to see the hyperlink destination in the browser's status bar.

Clicking on links was disabled for all tests by appropriate Javascript code. Note that there were two exceptions where we included a screenshot instead of embedding the link in the HTML page. These were attacks that required a visual context surrounding the link (e.g., the Facebook wall post shown in Figure 1).

3.3. Email-Based Attacks

In the email-based attacks suite, we showed the participants screenshots of emails together with full header information (see Figure 2 for an example). The suite included a PayPal phishing email, a spam email suggesting to click on a suspicious IP link, an ordinary E-Bay newsletter, a fake prize-giveaway notification, an email with a malicious attachment, an advance-fee fraud with the classic Nigerian connection text (e.g., [55]), an innocuous Amazon advertisement, and a phishing email crafted to look like it was sent from a bank. Similar to the web-based attacks test suite, we asked the participants to rate the risk they perceived on a 5-point Likert scale. We also inquired whether they would react to the email, for example, by downloading the attachment, and asked them to state the reasons behind their decisions.

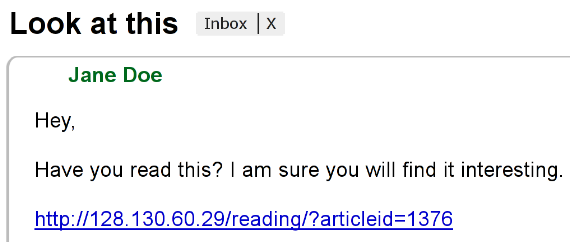

Look at this

Jane Doe

Hey,

Have you read this? I am sure you will find it interesting.

http://128.130.60.29/reading/?articleid=1376

Figure 2. A spam email with generic text. The message attempts to trick the user into following a suspicious link.

Again, each email and the related questions were shown on a separate page accompanied by a short text for providing context, explaining to the participants the scenario of how and when they received the shown email. In order to avoid overwhelming the non-technical participants with complicated email header information, prior to every question, the participants were informed that if they did not know how to interpret the header information, they should ignore them and only focus on the email content.

3.4. File Sharing-Related Attacks

In the file sharing test suite, our aim was to confront the participants with typical (but potentially risky) file sharing scenarios (e.g., the download of an executable file disguised as a movie). Obviously, making the participants go through a number of file sharing questions if they did not possess prior experience in the area would not have produced useful results. Therefore, prior to these tests, we asked the participants whether they know what BitTorrent or one-click-hosting are, and whether they have had experience with the websites we used in our tests. The participants who responded negatively were not asked to complete the tests in this suite.

We split the file sharing test suite into two parts: BitTorrent-related tests and one-click-hosting-related tests. In the BitTorrent tests, in order to provide the participants with a concrete context, we walked the participants through a scenario in which they were trying to download their favorite movie. We showed the participants a set of screenshots of torrent search results and torrent detail listings, as presented by three popular BitTorrent hosters/meta-search engines: The Pirate Bay [11], isoHunt [6] and Torrentz [13]. Each of these pages included cues to the legitimacy of the movie file (or the lack thereof) such as file extensions, file sizes, torrent contents, number of people sharing the file, reputation of the uploader, torrent descriptions, and warn-

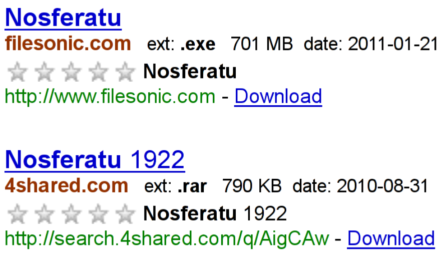

Nosferatu
filesonic.com ext: .exe 701 MB date: 2011-01-21
☆☆☆☆☆ Nosferatu
http://www.filesonic.com - Download

Nosferatu 1922
4shared.com ext: .rar 790 KB date: 2010-08-31
☆☆☆☆☆ Nosferatu 1922
http://search.4shared.com/q/AigCAw - Download

Figure 3. Search results for a movie in Filestube. The first hit has a bad file extension and the second has a suspiciously small file size.

ings in user submitted comments.

We then asked the participants to rate the risk they perceived for each search result we presented to them on a 5-point Likert scale, and decide whether they would proceed to download the file. Furthermore, we asked the participants which cues they used when they made the decision. While the screenshots showing the details of a single torrent were presented on a separate page, the search results were all given in the same page in order for the participants to be able to compare the search hits to each other.

Once they answered these questions, we took the participants to fully interactive torrent download pages carefully reproduced from the ones of The Pirate Bay and isoHunt. We then asked them to click on the correct download button among the various advertisement banners disguising themselves as legitimate download links. Note that we did not introduce any artificial web banners for the purpose of our study. Instead, we faithfully copied the original content from the corresponding websites for realistic observations.

We designed the one-click-hosting tests in a similar fashion. However, this time, we showed screenshots and reproduced pages from the popular websites Filestube [4], iFolder [5], Megaupload, [9] and Megavideo [10] (see Figure 3 for an example).

4. Analysis of the Test Results

In this section, we explain our strategy to evaluate the collected data, and we present the results obtained through the experiments. In Section 5, we interpret these results, and summarize the insights we distilled.

4.1. Demographics and Diversity

The test participants were 31.1% female and 68.9% male, their ages ranged from 19 to 69 (mean=26.52, s.d.=8.76, variance=76.79), and their nationalities spanned 17 different countries. 11.6% of the participants held a doctoral degree, 10.4% a master's degree, 45.7% a bachelor's degree, and 1.2% an associate's degree. 72.6% of our participants were continuing students of which 5.5% were pursuing a doctoral degree, 37.2% a master's degree, and 29.9% a bachelor's degree. 33.5% of them were employed. The majority of the participants reported being comfortable with common computer tasks and concepts such as surfing the web, using email services and social network applications, doing entertainment activities such as watching movies and playing games, performing e-commerce and e-banking operations, installing and using applications, and doing basic system maintenance by configuring their operating system and recovering from simple errors (see Figure 4 for a summary).

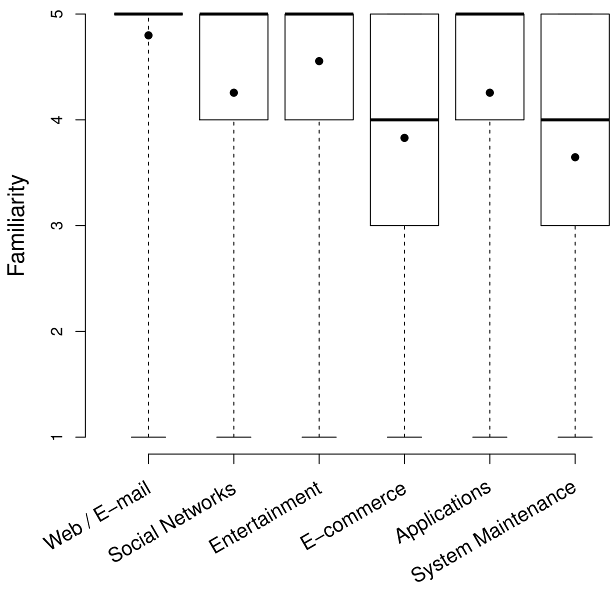

Figure 4. Five-number summaries for participant familiarity with typical computer tasks and concepts (5: Familiar, 1: Not familiar). Mean values are displayed by the black dots.

65.8% of the test participants used Windows as their primary operating system, 17.7% used Linux, 15.8% used Mac OS X, and one participant used a BSD-variant. 41.5% of them preferred Firefox, 34.8% preferred Chrome, 14.6% preferred Internet Explorer, 6.1% preferred Safari, and the remaining 3.0% used various other web browsers.

Based on the participants' responses to the questions about their security background, we divided them into three expertise groups:

- *Non-techies* use computers at home or work on a regular basis. They are comfortable using basic applications to perform everyday tasks. Their professions are in non-technical fields, and they have little, or no programming experience. 42.7% of the test participants fall into this group.

- *Techies* are either computer scientists, or otherwise involved in a closely-related field of study or profession (such as engineering disciplines dealing with technology). These participants are knowledgeable on the intricacies of computer systems. However, their technical training does not focus on computer security. Techies constitute 19.5% of the participants.

- *Experts* are computer security professionals. In their studies or professions, they specialize in securing computer systems. They claim to have a deep understanding of security fundamentals, and have some practical experience in the field. 37.8% of the participants are experts.

4.2. Security and Risk Perception

In order to quantify the performance of the participants in the security tests, we computed two global scores: a *security score* and a *risk perception score*.

The *security score* is a measure of how good a participant is at averting attacks, while also refraining from erroneously discrediting non-threats as being dangerous. We compute the security score as the total number of questions answered correctly in this manner. We then normalize it to account for the participants who skipped any of the file sharing tests, and scale it to a value between 0 and 100.

The *risk perception score* is a measure of a participant's ability to recognize the severity and consequences of each situation. We compute it based on the 5-point Likert [43] scale, with questions that ask the participants how dangerous they think each scenario is. For an obvious threat, the participants who respond with 'definitely dangerous' receive 5 points, while those who answer 'definitely safe' only receive 1 point. For benign items, we reverse the scores accordingly. We then normalize and scale the score in the same way we calculate the security score.

Considering all the participants in all tests, the security scores ranged from 46.43 to 96.97 (mean=70.21, s.d.=10.84, variance=117.53, 1st quartile=63.64, median=69.70, 3rd quartile=78.57) and risk perception scores ranged from 48.18 to 90.59 (mean=70.48, s.d.=7.35, variance=54.09, 1st quartile=66.29, median=70.45, 3rd quartile=75.19). The distributions of scores for the three groups are summarized in the plots in Figure 5. Note that, since scores show what percentage of the questions were answered correctly, values close to the middle of the scale (i.e., 50) could possibly indicate no security awareness but merely random guesses.

Out of all the questions the participants answered incorrectly, 56.1% were benign samples misjudged as being malicious. 43.9% were attacks confused as being benign. This slight imbalance could be explained by the observation that participants were over-careful, expecting to encounter attacks in the tests.

The Kruskal-Wallis one-way analysis of variance tests [40] showed that both the security and risk perception scores differed significantly among the three participant groups (i.e., $H = 26.89$, $df = 2$, $p = 1.45 \times 10^{-6}$ for security scores and $H = 37.36$, $df = 2$, $p = 7.71 \times 10^{-9}$ for risk perception scores). Following these with multiple comparison post-hoc tests revealed that non-techies and experts differed significantly in both scores, while non-techies and techies, or techies and experts did not. This means that both security and risk perception considerably increased with security expertise.

We also analyzed the scores for each test suite separately (see Table 1), and checked the scores once again for potential differences among participant groups. Similarly, the risk perception scores differed significantly between non-techies and experts in every test suite. However, we observed that the security scores only differed significantly for the web-based attacks suite between non-techies and experts (i.e., $H = 25.42$, $df = 2$, $p = 3.01 \times 10^{-6}$). We did not observe a statistically significant difference in security scores for the email and the two file sharing test suites (i.e., $p > 0.54$, $p > 0.23$ and $p > 0.67$, respectively).

When we investigated the relationship between risk perception and security scores in each participant group (see Figure 6), an analysis with Spearman's rank correlation [51] revealed that the two types of scores are positively correlated for each group. In other words, for higher risk perception scores, the security scores show an increasing trend as well. However, this correlation is considerably weaker for non-techies (i.e., $\rho = 0.50$, $p = 9.99 \times 10^{-6}$) compared to techies and experts combined (i.e., $\rho = 0.70$, $p = 6.51 \times 10^{-15}$).

In our study group, we did not see a significant correlation between age and education level, and either of the scores. The score medians differed significantly with sex, where the females scored significantly lower compared to the males. However, this was largely due to the fact that the females were concentrated in the non-techies. When testing for each group separately there was no significant difference in the score medians within the groups.

4.3. Test-Specific Results

Some of the test suites contained questions that cannot be quantified with the previous scoring approach. Hence, we used case-specific evaluation strategies for these questions, as described in the following subsections.

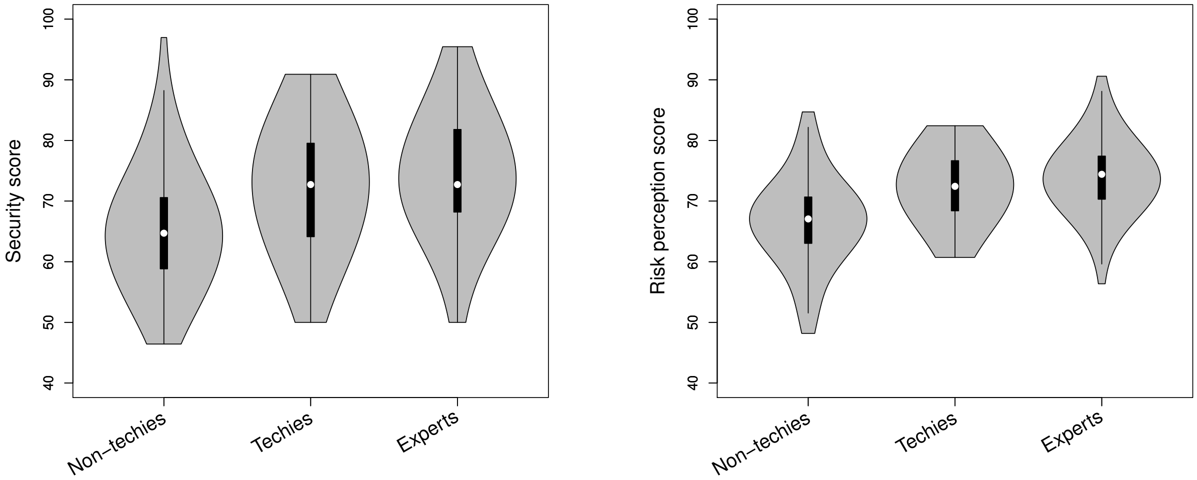

Figure 5. Five-number summaries and probability densities for total security and perception scores.

		Security Scores					
		Min	1^{st} Quartile	Median	Mean	3^{rd} Quartile	Max
Web Tests	Non-techies	22.2	44.4	55.6	56.7	66.7	100.0
	Techies	44.4	55.6	66.7	68.0	77.8	100.0
	Experts	33.3	66.7	77.8	74.2	88.9	100.0
Email Tests	Non-techies	50.0	75.0	75.0	77.5	87.5	100.0
	Techies	62.5	75.0	87.5	85.2	100.0	100.0
	Experts	50.0	75.0	87.5	84.7	100.0	100.0
BitTorrent Tests	Non-techies	36.4	54.5	54.5	58.7	63.6	90.9
	Techies	36.4	45.4	63.6	63.6	75.0	90.9
	Experts	36.4	54.5	63.6	63.6	72.7	81.8
One-click Hosting Tests	Non-techies	20.0	60.0	60.0	64.8	80.0	100.0
	Techies	40.0	65.0	80.0	79.1	100.0	100.0
	Experts	40.0	60.0	80.0	76.1	100.0	100.0

		Risk Perception Scores					
		Min.	1^{st} Quartile	Median	Mean	3^{rd} Quartile	Max
Web Tests	Non-techies	46.7	60.0	63.3	64.1	68.9	82.2
	Techies	53.3	64.4	71.1	71.1	75.6	91.1
	Experts	53.3	68.9	75.6	74.3	79.4	91.1
Email Tests	Non-techies	40.0	65.0	72.5	71.7	80.0	97.5
	Techies	55.0	75.0	80.0	80.8	87.5	97.5
	Experts	60.0	75.6	81.2	81.4	90.0	92.5
BitTorrent Tests	Non-techies	49.1	56.4	60.9	60.9	63.6	76.4
	Techies	56.4	61.8	66.4	65.8	67.7	78.2
	Experts	54.5	61.8	67.3	66.4	70.9	80.0
One-click Hosting Tests	Non-techies	48.0	60.0	68.0	67.6	72.0	100.0
	Techies	48.0	69.0	76.0	74.5	84.0	92.0
	Experts	56.0	68.0	72.0	73.1	80.0	92.0

Table 1. Five-number summaries and mean values for each test suite and participant group.

4.3.1 IP addresses and shortened URLs

In the web-based attacks suite, the legitimacy of the IP and shortened URL links (e.g., TinyURL) cannot be determined just by looking at the URL. Hence, we did not compute scores for them. Instead, we investigated how many participants were able to successfully verify the link destinations. For example, the participants could have fetched HTML headers, performed WHOIS and reverse DNS look-ups, or utilized URL expansion tools. A summary of these results is given in Table 2.

		Blindly Follow	Ignore	Technically Verify	Depends on Source	Not Familiar
TinyURL	Non-tech	17.1 %	74.3 %	0.0 %	8.6 %	35.7 %
	Tech	21.9 %	31.3 %	18.7 %	28.1 %	15.6 %
	Expert	16.1 %	32.3 %	32.3 %	19.3 %	4.8 %
Raw IP	Non-tech	15.7 %	80.0 %	0.0 %	4.3 %	28.6 %
	Tech	18.8 %	43.7 %	25.0 %	12.5 %	6.3 %
	Expert	17.7 %	38.7 %	33.9 %	9.7 %	3.2 %

Table 2. Summary of participants' decisions whether to follow a link or not, when shown a TinyURL and an IP link. The “Not Familiar” column is not mutually exclusive with the others; some participants chose to follow the links without knowing what shortened URLs or IP addresses are.

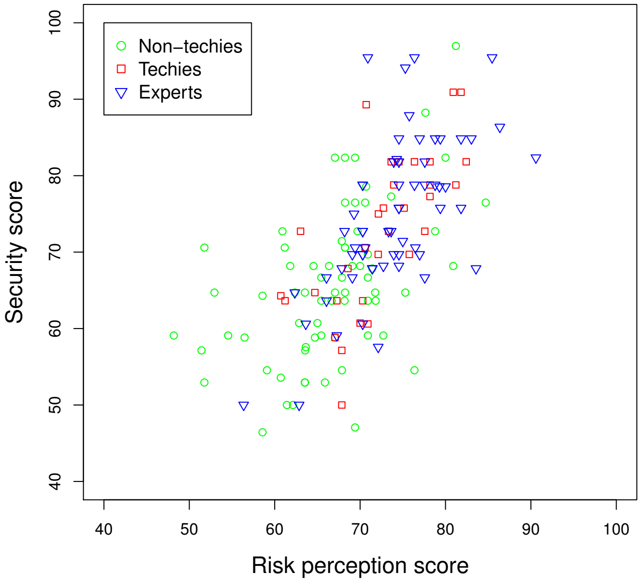

Figure 6. The relationship between risk perception and security scores. The correlation between the two scores is considerably weaker for non-techies.

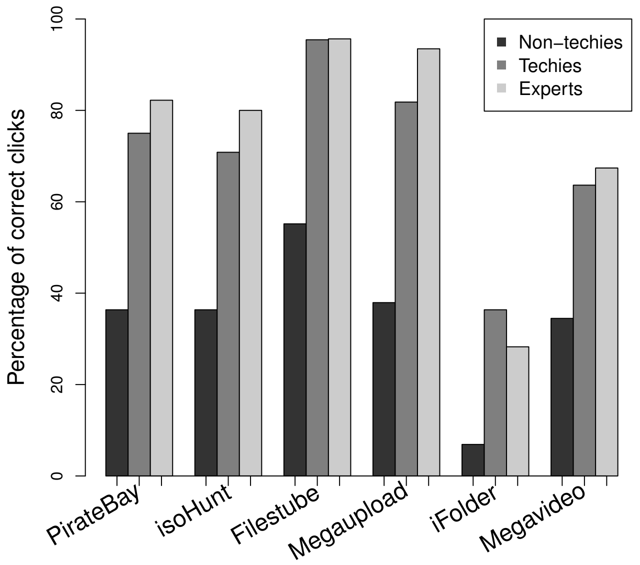

Figure 7. Number of clicks on correct download buttons, as opposed to banners, for the trick banner tests on pages reproduced from popular file sharing websites.

An interesting observation in the experiments was that none of the participants in the non-techie group said that they would attempt to verify the destination of either the IP, or the TinyURL link. Moreover, compared to the other groups, a considerably higher number of non-techies did not know what an IP address, or shortened URL was.

Note that many participants directly related to their previous surfing experience and were confused by similar links they had used in the past. In their answers, these participants demonstrated a completely wrong technical understanding of the use of IP addresses, or URL shortening services. For example, some answers we received indicated an IP to be an interface for a printer/router configuration screen, an index of photographs, or a “proxy code”. The TinyURL link was instead thought to be a YouTube video, a photo sharing service, a blog, or a music download website.

4.3.2 Trick banner tests

For the simulated websites in the file sharing test suite, we counted the number of clicks on the correct download links, as well as on deceptive banners crafted with the aim of luring Internet users to click on them. The percentage of clicks on the correct links for all participant groups are shown in Figure 7. The high error rate in the iFolder test was in part due to the fact that we reproduced the page in its original language (i.e., Russian). The aim was to observe the participants' navigation behavior when the website is in an unfamiliar language.

These results showed that while most experts and techies were able to recognize and avoid false banners, the majority of non-techies were deceived. That is, such participants did not realize that they were not clicking on the actual link (and were being tricked into clicking on potentially dangerous banners) even though they reported being familiar with the test website.

5. Discussion and Insights Gained

In this section, we provide detailed interpretations of the test results presented in Section 4, and list the insights we distilled from them.

5.1. Exposure to Threats and Risk Perception

As shown in Section 4.2, among the individual test suites, only in the web-based attacks did the security experts and techies get significantly higher security scores than non-techie participants. In other words, we did not see statistical proof that, when generalized to the whole population, security experts would perform better in email and file sharing security scenarios compared to technically unsophisticated users.

While our experimental setup is not designed to directly identify the causes of this effect, the data we have collected from the participants in order to estimate their security expertise prior to the tests suggest that increased exposure to security threats help non-technical participants avert common and less intricate attacks, such as email scams. Specifically, when asked for their familiarity with our attack scenarios, 95.7% of non-techies reported being exposed to spam and suspicious emails regularly, while only 48.6% said they recall encountering at least one malicious URL on the Internet.

Note that the fact that users are able to detect an attack does not necessarily mean that they also understand the way in which the attack works, or that they correctly *perceive* the risk involved. For example, in the email test suite, when we categorized the answers of non-techies by looking at their explanations, only 2.9% of the responses provided meaningful technical insights, while the remaining 97.1% were based purely on intuition and past experience. In contrast, 23.4% of the techies and 30.6% of the experts directly looked for technical cues of an attack (e.g., by investigating the email headers). Indeed, the results we presented in Section 4.2 indicate that the difference in risk perception scores between non-techies and experts is statistically significant even when the security scores do not significantly differ among these two groups. That is, non-techies and experts have different perceptions of the risk in a given situation (i.e., the risk perception scores differ significantly). Nevertheless, these groups reach similar conclusions, and act in a similar manner (i.e., the security scores do not differ significantly). For instance, in many cases, non-techies judged a malicious email (e.g., a PayPal phishing scam) as being "Definitely safe", or being "Most Probably Safe", stating that they cannot read email headers, and that they do not see anything wrong with the content. However, they chose to ignore the mail instead of clicking on the given link. One participant explained: "Looks good. But I don't trust it, I don't know why". Not being able to articulate the reasons behind a correct decision is a known indication of guesses based on intuition [25, 35].

This observation is also supported by the relatively weaker correlation between the risk perception scores and the security scores of non-techies, compared to techies and experts. Although such a correlation in no way implies a causality relationship between the two scores, it shows that for experts, a higher risk perception is associated with higher security, but much less so for non-technical participants.

All of these observations suggest that most of non-techie users can *instinctively* avoid common scams, even without having technical knowledge or perceiving the severity of a threat, possibly because of their high exposure to such attacks in everyday life. This observation is also in line with psychology literature which shows that individuals fall back on their intuition when faced with complex information that they cannot process, and the guesses based on intuition could be correct since they draw from vast previous experience [25, 35].

A notable implication of this observation is that security games and online test platforms that are tailored towards non-technical people in order to familiarize them with attack patterns (e.g., PhishGuru and Anti-Phishing Phil [41, 50]) could effectively be used to help achieve a similar effect to the one we observed in our tests. General psychology literature also supports the idea that intuition could be "taught" by repeated experience, and also by virtual simulations [35, 49].

5.2. Size Matters

When the participants did not have the technical knowledge to make an informed decision for a test and had to rely on their intuition, a very common trend was to make a guess based on the "size", the "length", or the "complexity" of the artifacts involved. For example, a benign Amazon link was labeled as malicious by non-technical participants on the basis that the URL contained a crowded parameter string. Some of the comments included: "*Too long and complicated.*", "*It consists of many numbers.*", "*It has lots of funny letters.*" and "*It has a very long name and also has* *some unknown code in it.*". Many of these participants later said they would instead follow a malicious PayPal phishing URL because "*It is simple.*", "*Easy to read.*", "*Clear obvious link.*" and it has a "*Short address*". One participant made a direct comparison between the two links: "*This is not dangerous, address is clear. [Amazon link] was dangerous because it was not like this.*". Interestingly, in some cases, the non-technical participants managed to avert attacks thanks to this strategy. For example, a number of participants concluded that a Facebook post containing a code injection attack was dangerous solely on the grounds that the link was "long" and "confusing".

Analogously, in the file sharing tests, the responses based on intuition mainly relied on arguments about the file size. For example, the participants who did not understand how BitTorrent works judged torrents merely on their expectations of a full-length movie's size. These participants often made misinformed decisions such as discrediting a 700MB RAR archive as being malicious as the size of the movie had not decreased after the compression (note that movie files are already heavily compressed), or a 790K file as being correct since it referred to a very old movie from 1922.

Again, these results underline the importance of familiarizing users with common security-related scenarios to increase their security awareness. When users are not able to make an informed decision about a possible threat, they fall back on judging the situation based on often misleading characteristics, such as an item's size and complexity.

5.3. URL Shortening Services and Tools

Our tests indicate that none of the non-technical participants attempted to verify the destination of a shortened URL (in our case, a TinyURL). As explained in Section 4.3.1, the majority of the non-techie group was not aware of the fact that a shortened URL could link to any destination on the web. Rather, they thought that TinyURL was the website that actually hosted the content. Even those participants who were aware of the risks stated that they did not know how to verify the destinations of these links.

A wide variety and number of URL shortening services are available on the Internet today. Their frequent use in social networks such as Twitter make them ubiquitous. Unfortunately, the prevalence of shortened URLs also make them an effective way to distribute malware and lure users into scams. A recent study by Grier et al. [30] states that over 2 million links posted on Twitter point to attack pages and that through the use of nested URL shortening, blacklisting solutions can be circumvented. Our results demonstrate that non-technical users are easily tricked by shortened URLs in practice.

There exists several online services (e.g., [3, 8]), and extensions for popular browsers (e.g., [2, 14]) that offer shortened URL expansion capabilities. While these tools would definitely help technically inclined people assess the risk before following a shortened URL, our experiments show that they are ineffective for non-technical users who do not have a firm grasp of the technology behind URL shortening.

Analogous to the recent integration of website blacklists and phishing detection heuristics into popular browsers (e.g., the anti-phishing features in IE as of Version 7), we believe that URL expansion and threat detection capabilities (e.g., [7]) need to be integrated into browsers as soon as possible.

5.4. Trick Banners

In the interactive tests featuring reproduced download websites, the false click rates for non-techies were considerably higher compared to experts. In 5 of the 6 tests, more than half of the participants clicked on a banner instead of the real link. That is, even if these participants were able to differentiate between a legitimate and a malicious search result displayed by the file sharing website, they still would not have managed to complete the download successfully. On another note, the correct click rates for security experts could get as low as 70% (in the Megavideo test), which indicates that trick banners are also effective to some extent at deceiving more knowledgeable users.

Using deceptive banners to trick Internet users into visiting a website is a well-known advertisement strategy [53]. However, there have also been recent attacks on the advertisement networks of popular websites where attackers have legitimately bought banner space [54, 56, 59] or exploited bugs in ad servers (e.g., in a recent attack against The Pirate Bay [57]). In such attacks, the attackers typically use banners to serve malware. Additionally, some malware have utilized trick banners for committing fraud [58]. Our study empirically confirms that trick banners are very effective (attack) techniques in influencing the click behavior of users. From a user's point of view, a possible defense technique in dealing with such tricks would be utilizing ad-blockers (e.g., [1]). Hence, it is important to inform and train users about the use of such tools, especially when visiting certain classes of websites.

6. Conclusions

In this paper, we described an experiment platform for observing the behavior of users when they are confronted with typical benign and malicious interaction scenarios on the Internet. We presented the results of a study we had conducted on 164 Internet users who possess diverse backgrounds and varying degrees of computer security knowledge. Our results confirm the general intuition that technical security knowledge has a considerable positive impact on a user's ability to assess risk and make correct security decisions, especially when the threats involve technically complex attacks. However, for relatively simple and common threats that users are frequently exposed to (e.g., well-known email scams), non-technical users can exhibit performance comparable to knowledgeable users by solely depending on their intuition and past experience.

We observed that many users consider unusual "size" and "length" characteristics of URLs and downloaded files as an indicator of risk. Moreover, we have also seen that users are often highly susceptible to attack strategies that exploit shortened URLs, raw IP addresses, and trick banners. Recently, URL-expansion tools such as Longshore [7] have been introduced that aim to assist users in revealing the real destinations of shortened URLs. Our findings suggest that such security services are largely ineffective for non-technical users since they are not able to use them, or they do not understand the concepts behind URL shortening services.

7. Acknowledgments

The research leading to these results was partially funded by the EU Seventh Framework Programme (FP7/2007-2013) under grant agreement nº 257007, the National Science Foundation (NSF) under grant CNS-1116777, COMET K1, the Austrian Research Promotion Agency (FFG) and Secure Business Austria.

References

[1] Adblock Plus. `http://adblockplus.org/en/`, 2011.

[2] ChromeMUSE - Multi-URL Shortener/Expander. `https://chrome.google.com/extensions/`, 2011.

[3] Clybs - Url expander. `http://www.clybs.com/urlexpander`, 2011.

[4] Filestube. `http://www.filestube.com/`, 2011.

[5] iFolder. `http://www.ifolder.com/`, 2011.

[6] isoHunt. `http://isohunt.com/`, 2011.

[7] Long-Shore. `http://long-shore.com/`, 2011.

[8] LongURL. `http://longurl.org/`, 2011.

[9] Megaupload. `http://www.megaupload.com/`, 2011.

[10] Megavideo. `http://www.megavideo.com/`, 2011.

[11] The Pirate Bay. `http://thepiratebay.org/`, 2011.

[12] TinyURL. `http://www.tinyurl.com/`, 2011.

[13] Torrentz. `http://torrentz.eu/`, 2011.

[14] Xpnd.it! short URL expander. `http://addons.mozilla.org/en-us/firefox/addon/xpndit-short-url-expander/`, 2011.

[15] M. Antonakakis, R. Perdisci, D. Dagon, W. Lee, and N. Feamster. Building a dynamic reputation system for dns. In *Proceedings of the 19th USENIX conference on Security*, USENIX Security'10, pages 18–18, Berkeley, CA, USA, 2010. USENIX Association.

[16] M. Bailey, J. Oberheide, J. Andersen, Z. M. Mao, F. Jahanian, and J. Nazario. Automated classification and analysis of internet malware. In *Proceedings of the 10th international conference on Recent advances in intrusion detection*, RAID'07, pages 178–197, Berlin, Heidelberg, 2007. Springer-Verlag.

[17] A. Barth, C. Jackson, and J. C. Mitchell. Robust defenses for cross-site request forgery. In *Proceedings of the 15th ACM conference on Computer and communications security*, CCS '08, pages 75–88, New York, NY, USA, 2008. ACM.

[18] L. Bilge, T. Strufe, D. Balzarotti, and E. Kirda. All your contacts are belong to us: automated identity theft attacks on social networks. In *Proceedings of the 18th international conference on World wide web*, WWW '09, pages 551–560, New York, NY, USA, 2009. ACM.

[19] S. Chiasson, P. C. van Oorschot, and R. Biddle. A usability study and critique of two password managers. In *Proceedings of the 15th conference on USENIX Security Symposium - Volume 15*, Berkeley, CA, USA, 2006. USENIX Association.

[20] J. Clark, P. C. van Oorschot, and C. Adams. Usability of anonymous web browsing: an examination of tor interfaces and deployability. In *Proceedings of the 3rd symposium on Usable privacy and security*, SOUPS '07, pages 41–51, New York, NY, USA, 2007. ACM.

[21] CNN. Amazon EC2 outage downs Reddit, Quora. `http://money.cnn.com/2011/04/21/technology/amazon_server_outage/index.htm`, 2011.

[22] G. Conti and E. Sobiesk. Malicious interface design: exploiting the user. In *Proceedings of the 19th international conference on World wide web*, WWW '10, pages 271–280, New York, NY, USA, 2010. ACM.

[23] L. F. Cranor. A framework for reasoning about the human in the loop. In *Proceedings of the 1st Conference on Usability, Psychology, and Security*, pages 1:1–1:15, Berkeley, CA, USA, 2008. USENIX Association.

[24] R. Dhamija, J. D. Tygar, and M. Hearst. Why phishing works. In *Proceedings of the SIGCHI conference on Human Factors in computing systems*, CHI '06, pages 581–590, New York, NY, USA, 2006. ACM.

[25] A. Dijksterhuis, M. W. Bos, L. F. Nordgren, and R. B. van Baaren. On making the right choice: The deliberation-without-attention effect. *Science*, 311:1005–1007, February 2006.

[26] S. Egelman, L. F. Cranor, and J. Hong. You've been warned: an empirical study of the effectiveness of web browser phishing warnings. In *Proceeding of the twenty-sixth annual SIGCHI conference on Human factors in computing systems*, CHI '08, pages 1065–1074, New York, NY, USA, 2008. ACM.

[27] D. Endler. The Evolution of Cross Site Scripting Attacks. Technical report, iDEFENSE Labs, 2002.

[28] B. Friedman, D. Hurley, D. C. Howe, E. Felten, and H. Nissenbaum. Users' conceptions of web security: a comparative study. In *CHI '02 extended abstracts on Human factors in computing systems*, CHI EA '02, pages 746–747, New York, NY, USA, 2002. ACM.

[29] B. Friedman, D. Hurley, D. C. Howe, H. Nissenbaum, and E. Felten. Users' conceptions of risks and harms on the web: a comparative study. In *CHI '02 extended abstracts on Human factors in computing systems*, CHI EA '02, pages 614–615, New York, NY, USA, 2002. ACM.

[30] C. Grier, K. Thomas, V. Paxson, and M. Zhang. @spam: the underground on 140 characters or less. In *Proceedings of the 17th ACM conference on Computer and communications security*, CCS '10, pages 27–37, New York, NY, USA, 2010. ACM.

[31] G. Gu, R. Perdisci, J. Zhang, and W. Lee. Botminer: clustering analysis of network traffic for protocol- and structure-independent botnet detection. In *Proceedings of the 17th conference on Security symposium*, pages 139–154, Berkeley, CA, USA, 2008. USENIX Association.

[32] G. Gu, P. Porras, V. Yegneswaran, M. Fong, and W. Lee. Bothunter: detecting malware infection through ids-driven dialog correlation. In *Proceedings of 16th USENIX Security Symposium on USENIX Security Symposium*, pages 12:1–12:16, Berkeley, CA, USA, 2007. USENIX Association.

[33] C. Herley. So long, and no thanks for the externalities: the rational rejection of security advice by users. In *Proceedings of the 2009 workshop on New security paradigms workshop*, NSPW '09, pages 133–144, New York, NY, USA, 2009. ACM.

[34] J. T. Ho, D. Dearman, and K. N. Truong. Improving users' security choices on home wireless networks. In *Proceedings of the Sixth Symposium on Usable Privacy and Security*, SOUPS '10, pages 12:1–12:12, New York, NY, USA, 2010. ACM.

[35] R. M. Hogarth. *Educating Intuition*. Univ. of Chicago Press, 2001.

[36] S. Institute. Top Cyber Security Risks, September 2009. `http://www.sans.org/top-cyber-security-risks/summary.php`.

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

[38] T. N. Jagatic, N. A. Johnson, M. Jakobsson, and F. Menczer. Social phishing. *Commun. ACM*, 50:94–100, October 2007.

[39] E. Kirda, C. Kruegel, G. Banks, G. Vigna, and R. A. Kemmerer. Behavior-based spyware detection. In *Proceedings of the 15th conference on USENIX Security Symposium - Volume 15*, Berkeley, CA, USA, 2006. USENIX Association.

[40] W. Kruskal and W. A. Wallis. Use of ranks in one-criterion variance analysis. *Journal of the American Statistical Association*, pages 583–621, 1952.

[41] P. Kumaraguru, S. Sheng, A. Acquisti, L. F. Cranor, and J. Hong. Teaching johnny not to fall for phish. *ACM Trans. Internet Technol.*, 10:7:1–7:31, June 2010.

[42] E. Lin, S. Greenberg, E. Trotter, D. Ma, and J. Aycock. Does domain highlighting help people identify phishing sites? In *Proceedings of the 2011 annual conference on Human factors in computing systems*, CHI '11, pages 2075–2084, New York, NY, USA, 2011. ACM.

[43] R. Linkert. A Technique for the Measurement of Attitudes. In *Archives of Psychology*, volume 140, 1932.

[44] M.-E. Maurer, A. De Luca, and H. Hussmann. Data type based security alert dialogs. In *Proceedings of the 2011 annual conference extended abstracts on Human factors in computing systems*, CHI EA '11, pages 2359–2364, New York, NY, USA, 2011. ACM.

[45] S. Motiee, K. Hawkey, and K. Beznosov. Do windows users follow the principle of least privilege?: investigating user account control practices. In *Proceedings of the Sixth Symposium on Usable Privacy and Security*, SOUPS '10, pages 1:1–1:13, New York, NY, USA, 2010. ACM.

[46] F. Raja, K. Hawkey, S. Hsu, K.-L. Wang, and K. Beznosov. Promoting a physical security mental model for personal firewall warnings. In *Proceedings of the 2011 annual conference extended abstracts on Human factors in computing systems*, CHI EA '11, pages 1585–1590, New York, NY, USA, 2011. ACM.

[47] P. Ratanaworabhan, B. Livshits, and B. Zorn. Nozzle: a defense against heap-spraying code injection attacks. In *Proceedings of the 18th conference on USENIX security symposium*, SSYM'09, pages 169–186, Berkeley, CA, USA, 2009. USENIX Association.

[48] S. E. Schechter, R. Dhamija, A. Ozment, and I. Fischer. Emperor's new security indicators: An evaluation of website authentication and the effect of role playing on usability studies, 2007.

[49] M. E. P. Seligman and M. Kahana. Unpacking intuition: a conjecture. *Perspectives on Psychological Science*, 4:399–402, July 2009.

[50] S. Sheng, B. Magnien, P. Kumaraguru, A. Acquisti, L. F. Cranor, J. Hong, and E. Nunge. Anti-phishing phil: the design and evaluation of a game that teaches people not to fall for phish. In *Proceedings of the 3rd symposium on Usable privacy and security*, SOUPS '07, pages 88–99, New York, NY, USA, 2007. ACM.

[51] C. Spearman. The proof and measurement of association between two things. *American Journal of Psychology*, 15:88–103, 1904.

[52] J. Sunshine, S. Egelman, H. Almuhimedi, N. Atri, and L. F. Cranor. Crying wolf: an empirical study of ssl warning effectiveness. In *Proceedings of the 18th conference on USENIX security symposium*, SSYM'09, pages 399–416, Berkeley, CA, USA, 2009. USENIX Association.

[53] M. Terms. Trick Banner. `http://www.marketingterms.com/dictionary/trick_banner/`.

[54] ThreatPost. Major Ad Networks Found Serving Malicious Ads. `https://threatpost.com/en_us/blogs/major-ad-networks-found-serving-malicious-ads-121210`, December 2010.

[55] C. Tive. *419 Scam: Exploits of the Nigerian Con Man*. iUniverse.com, 2006.

[56] TorrentFreak. Yahoo! pimping malware from banner ads. `http://www.theregister.co.uk/2008/04/28/yahoo_serves_rogue_ads/`, April 2008.

[57] TorrentFreak. Hackers Target and Exploit Pirate Bay Ad Server. `http://torrentfreak.com/hackers-target-and-exploit-pirate-bay-ad-server-100913/`, September 2010.

[58] Trusteer. Zeus Adds Investment Fraud to its Bag of Tricks. `http://www.trusteer.com/blog/zeus-adds-investment-fraud-its-bag-tricks/`, April 2011.

[59] Wired. Rogue Anti-Virus Slimeballs Hide Malware in Ads. http://www.wired.com/epicenter/2007/11/doubleclick-red/, November 2007.

[60] M. Wu, R. C. Miller, and S. L. Garfinkel. Do security toolbars actually prevent phishing attacks? In *Proceedings of the SIGCHI conference on Human Factors in computing systems*, CHI '06, pages 601–610, New York, NY, USA, 2006. ACM.

[61] Y. Xie and A. Aiken. Static detection of security vulnerabilities in scripting languages. In *Proceedings of the 15th conference on USENIX Security Symposium - Volume 15*, Berkeley, CA, USA, 2006. USENIX Association.

[62] H. Yin, D. Song, M. Egele, C. Kruegel, and E. Kirda. Panorama: capturing system-wide information flow for malware detection and analysis. In *Proceedings of the 14th ACM conference on Computer and communications security*, CCS '07, pages 116–127, New York, NY, USA, 2007. ACM.