

## STATEMENT OF RESEARCH INTERESTS

SERGE EGELMAN

Human error is responsible for most software security problems, and ad-hoc solutions that do not account for human factors will be ineffective over time. Thus, my research goal is to solve software security problems by addressing human factors. I do this by designing, conducting, and analyzing data from human subjects experiments in order to construct mental models, and then I use these mental models to develop user-centered security solutions. In earlier work I performed studies on web browser security warnings, website privacy indicators, authentication systems, and account models to understand how average computer users' perceptions of risk influence their behaviors when they make trust decisions. In current and future work, I am combining human-computer interaction and behavioral economics techniques to determine how risk perceptions can be altered to better align them with reality, in order to aid users in making rational trust decisions.

### THESIS WORK

In my thesis work, I created and tested several design patterns for online privacy and security indicators by conducting several laboratory and field studies. I examined two areas: critical security warnings and passive indicators that convey contextual information. My work on critical warnings focused on phishing and SSL warnings [3, 8]. Figure 1 depicts an example design pattern. These contributions have been incorporated into Internet Explorer, which is used by millions of people. Ultimately this work will help web browser developers convey security information more effectively so that users will make more informed trust decisions.

<p><b>Design Pattern:</b> Warnings for a high level of severity need to be easily distinguishable from warnings for a lower level of severity.</p> <p><b>Implications:</b> If all warnings are designed similarly, user may become habituated to seeing a warning in a low-risk situation and then become habituated to seeing a similar-looking warning in a high-risk situation.</p> <p><b>Experiment:</b> I conducted a laboratory study to test the Internet Explorer and Firefox phishing warnings by simulating a phishing attack. I analyzed the results using a model from the warning sciences—a branch of ergonomics.</p> <p><b>Results:</b> Users of Internet Explorer confused the phishing warnings with error messages that they frequently saw on legitimate websites, therefore ignoring the phishing warnings. This was not a problem for the unique Firefox phishing warnings.</p> <p><b>Impact:</b> I received a CHI 2008 best paper nomination [3]. I also helped redesign the security warnings in Internet Explorer 8.</p>
--

FIGURE 1. An example design pattern for critical online security warnings.

My work on passive indicators focused on icons representing website privacy policies. Throughout graduate school I worked on the W3C's Platform for Privacy Preferences (P3P). I created a search engine, Privacy Finder,<sup>1</sup> which annotates search results with privacy information [1]. I conducted several user studies to examine how users take privacy into account when making online purchases [4, 9]. After controlling for several factors, I found that users made better privacy decisions when search results were annotated (Figure 2). This research aids developers in providing users with more effective contextual information.

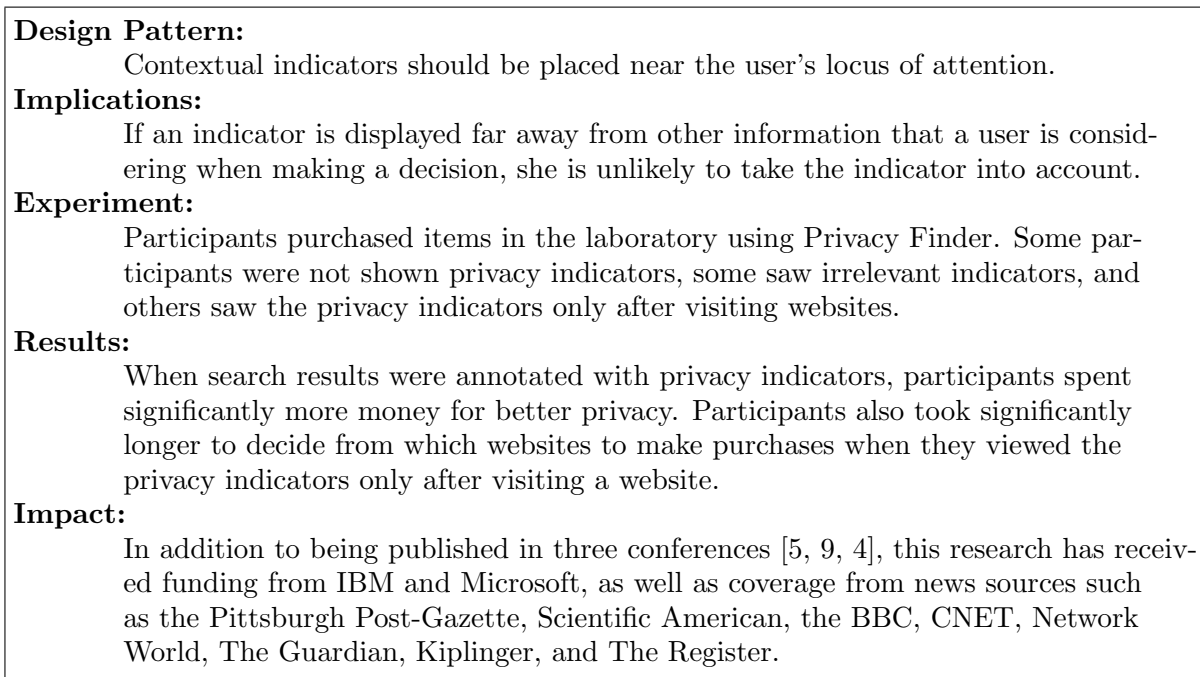


FIGURE 2. An example design pattern for passive contextual indicators.

#### OTHER PREVIOUS WORK

**Computing at Home.** I worked with A.J. Brush and Kori Inkpen at MSR to develop a new user account model for shared home computers. When configuring shared home computers, families must decide between a single shared account or multiple individual accounts. Yet users of individual accounts complain that it is difficult to share files, while users of shared accounts complain that they cannot personalize settings. The current account models are simply not designed for this environment: privacy from other users is less of a concern and the home computer is frequently used for ad-hoc tasks that do not require personalization, much like a kiosk. We created a new account model designed for the home environment where files are shared by default but can be made private, unlike the traditional model where files are private by default. Users also do not need to decide whether to switch accounts at the beginning of the session, they could switch profiles at any time without suspending applications. We conducted a laboratory study and found that our model provided both personalization and sharing, and that these needs heavily vary based on context [2]. This work is increasingly relevant as more and more shared computers appear in homes, especially in the form of set-top boxes.

<sup>1</sup><http://www.privacyfinder.org/>

**Backup Authentication.** I have also enjoyed a fruitful and ongoing collaboration with Stuart Schechter at MSR. We examined “last resort” authentication systems: if a user forgets her webmail password, she must answer a personal question to reset her password. We performed a study to show that these questions are commonly forgotten by the account owners and easily guessable by attackers [6]. As a follow up, we created a new backup authentication system that specifically addresses the human factors and tested it in the laboratory and the field. We created a *social authentication* system whereby account owners assigned account “trustees”—a set of four people who they could contact if they ever forgot their passwords. Once the reset process was initiated, each trustee received a secret code. The account holder was required to retrieve a subset of these codes from her trustees in order to reset her account password. We conducted a usability study to examine whether participants would be able to use this system successfully and how susceptible it was to several likely attacks [7]. The resulting authentication mechanism is being considered for deployment by Hotmail.

### CURRENT AND FUTURE WORK

My main research area has focused on and will continue to focus on improving the interfaces that users interact with when confronted with security and privacy decisions. Due to the broad nature of my research, I believe I am well suited for a variety of research topics related to how humans make privacy and security decisions. Currently, as a postdoctoral researcher at Brown University, I have been studying how access control interfaces lead to security breaches when they fail to account for human factors, and how these interfaces can be improved to help people make better choices by better understanding risks. My planned future work focuses on using economic analyses to examine how people perceive and react to security threats.

**Access Control Interfaces.** Previous work has shown that unusable access control interfaces can lead to either confidential data being released to unauthorized parties or legitimate parties being denied from accessing resources they need to complete a task. Researchers have attempted to improve these interfaces by creating new ways of visualizing effective permissions. However, many of these solutions have simply redesigned the interfaces without a firm understanding of how most policy errors are introduced. To gain this understanding, I am in the process of conducting an ethnographic study of workplace computer users. My goal is to better understand how they share files with each other, the types of errors they encounter, and how they mitigate those errors. The results of this research will be used to improve a new policy authoring interface that I am developing along with Shriram Krishnamurthi that minimizes access control errors by creating an interactive experience that detects and corrects ambiguities. In the next two years I expect to conduct user studies on this new interface across various domains, ranging from workplace file sharing to privacy controls on social networking websites.

**Rational Security Decisions.** In my previous research on security warnings, I discovered that risk perceptions were the biggest factor in predicting a user’s behavior; users who have incorrect risk perceptions are likely to make poor trust decisions. Over the next five years I expect to perform research to understand how users assess privacy and security risks, as well as how developers can influence these risk assessments to aid users in making better trust decisions. In the first study that I have planned, I am collaborating with researchers specializing in the economics of network security, behavioral economics, and computer security to examine whether hyperbolic discounting applies to computer security. We are specifically examining whether people will accept an immediate performance cost when they believe it is for “security purposes,” and the extent to which their decisions are a result of bounded rationality. This first study will be followed up by additional studies over the course of two years to examine whether people will accept security

delays when given the choice, and the extent to which they will tolerate those delays when they do not understand the underlying threat models from which they believe they are being protected.

As a corollary to studying how end-users make security decisions, I am also planning a study on how software developers make security decisions, as well as the economic impacts of these decisions. Information asymmetries exist between developers and users, which results in developers offloading security decisions to the users because they believe the users are in a better position to make these decisions (e.g., accepting a self-signed certificate, determining if a website is fraudulent, etc.). However, in practice we know that this is not the case: while users may have access to more information, they are usually unqualified to properly evaluate that information to arrive at a better decision than the developer could have made with less information. I plan to conduct a series of studies to examine the economics of these types of decisions and the externalities that they create. The goal of this research is to help developers and users alike make optimal risk assessments.

**Additional Areas.** In addition to the above future research directions, I believe my background and interests give me great flexibility for future research. In the past I have collaborated with faculty and industry researchers in such diverse areas as human-computer interaction, behavioral economics, social psychology, information retrieval, networking, and computer security. In the future I expect to build and maintain collaborations across these and other disciplines.

#### REFERENCES

- [1] L. F. Cranor, S. Egelman, S. Sheng, A. M. McDonald, and A. Chowdhury. P3P Deployment on Websites. *Electronic Commerce Research and Applications*, 7(3):274–293, Autumn 2008.
- [2] S. Egelman, A. J. B. Brush, and K. M. Inkpen. Family accounts: A new paradigm for user accounts within the home environment. In *Proceedings of the 2008 Computer Supported Cooperative Work Conference (CSCW'08)*, San Diego, CA, 2008.
- [3] S. Egelman, L. F. Cranor, and J. Hong. You've been warned: an empirical study of the effectiveness of web browser phishing warnings. In *CHI '08: Proceeding of the 26th SIGCHI Conference on Human Factors in Computing Systems*, pages 1065–1074, New York, NY, USA, 2008. ACM.
- [4] S. Egelman, J. Tsai, L. Cranor, and A. Acquisti. Timing is Everything? The Effects of Timing and Placement of Online Privacy Indicators. In *CHI '09: Proceeding of the 27th SIGCHI Conference on Human Factors in Computing Systems*, New York, NY, USA, 2009. ACM Press.
- [5] J. Gideon, S. Egelman, L. Cranor, and A. Acquisti. Power Strips, Prophylactics, and Privacy, Oh My! In *Proceedings of the 2006 Symposium on Usable Privacy and Security*, pages 133–144, 12-14, July 2006.
- [6] S. Schechter, A. B. Brush, and S. Egelman. It's no secret. measuring the security and reliability of authentication via. In *Proceedings of the 2009 IEEE Symposium on Security and Privacy*, pages 375–390, Los Alamitos, CA, USA, 2009. IEEE Computer Society.
- [7] S. E. Schechter, S. Egelman, and R. Reeder. It's not what you know, but who you know: A social approach to last-resort authentication. In *CHI '09: Proceeding of the 27th SIGCHI Conference on Human Factors in Computing Systems*, New York, NY, USA, 2009. ACM Press.
- [8] J. Sunshine, S. Egelman, H. Almuhiemedi, N. Atri, and L. F. Cranor. Crying wolf: An empirical study of ssl warning effectiveness. In *Proceedings of the 18th USENIX Security Symposium*, 2009.
- [9] J. Tsai, S. Egelman, L. Cranor, and A. Acquisti. The effect of online privacy information on purchasing behavior: An experimental study. In *Proceedings of the 2007 Workshop on the Economics of Information Security (WEIS'07)*, Pittsburgh, PA, USA, 2007.