

Social Cybersecurity: Applying Social Psychology to Cybersecurity

Jason Hong, Sauvik Das, Tiffany Hyun-Jin Kim, Laura Dabbish
Human Computer Interaction Institute, Carnegie Mellon University

Summary

While there are certainly cases where new theories are needed, we have not yet taken full advantage of existing theories. As one example, our team has recently been taking a strong theory-driven approach in looking at how to apply ideas from social psychology to improve cybersecurity. Past research investigating the human side of cybersecurity has treated people as isolated individuals rather than as social actors acting within a web of social influence. In our research we are (a) investigating how social influence can and does affect cybersecurity, and (b) developing and evaluating novel social influence techniques to improve people's awareness and knowledge of cybersecurity, as well as motivation to act securely. We use our experiences as a lens for discussing the ongoing role of theory in HCI.

Introduction

To date, the majority of cybersecurity research has focused on the computer itself. Recently, there has been a growing work in human factors of cybersecurity, but this work primarily views users as isolated individuals. Social psychology offers a complementary view, with a rich body of work documenting how an individual's attitudes and behaviors are strongly affected by others. Subtle and powerful social forces, such as social norms, can have outsized influences on people's behaviors and perceptions of risk.

For example, normative group influences motivate individuals to conform to group standards of behavior in order to 'fit in'. There is ample evidence that we can leverage social influence to encourage positive behavior change. In one powerful example, Nolan et al. [4] found that telling people that many of their neighbors were saving power increased conservation far more than other non-social interventions (e.g. telling them that saving power was good for the environment, that it would benefit society, or that it would save them money). This result held in both the short and long term. Goldstein et al. [5] found that simply by telling hotel guests that the last person who had stayed in that room opted to reuse their towels increased hotel room towel reuse rate by 28%.

The interventions above are simple but significantly changed people's behavior. In our own work, we are investigating how to design technology that leverages social influence to encourage safer cybersecurity behaviors. We are developing a new science of social cybersecurity that can be used to (a) make people more *aware* of cybersecurity problems, (b) help them understand what *actions* they can take to mitigate security problems, and (c) *motivate* them to take action to be more secure in practice.

Our Research on Social Cybersecurity

We have been conducting three complementary streams of research to probe the social dimension of cybersecurity: 1) analyzing social influences on current cybersecurity behavior through interviews, surveys, and log analysis 2) conducting experiments on theoretically motivated design interventions, and 3) developing novel social cybersecurity tools. Due to space, we will only discuss the first two.

Social Influences on Current Cybersecurity Behavior

We are conducting interviews and surveys to better understand *how people learn about new security behaviors* from their social network. In our first study of this type we used semi-structured interviews to ask people of various ages and backgrounds about their security-related behavior changes and communications [1]. We identified a set of reference security behaviors and then asked people a series of questions about each: if they engage in the behavior (or not) and why, how they first learned about it, and which of their friends have similar (or dissimilar) behaviors. For example, multiple surveys have found that only about half of people use any kind of authentication on their smartphones. For people that use authentication, we probed about why they use it and how they first learned about it.

We found that social factors were key drivers of security-related behavior change, accounting for nearly half of all reported behavior changes (e.g., using a PIN on one's phone or enabling a Facebook security tool). The most prevalent social catalyst for security-related behavior change was observing friends—people often started using security tools after observing friends and/or strangers use those same tools. Unfortunately, few security tools are built for this form of passive observability, and are thus unable to spread in this powerful and social way. We also learned that communications about security are scarce,

with many participants, including several security experts, reporting that they did not talk much about security at the risk of being boring or sounding preachy. The few conversations about security that did happen, though, were mostly to warn others of security threats or to teach others about how to protect themselves from a threat. Taken together, it seems that people feel accountable for the security of their loved ones but have few options to act on these feelings. We are complementing these interviews with surveys that validate our findings with a broader sample of participants.

Our results offer a better understanding of why people use certain security features and how this use is influenced by their social connections. It also provides a baseline for how much people employ different security behaviors. With this information, we can design social manipulations that can improve desired behaviors. For example, we find that people often learn about cybersecurity through stories from friends. This finding suggests we could encourage authentication by presenting a story about a friend who had their account compromised because they did not use phone authentication.

Examining Diffusion of Cybersecurity Behavior in the Field

Based on social influence theory and the results of our first study, we tailored interface designs in the Facebook system to manipulate social influence dimensions and increase the likelihood of secure behavior. We have conducted a series of experiments integrating social influence tactics into the design of interactive systems and examined their influence on cybersecurity behavior.

We analyzed how the adoption of three Facebook security tools—Login Notifications, Login Approvals and Trusted Contacts—diffused through the social networks of 1.5 million people [2]. We employed a “matched propensity sampling analysis” to determine whether exposure to more friends who used a security tool predicted for whether one was more or less likely to adopt that security tool. The social psychology literature suggests that people who are exposed to more friends who use security tools should be more likely to use those tools themselves. Our results, summarized in Figure 1, suggest that social influence indeed affects the adoption of security tools, but not always in a positive direction.




Figure 1. Matched propensity sampling analysis results. All differences significant ($p < 2e-16$).

Specifically, using Rogers’ Diffusion of Innovations theory and our own prior qualitative work as a lens, we hypothesized that Trusted Contacts had two advantages in its potential for social spread, over Login Notifications and Login Approvals. First, its use is more observable. Whereas Login Notifications and Login Approvals are private, Trusted Contacts requires a user to specify 3-5 friends to help with account recovery. These “trusted contacts” are, in turn, notified that they have been entrusted with this role and thus its use is broadcast. Second, Trusted Contacts is more socially compatible. Whereas Login Notifications and Login Approvals are used to exclude others from access and may thus be indicative of distrust [1], Trusted Contacts lets friends provide security, and may thus be more indicative of trust. To summarize, it seems that security tool adoption does depend on social influence, but only positively for tools that are observable, socially compatible, and/or widely adopted among one’s own social network.

Experiments on Social Cybersecurity

We are also conducting experiments to enumerate and empirically validate the design dimensions of security tools that affect its potential for social spread. The outcome will be a “social checklist” for security tool designers to evaluate whether their tools adequately consider these social design dimensions. In these experiments we are testing designs motivated by diffusion of innovation theory. Our experiments on cybersecurity diffusion demonstrate application of theory for practical benefit.



Figure 2. One of the tested social announcements, showing viewers the number of their friends who already used security tools on Facebook.

We collaborated with Facebook’s Site Integrity team to put the theoretical findings from our prior work into practice. We tested whether increasing observability of security tool usage would increase awareness and adoption of security tools [3]. Initially, the Site Integrity team had planned to increase security tool adoption by showing some Facebook users a simple message: “You can use security settings to protect your account and make sure it can be recovered if you ever lose access.” This announcement contained a call-to-action button for activating Login Notifications, Login Approvals, and Trusted Contacts. For a subset of 50,000 people, we added some social information that increased the observability of a viewer’s friends’ security tool usage. Figure 2 shows one such modification, where we prepended the number of the viewer’s friends who used one of the promoted security tools. We then measured whether the addition of this social information increased clicks on the announcements and tool adoptions.

The results were clear: the social announcements both significantly and substantially increased both clicks and adoptions relative to the announcement with no social information. The best performing social condition (Figure 2) increased clicks by 37% (14.4% vs. 10.5% of viewers) and adoptions of any one of the three promoted security tools, in the 7-day period after the announcement was shown, by 30% (4.8% vs. 3.7% of viewers). In other words, this simple example of making security more social significantly increased awareness and adoption of security tools on Facebook.

In future tool-building work, we also plan to create security tools that encourage social accountability—e.g., tools to audit the security of friends and loved ones. We also want to explore creating security tools that are trusting of friends and loved ones by making it easy for them to gain *limited* access to one’s data and devices. Finally, we want to explore if providing social information *in context* will bolster the effect of social influence on security tool adoption—that is, providing information about friends who use security tools in the wake of a security breach.

Reflecting on the Role of Theory in HCI

Here, we reflect more on the role of theory in HCI. First, we can clearly improve user interfaces by doing hundreds of A/B tests without any grounding in theory. However, this doesn’t address the issue of what should be tested, nor does it offer insight about broader design principles. The famed mathematician William Thurston opined that the controversy around the computation-oriented proof of the 4-color map theorem was not about the proof itself, but rather the desire for deeper human understanding [6]. In many ways, the push for theory in HCI is the same: we want (and need) deeper and generalizable ways of understanding human behavior and designing interactive systems, to avoid doing a blind search in large design spaces. Note that blind searches are a perfectly fine strategy when initially exploring a new design space. We often see this at CHI in the form of point solutions, which push the bounds of what is possible and help define the design space. We also note that the role of theory for HCI systems work is unclear.

Second, existing predictive theories can be a source of inspiration for longstanding problems. Every discipline embodies a set of methods and heuristics, and after enough time (and dollars!), we will hit diminishing returns repeatedly applying them. Thus, it makes sense for theory-driven researchers to apply their ideas to other domains, and for researchers in a specific domain to draw on new theories. However, to really succeed, it takes a deep interdisciplinary approach. It’s one thing to know about a theory, it’s quite another to apply it. This of course leads to questions about how we train and organize researchers.

Third, HCI is in the unusual position of developing both better understanding of phenomena and better ways of building things. We argue that the best theories can push in both of those dimensions, in a manner that is easy and robust enough for practitioners to apply. In many ways, this argument is similar to Pasteur’s Quadrant [7], that research should lead to good science and good applications. Perhaps a good rule of thumb is: could we (and would we want to) teach it to our undergrads?

About us

Jason Hong is an associate professor in the Human Computer Interaction Institute, part of the School of Computer Science at Carnegie Mellon University. Jason specializes in usable privacy and security, and mobile and ubiquitous computing. Laura Dabbish is an assistant professor in the Human Computer Interaction Institute. Laura specializes in social computing and computer-mediated social interaction. Sauvik Das is a PhD student in the Human Computer Interaction Institute, and Tiffany Hyun-Jin Kim is a researcher at HRL Laboratories.

References

1. Das, S., Kim, T.H.J., Dabbish, L., & Hong, J.I (2014). The Effect of Social Influence on Security Sensitivity. In Proc. SOUPS'2014.
2. Das, S., Kramer, A.D.I., Dabbish, L., & Hong, J.I. (2015). The Role of Social Influence in Security Feature Adoption. To appear In Proc. CSCW'2015.
3. Das, S., Kramer, A.D.I., Dabbish, L., & Hong, J.I. (2014). Increasing Security Sensitivity With Social Proof: A Large Scale Experimental Confirmation. In Proc. CCS'2014.
4. Nolan, J.M., et al. Normative Social Influence is Underdetected. *Personality and Social Psychology Bulletin* 2008; 34; 913. <http://www.csom.umn.edu/assets/118360.pdf>
5. Goldstein, N., Cialdini, R., & Griskevicius, V. (2008). A Room with a Viewpoint: Using Social Norms to Motivate Environmental Conservation in Hotels, *J. Consumer Research*, 35. <http://www.csom.umn.edu/assets/118359.pdf>
6. Thurston, W. On Proof and Progress in Mathematics. *Bulletin of the American Mathematical Society*; 30(2); 161-177. <http://arxiv.org/pdf/math/9404236v1.pdf>
7. Stokes, D.E. *Pasteur's Quadrant: Basic Science and Technological Innovation*. Brookings Institution Press. 1997.