Bridging the Gap in Computer Security Warnings

A Mental Model Approach

Computer security warnings are intended to protect users and their computers. However, research suggests that users frequently ignore these warnings. The authors describe a study designed to gain insight into how users perceive and respond to computer warnings.

Cristian Bravo-Lillo, Lorrie Faith Cranor, Julie S. Downs, and Saranga Komanduri
Carnegie Mellon University

Warnings are a form of communication designed to protect people from harm.[1] Psychologists have studied physical-world warnings since the early 20th century.[2] The psychological processes involved in paying attention to warnings, grasping their meaning, and deciding to comply with them haven't changed substantially, even in the digital realm. An effective physical warning clearly communicates risk, consequences of not complying, and instructions to comply (although some of this information can be omitted if the risk is obvious or the consequences can be deduced from the warning).[1] However, many of the most common computer warnings fail to follow one or more of these guidelines. For example, in Figure 1, the warning dialog doesn't explain the risk (the file might be infected with malware) or consequences (information might get corrupted, erased, or disclosed to third parties), and it doesn't instruct users on how to avoid the risk (either delete attachment or save it on your hard disk and scan it with your antivirus software).

Besides protecting people from harm, warnings are also intended to modify behavior to comply with existing safety regulations; to decrease the likelihood of health problems, accidents, or property damage; and to serve as reminders. However, warnings aren't the most effective method for protecting people from hazards, and should be used only as a third line of defense, after considering ways of designing out hazards and guarding against them.[1] Consider a hazardous broken sidewalk. You could repair (design the risk out) or put a barricade around it (guard against the risk). You could post warning signs as an interim solution, but they shouldn't be the only safeguard.

However, in some situations, designing out a hazard or guarding against it might not be feasible. For example, the sharp edge of a knife can't be designed out without making the knife useless, and guarding against the risk of cutting oneself isn't practical. Similarly, the risk of being phished by a malicious website can't be completely designed out, although users could employ guarding strategies such as automatically detecting and removing suspicious links from email.

An additional difficulty with digital risks is that they're less understood, and analogies to the physical world can be incomplete or misleading. If users don't know what phishing is, they won't be able to assess whether they're at risk. Here, users rely on a very rudimentary form of learning. If choosing a particular option lets them continue their work unhindered, then they might choose that option every time, especially if they don't really understand the risk that they're being warned about. Even experts dismiss warnings when they understand them because they find them unimportant.

Regardless of people's expertise level, warnings don't seem to be doing what they are supposed to be doing—stopping people from engaging in unsafe behaviors. To improve users' understanding of warnings, we first need to determine how users

 1540-7993/11/$26.00 © 2011 IEEE

process the information in them, that is, how they think about warnings. For this purpose, we conducted 30 interviews—10 with advanced users in security and privacy and 20 with novice users. We categorized and coded their answers and used these codes to create a mental model diagram that illustrates the knowledge gap between novice and advanced users.

Study Methodology

We collected examples of 29 security warnings from popular operating systems and application software and categorized them into four warning types: *information deletion or loss*, *information disclosure*, *execution of malicious code*, and *trust in malicious third parties*. We picked one to two warnings from each category: a disk space warning, an email-encryption warning, an address book disclosure warning, an email attachment warning (see Figure 1), and a certificate warning. We created at least one scenario per warning in which we briefly described a situation that provided context for the warning's appearance.

Recruiting Process

Our mental model studies typically include 20 to 30 participants. This sample size is large enough to be likely to reveal at least once any belief held by 10 percent or more of the population.[3] In this study, we don't make inferential statements about quantitative differences between groups, and thus we don't need a formal power analysis. We plan to conduct follow-up studies with larger numbers of participants to test hypotheses that emerge from this study.

We recruited our 10 advanced users by direct email invitations sent to two mailing lists at Carnegie Mellon University. Participants were between 22 and 63 years old (average = 30.7, σ = 11.8), and included two faculty members, five computer security doctoral students, two research programmers, and one information security researcher. Advanced users were considered as such if they had either taken at least one computer security or privacy graduate course or had worked on computer security or privacy projects for at least one year. Most of our advanced participants had multiple years of security course work or experience. Past studies have found that even lower levels of expertise are sufficient for making significantly better security decisions, for example, in the context of phishing.[4] We gave all advanced users US$10 and a chocolate bar as compensation for their time.

We recruited novice users through messages posted on Craigslist and flyers posted in bus stops around the university, which directed respondents to an online screening survey. We excluded those who worked in any field related to computer security or privacy or who had taken at least one college-level course in computer security. We selected 20 participants for our interviews. Their ages ranged from 18 to 57 years old (average = 32.6, σ = 11.6), and their occupations were diverse: seven students, six employees or supervisors in different industries, three professional musicians, two self-employed, and two unemployed persons. All novice users received US$20 as compensation for their time.

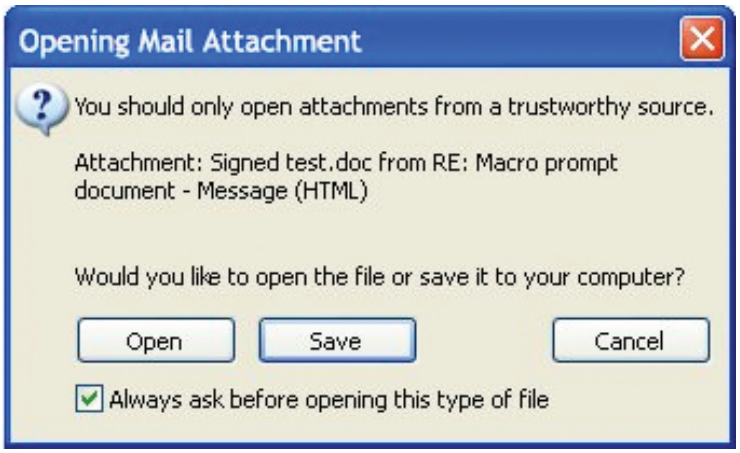

Figure 1. Attachment warning. The dialog doesn't explain the risk or consequences, or contain instructions on how to avoid the risk.

Interviews

We conducted one-on-one, open-ended interviews with advanced and novice users. At the beginning of each interview, we told all participants that we were interested in knowing how they made use of computers and software, and that we weren't looking for any particular answer. Interviews had seven segments: a brief general section about computer use, five sections that asked about warning reactions, and a final segment about demographics. In each warning segment, we showed a warning dialog and read aloud a brief scenario that described a nontechnically savvy friend asking the participant for help. We then asked the following main questions (and other questions not shown):

- Could you tell me what this message is?
- What do you think will happen if your friend clicks on X? (We asked for all the options present in the warning.)
- What do you think your friend should do?

To understand their thought processes, we asked participants to explain their reasoning and any terms that they used and followed up on any interesting observations. We transcribed the audio recordings of the interviews verbatim.

Two investigators read five advanced users'

Related Work in Computer Warnings

In the physical world, people pay sporadic attention to warnings and are particularly likely to ignore those that don't map well onto a clear and understandable course of action.[1] Similarly, evidence from experimental studies indicates that most people don't read computer warnings,[2,3] don't understand them,[4] or simply don't heed them,[5] even when the situation is clearly hazardous. Researchers have offered variety of explanations for this behavior. For one, people's trust in computer systems might cause them to underestimate the risks.[6] They might also be unaware of some risks,[4] or they might not understand the risks behind a warning. In one study, 32 percent of people who heeded a phishing warning attributed the warning to a Web problem and still believed that phishing emails sent to them were legitimate.[2] The authors suggest that participants "had very inaccurate mental models of phishing," in part because they didn't understand that the email that took them to the phishing website could have been spoofed. Another explanation is that users do understand presented choices but that they also consider cues that are external to the system; for example, the trust they put in the sender of an email with a potentially dangerous attachment. Further, they weigh the trust they put into the system against their desire to continue with their primary task.[6] Whatever the reasons, it's clear that we must understand what users think and believe about warnings to help them make safer choices.

Despite the problems we described, software designers sometimes rely on users to perform important security tasks, including judging whether or not to heed a warning. Lorrie F. Cranor has proposed a general model to aid in understanding and designing out security problems that might arise from the interaction between humans and software systems—the human-in-the-loop (HITL) framework.[7] This framework is based on a more general model of human cognition, the communication-human information processing (C-HIP) model developed by Michael Wogalter to describe the sequential processing that occurs when users encounter warnings.[8] Both models describe a set of sequential stages with feedback loops that users might experience, with flow of information or processing from one stage to the next, until a change of behavior attributable to the warning occurs.

Cranor distinguishes five different forms of security communications in her model: warning dialogs, notices, status indicators, training, and policies. The HITL model applies to all five, but our research in this paper focuses specifically on warning dialogs. In Figure A, a communication is delivered to users in the form of a warning.[7] Assuming that the communication hasn't been interfered with or distorted before reaching the users, the warning is processed in several steps. Users might or might not switch their attention to the warning. If they do, they must attend to the warning long enough to comprehend its meaning. If they grasp the meaning, they must acquire and retain the warning information and apply it to the current situation. The process ends when some behavior attributable to the warning is observed. The whole sequence can be modulated or even completely overridden by the users' intentions, capabilities, or personal variables, which include previous knowledge and past experience.

Existing literature about computer warnings addresses only some of the dimensions in the HITL model. For example, a study by José Brustoloni and Ricardo Villamarín-Salomón used two approaches to encourage users to make safer choices when managing their email and attachments.[9] The first approach was to warn users that their actions were being audited by a human observer who might impose penalties on them. In this condition, people made better security judgments related to their email, showing that increasing users' motivation is possible (see the intentions box in Figure A). However, auditing people's actions is resource consuming and requires an organizational context that home users don't have.

The second approach was to randomly reorder the warning options in each presentation, forcing the user to actually read the options presented, thus increasing attention. Although this technique can be applied to home users, it neither improves warnings' quality nor puts the user in a better position to make the safest choice. In addition, the study doesn't tell us much about what happens to participants' comprehension and the subsequent stages of the model.

Research conducted by Serge Egelman and colleagues reviews most stages of the C-HIP model to explain why many participants were fooled by spear phishing messages sent to them in the lab. These targeted phishing messages appeared to have been sent from the online vendor from which participants had just made a purchase. Although this study identified specific stages of the model at which participants failed, it didn't directly address why they failed.[2] A mental model approach can be used to gain a clearer answer to this question.

The mental models approach has been used in areas such as nuclear waste management, radon pollution, and sexual disease transmission.[10] Only a few studies have applied mental models to computer security or privacy risk communication.

transcripts independently, identified common ideas, and assigned a unique code to each idea. We then compared the code lists and resolved the differences to create a single code list to be used with the remaining transcripts. Each investigator then read a new set of transcripts independently and coded new common ideas, repeating this process until no new ideas emerged (that is, no new codes were generated), which happened after having read seven advanced transcripts and 10 novice transcripts. Finally, the investigators reread all transcripts and coded them with the agreed-upon code list. They also identified

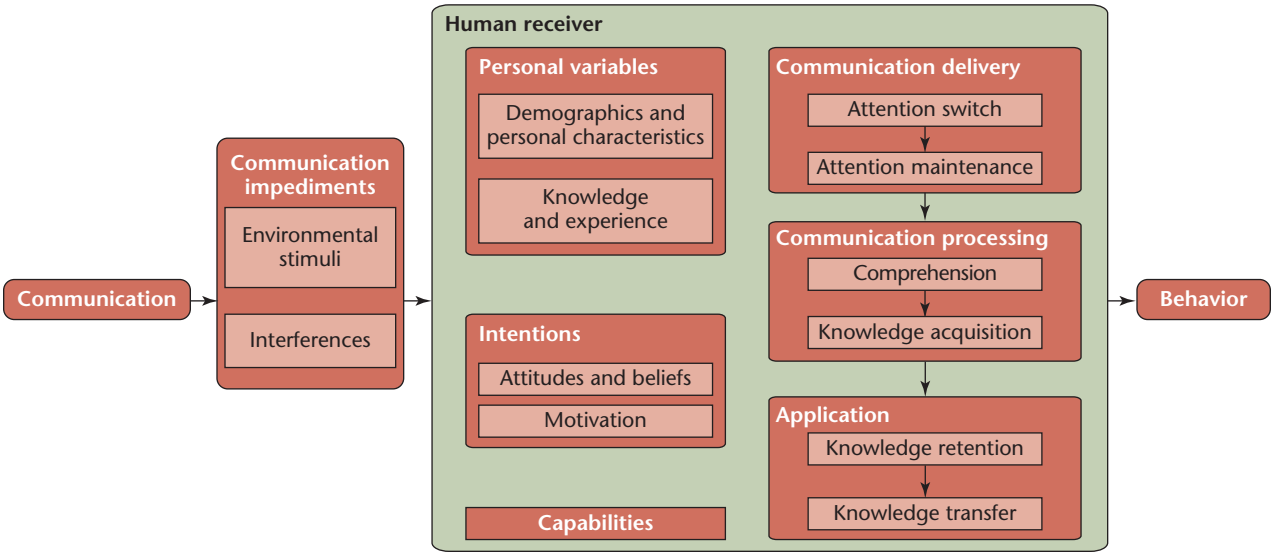

Figure A. The human-in-the-loop framework. Security communications might be subject to communication impediments before they're delivered to human receivers. The receivers must process the communication and determine how to apply it. However, the receivers' intentions, capabilities, and personal variables might impact their behavior.

L. Jean Camp describes five generic mental models that might help with delivering computer risk communication to lay users, and concludes that these models "can be used to improve risk communication," acknowledging that a user study should be performed to test these models.[11] Recently, Rick Wash identified four mental models about the notion of *hacker*, and another four about *virus*, through open-ended interviews with a similar methodology to the one we use in our study.[12] Our study focuses on people's reactions and beliefs about computer warnings.

References

1. K. Witte, "Putting the Fear Back into Fear Appeals: The Extended Parallel Process Model," *Comm. Monographs*, vol. 59, no. 4, 1992, pp. 329–349.
2. S. Egelman, L.F. Cranor, and J.I. Hong, "You've Been Warned: An Empirical Study of the Effectiveness of Web Browser Phishing Warnings," *Proc. 2008 Conf. Human Factors in Computing Systems* (CHI 08), ACM Press, 2008, pp. 1065–1074.
3. J. Sunshine et al., "Crying Wolf: An Empirical Study of SSL Warning Effectiveness," *Proc. 18th Usenix Security Symp.* (SSYM 09), Usenix Assoc., 2009; http://lorrie.cranor.org/pubs/sslwarnings.pdf.
4. J.S. Downs, M.B. Holbrook, and L.F. Cranor, "Decision Strategies and Susceptibility to Phishing," *Proc. 2nd ACM Int'l Symp. Usable Privacy and Security*, (SOUPS 06), vol. 149, ACM Press, 2006, pp. 79–90.
5. S.E. Schechter et al., "The Emperor's New Security Indicators," *Proc. 2007 IEEE Symp. Security and Privacy* (SP 07), IEEE CS Press, 2007, pp. 51–65.
6. C. Nodder, "Users and Trust: A Microsoft Case Study," *Security and Usability: Designing Secure Systems that People Can Use*, L.F. Cranor and S.L. Garfinkel, eds., O'Reilly Media, 2005, pp. 589–606.
7. L.F. Cranor, "A Framework for Reasoning about the Human in the Loop," *Proc. 1st Conf. Usability, Psychology, and Security* (UPSEC 08), Usenix Assoc., 2008; www.usenix.org/event/upsec08/tech/full_papers/cranor/cranor.pdf.
8. M.S. Wogalter, "Communication-Human Information Processing Model," *Handbook of Warnings* (Human Factors/Ergonomics), M.S. Wogalter, ed., Lawrence Erlbaum Associates, 2006, pp. 51–61.
9. J.C. Brustoloni and R. Villamarín-Salomón, "Improving Security Decisions with Polymorphic and Audited Dialogs," *Proc. 3rd ACM Int'l Symp. Usable Privacy and Security* (SOUPS 07), vol. 229, ACM Press, 2007, pp. 76–85.
10. G.M. Morgan et al., *Risk Communication: A Mental Models Approach*, Cambridge Univ. Press, 2001.
11. L.J. Camp, "Mental Models of Privacy and Security," *Technology and Society Magazine*, vol. 28, no. 3, 2009, pp. 37–46.
12. R. Wash, "Folk Models of Home Computer Security," *Proc. 6th Symp. Usable Privacy and Security* (SOUPS 10), ACM Press, 2010, pp. 1–16.

semantic relationships between ideas. Figure 2 presents the resulting mental model as a diagram of these relationships.

The Mental Model

The arrow to the left of Figure 2 shows the model's main stages—three sets of tasks that users perform after the dialog pops up. In the first set, users observe and consider any of several factors and events (variables) to try to understand what the warning message is communicating. After these observations, users attempt to diagnose the cause of the warning,

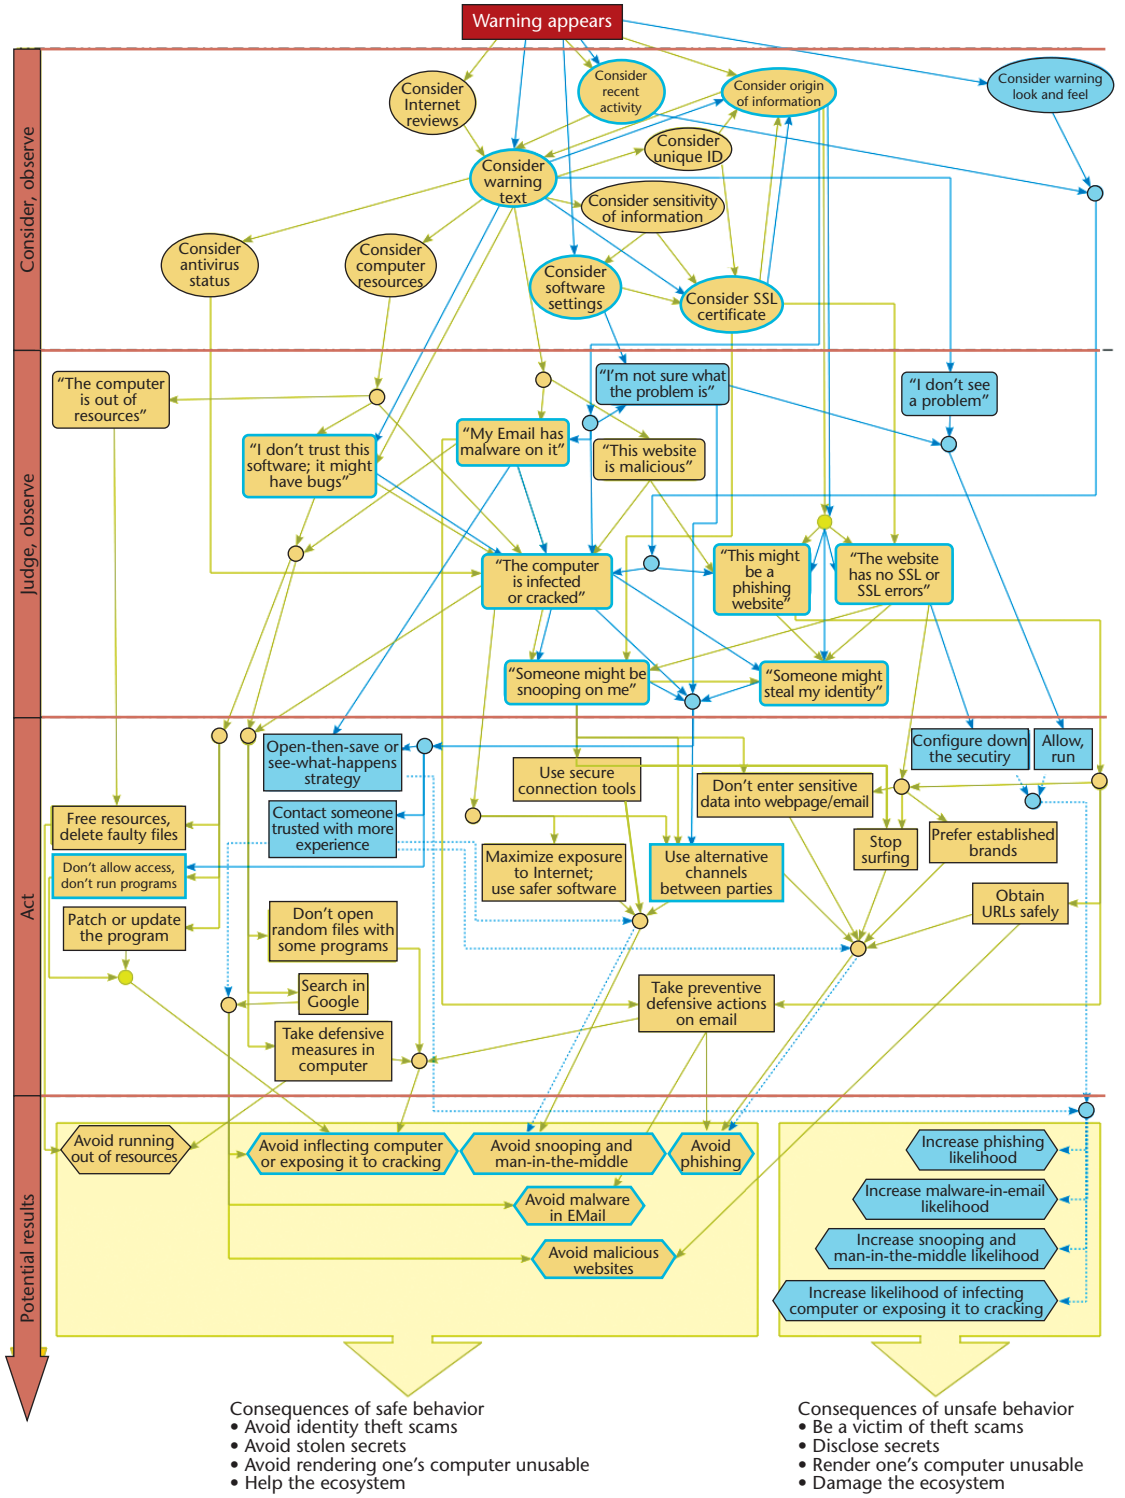

Figure 2. Our detailed mental model of warning response behaviors. Yellow items indicate advanced users' responses, and blue items represent novice users' responses. Yellow items with a blue outline were mentioned by both.

or judge which problem, out of several potential options, they think they're dealing with. Then, they take one or more actions to attempt to address the perceived problem. If the diagnosis was correct and the behavior was appropriate, then the problem is solved. Otherwise, it might persist, or another

problem might arise.

Advanced users' tasks are in yellow, and novice users' tasks in blue. Arrows depict an observed and likely relationship between two tasks. For example, immediately after a dialog pops up, novice users often consider the warning's look and feel; if they find it suspicious, then they often judge that their computer might be infected or cracked, or if they're visiting a website, that the website might be a phishing attempt. In contrast, advanced users often consider recent actions that might have prompted the warning and will search for the warning text on the Internet to determine its legitimacy.

The mental model represents a set of common lines of reasoning about computer warnings. As such, we can use it to better understand the differences between how advanced and novice users would think about a particular warning. This mental model can inform and improve a warning's design in several ways.

One way is to determine under which conditions a certain belief must be addressed before showing a warning. For example, unknown applications are hazardous due to at least two risks: the application might be a virus, and the application might access and misuse users' personally identifiable information. A smartly designed interaction would discard the first alternative by executing an antivirus program first. If the antivirus program reports that the application is free from known malware, then the warning text can be tailored toward the second alternative, and should mention that the application has been checked and is free from known viruses. Based on our interviews, this would be very helpful to novice users because they tend to relate all warnings to viruses. As Figure 2 illustrates, novice users tend to consider a warning and determine either that there isn't actually a problem, or that their computer has been infected. Although relying on an antivirus program isn't a perfect solution, it illustrates the potential of using a more holistic warning design approach to make warnings more informative and less generic.

We can also use this mental model to prioritize and deliver different messages in a risky situation. For example, at the bottom of Figure 2, most of the consequences of unsafe behavior are caused by three actions of novice users:

- configuring down the computer's global security level,
- letting unknown programs run, and
- performing a set of simple strategies we categorize as *open-then-save* or *see-what-happens*.

If the mental model can trace several possible paths and shows that one path might lead to an unsafe action, it's more important to discourage the user from taking this path than others.

Finally, insights from the model can inform the warning's content. As the top sector of Figure 2 illustrates, novice users often don't consider the sensitivity of the information they enter into emails or websites, which makes them more likely to be victims of phishing or identity theft. This suggests that the focus of phishing warnings should be on the sensitivity of the information entered into an unknown website, and not merely a vague warning that the current site might be a phishing site.

All the design insights discussed here are suggestions derived from the mental model study. However, each should be evaluated in a larger study to determine their effectiveness and gain insights into warning response behaviors.

How Advanced and Novice Users Differ

We found consistent differences between advanced and novice users' behavior. One interesting difference is that the groups observe different cues and arrive at different conclusions about the risks they might be facing. Therefore, they will take different actions that ultimately produce different outcomes.

We also observed more specific differences. For example, novice users assess the safety of an action *after* engaging in it, whereas advanced users judge how safe actions are *a priori*. Changing this behavior is probably unrealistic. However, in many cases, warnings can include a brief description of both the risks involved and the consequences of each option, thus cueing the novice user to consider this information in advance. The warning should present information in a manner that makes it available to novice users but doesn't burden advanced users.

Also, novice users consider fewer factors and perform fewer tasks to ensure their safety, whereas advanced users perform actions such as looking for vulnerabilities in public expert forums, regularly patching and updating their software, using "safe URLs" (for example, using personal bookmarks, recovering URLs with autocompletion in the browser surfing history, or typing them directly in the location bar), and taking proactive measures (for example, maintaining antivirus programs and installing security plug-ins in their browsers). Although asking a novice user to perform all these actions is unrealistic, we can gather and display useful information when necessary. For example, all warnings triggered by an email client might include a link to the online support forum maintained by the product's software vendor, rather than a link to generic help text. This would empower advanced users to contribute their

solutions to the problem that triggered the warning, and these solutions would be invaluable to novice users. A warning might check for available patches automatically or make use of heuristics to determine whether a typed email contains sensitive information. These and other strategies should be tested to evaluate their effectiveness.

A warning might check for available patches automatically or make use of heuristics to determine whether a typed email contains sensitive information.

Although some might believe that novice users simply "hate" warnings, this wasn't the case in our study: approximately half of our novice participants considered the presented warnings as a "good thing," regardless of their understanding, whereas the other half were neutral. None declared that presented warnings were "bad" or "not useful."

Misconceptions and Problems

Our novice participants' responses revealed several misconceptions about security. These misconceptions illustrate the importance of understanding users when designing security solutions. For example, six novice users reported that their interactions with banks' websites ought to be safe simply because banks have good security. In the scenario in which a bank's website produced an SSL warning, advanced participants strongly counseled against proceeding, but novice participants said (names have been changed):

> Elizabeth: I would hit yes, yes … I mean, assuming he trusts his bank. It's just, you know, the security certificate, you know, everything is valid about it, it's just you haven't elected to trust it yet, so I would feel better about hitting yes to that.

> Michael: Their site should automatically be secure because it's a bank. They're dealing with peoples' sensitive, private information like checking accounts, savings accounts, credit-card information, social security information. That stuff is sensitive, so most banks should ideally have really complex security.

Two novice participants wanted to adjust the security settings on the computer to prevent this warning from appearing because they were so sure that a bank's website would be safe.

Although advanced users agreed that banks will have good security, they wouldn't proceed to a bank's website if presented with an SSL certificate error. Advanced users were more likely to recognize the possibility that they're not truly at a bank website, whereas novice users relied on what might be fraudulent cues. As Min Wu and colleagues observed, novice users often make security judgments based on look and feel, and our data supports this.[5] When we asked a novice participant how he could tell that a warning was authentic, he said,

> James: I guess the message looks authentic in terms of just the design, the icon used, and the font and the text and the gradient for the bar up top.

Eight of our novice participants cited the warning's appearance as a factor in deciding to trust it, in contrast to advanced users, who advised that appearance should be used only to decide *not* to trust a warning, and never to confirm trust.

Another misconception had to do with opening and saving files. When a warning dialog presented a choice between opening or saving a file, advanced users felt that saving the file was safer because it can be scanned for malware before execution. By contrast, seven of our novice users felt that saving the file was more dangerous, because this permanently stored the file on the computer. They thought that opening the file only displayed a preview and was safe:

> Melissa: I would actually advise him to press Open if he really wanted to see the chain email because if you save a file that you're not sure would be safe or reliable, it's safer to open instead of save when you're dealing with something that you're not sure is reliable.

Four additional novice users perceived no difference between opening and saving files. These users felt that malware would activate either way. Joseph compared a suspicious email attachment to a time bomb:

> Joseph: Okay, a bomb or anything, I'm saying, okay, explode. Saving something, maybe I'm asking it to explode later.

Technical jargon is a common problem in computer warnings. Novice users often don't understand technical terms, and this certainly impedes their comprehension of warnings. We used warning dialogs containing terms such as *startup disk*, *encryption*, *virus*, *attachment*, *macro*, and *certificate*. Our participants had heard of, but not understood, these terms and struggled to make sense of them:

> Stephanie: I don't know whether if you send

> something that's unencrypted, does that mean that they can get into your whole computer and see everything? I don't, I don't know that. Can they see all your passwords and everything, everywhere you've been? You mean if something's unencrypted, is it just the message or is it your whole computer that's kind of see-through? I don't know.

SSL certificates turned out to be the most confusing concept in our study—16 novice users made incorrect statements about them. The SSL certificate warning in this study indicated a website with a certificate that couldn't be verified. However, novice users associated this warning with antivirus software, security updates, or website certifications about being "virus-free":

> Michael: Certificates are if you want to see or view how strong someone's computer security is from viruses. Basically, certificates say this is how you have programs like McAfee and all these different, like, Norton Antivirus programs. … They're basically kind of a security guard against viruses.
>
> John: Oh, just, like, it has a valid name, a valid website, and it won't contain any harmful software, virus, or something else, and it could be trusted by any user or any other website.
>
> Robert: It is like, almost like a credential or like a plug-in that allows you to use software, and it means your security is up to date on your computer.
>
> Melissa: I guess it just proves how authentic a website is, whether (pause) I don't know how much the government plays, like, how much it monitors websites, but I'm expecting that it's a certificate from the government or company that says that website doesn't have any viruses, or that it's run by respectable people.

Neither the warning dialogs we showed to our participants nor the brief scenarios we presented along with each dialog contained the word *virus* or *security*. We believe novice participants used an availability heuristic, assuming that viruses must be involved in any computer security context.[3,6]

This study provides qualitative insights into how novice and advanced users make sense of warnings. When presented with a warning, advanced and novice users observed different sets of cues, came to different diagnoses of the underlying risks, and consequently responded in very different ways.

This study suggests that to improve warnings' design, developers should consider all steps of warning processing. Previous studies have considered factors such as attention and motivation, which might improve users' awareness of different cues. However, our findings suggest that warnings should also deal with wrong diagnoses by indicating, for example, when a specific condition wasn't produced by a certain problem (for example, novice users tend to overdiagnose virus problems).

There is a trade-off between the amount of information included in a warning and the added likelihood that this new information might help users make the appropriate decision.[1,7] We observed that participants in our study often didn't thoroughly read warnings; giving them more text to read might worsen the problem.

To ensure that warnings are presented only when necessary, and then with only the necessary information, we should insist on applying the fundamental warning design principle: only present a security warning prompt when designing out or guarding against the risk are infeasible. In addition, we should only present warnings in situations in which the best course of action depends on details of the situation that are known to the user.

We used a mental model to highlight more effective ways to convey security information to the average user in response to immediate problems. However, it's also possible to make a more proactive use of these models to determine, for example, how to better employ users' time in security education or training. Mental models have myriad different possible and unexploited applications in the study of usable security; one unexplored possibility is their use as Bayesian belief networks to automatically estimate the probabilities of the different possible risks. This might help developers in implementing heuristics to determine when a warning is likely to be helpful. □

Acknowledgments

This research was funded in part by NSF grant CNS0831428. We thank Mandy Holbrook, Greg Norcie, and Manya Sleeper for their assistance with this study.

References

1. M.S. Wogalter, "Purposes and Scope of Warnings," *Handbook of Warnings* (Human Factors/Ergonomics), M.S. Wogalter, ed., Lawrence Erlbaum Assoc., 2006, pp. 3–9.
2. S.R. Bohme and D. Egilman, "A Brief History of Warnings," *Handbook of Warnings* (Human Factors/Ergonomics), M.S. Wogalter, ed., Lawrence Erlbaum Assoc., 2006, pp. 11–20.
3. G.M. Morgan et al., *Risk Communication: A Mental*

Models Approach, Cambridge Univ. Press, 2001.
4. S. Sheng et al., "An Empirical Analysis of Phishing Blacklists," *6th Conf. Email and Anti-Spam*, 2009; http://ceas.cc/2009/papers/ceas2009-paper-32.pdf.
5. M. Wu, R.C. Miller, and S.L. Garfinkel, "Do Security Toolbars Actually Prevent Phishing Attacks?" *Proc. Conf. Human Factors in Computing Systems* (CHI 06), ACM Press, 2006, pp. 601–610.
6. L.J. Camp, "Mental Models of Privacy and Security," *Technology and Society Magazine*, vol. 28, no. 3, 2009, pp. 37–46.
7. M.S. Wogalter, "Communication-Human Information Processing Model," *Handbook of Warnings* (Human Factors/Ergonomics), M.S. Wogalter, ed., Lawrence Erlbaum Assoc., 2006, pp. 51–61.

Cristian Bravo-Lillo *is a computing engineer and a PhD student in engineering and public policy at Carnegie Mellon University. He is currently researching software usability and security. His interests range from usable security to software engineering education and free software. Bravo-Lillo has a BS in computer science from Universidad de Chile. Contact him at cbravo@cmu.edu.*

Lorrie Faith Cranor *is an associate professor of computer science and of engineering and public policy at Carnegie Mellon University, where she is director of the CyLab Usable Privacy and Security Laboratory (CUPS). She is also chief scientist of Wombat Security Technologies. Her research interests include usable security and privacy. Cranor has a DSc in engineering and policy from Washington University in St. Louis. She's a senior member of IEEE and the ACM. Contact her at lorrie@cs.cmu.edu.*

Julie S. Downs *is the director of the Center for Risk Perception and Communication at Carnegie Mellon University's Department of Social and Decision Sciences. Her research interests include how social influences affect decision-making and how people can make better decisions by understanding the nature of these influences. Downs has a PhD in social psychology from Princeton University. Contact her at downs@cmu.edu.*

Saranga Komanduri *is a doctoral student at Carnegie Mellon University's School of Computer Science. His research covers a broad spectrum of security-related topics, including authentication, usable security, and warnings. Komanduri has an MS in computer science from Bowling Green State University. He's a member of the ACM. Contact him at sarangak@andrew.cmu.edu.*