

Influencing Mental Models of Security: A Research Agenda

Rick Wash
Communication Arts and Sciences
Michigan State University
East Lansing, MI USA
wash@msu.edu

Emilee Rader
Communication Arts and Sciences
Michigan State University
East Lansing, MI USA
emilee@msu.edu

ABSTRACT

Over 80 million households in the United States have a home computer and an Internet connection. The vast majority of these are administered by people who have little computer security knowledge or training, and many users try to avoid making security decisions because they feel they don't have the knowledge and skills to maintain proper security. Nevertheless, home computer users still make security-related decisions on a regular basis—for example, whether or not to click on a shady link in an email message—without even knowing that's what they are doing. Their decisions are guided by how they think about computer security, or their “mental models,” which do not have to be technically correct to lead to desirable security behaviors [44]. In other words, *sometimes even “wrong” mental models produce good security decisions*. By eliminating the constraint that non-technical users must become more like computer security experts to properly protect themselves, we believe that we can create more effective ways of helping home computer users make good security decisions. To that end, we propose a research agenda that will help us learn how to shape the mental models of regular non-technical computer users.

Categories and Subject Descriptors

K.6.5 [Management of Computing and Information Systems]: Security and Protection—*Invasive software, unauthorized access*; K.4.2 [Computers and Society]: Social Issues—*abuse and crime involving computers*; H.5.2 [Information Systems Applications]: User Interfaces—*user-centered design*

General Terms

Security, Human Factors

Keywords

home computer security, mental models, human subjects, socio-technical system

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

NSPW'11, September 12–15, 2011, Marin County, CA, USA.
Copyright 2011 ACM 978-1-4503-1078-9/11/09 ...\$10.00.

1. INTRODUCTION

There are over 80 million computers located in households in the United States [10]. A large proportion of these home computers frequently fall victim to various security threats, including contracting viruses, becoming botnet zombies, and being compromised by phishing scams. Threats to home computer security comprise a major epidemic right now; this insecurity is causing thousands of hours and millions of dollars in lost time and energy, and is dramatically increasing the difficulty of having a computer in the home. And, due to recent innovations by hackers (botnets and DDoS), these security problems are now spilling over and causing problems for many other sectors of the economy. Just recently, a hacker group calling itself “Anonymous” used DDoS attacks against PayPal and Amazon in retaliation for dropping support for Wikileaks. It would not be surprising if many of the bots used in these attacks were compromised home computers.

Different security researchers provide different answers to why this problem persists. Software engineers proclaim the difficulty of writing bug-free code. Many computer scientists complain about “stupid users.” Some usable security researchers believe that the software is simply too complex to operate securely [11]. As social scientists, we have a different take on this problem. We don't believe that the majority of people are stupid, or computer illiterate. Rather, we believe that users are intentionally choosing actions that leave them insecure. This is not because they are being tricked by social engineering (though that sometimes happens), but rather because people honestly believe that they are doing what is necessary to protect their computers.

In a previous study, Rick Wash found that home computer users have a variety of different “mental models” of security threats [44]. Mental models describe how a user thinks about a problem; it is the model in the person's mind of how things work. People use these models to make decisions about the effects of various actions [22]. For example, some believed that hackers are mischievous teenagers showing off for their friends. Others believed that hackers are criminals out to steal financial and identity information. All of the respondents he interviewed were motivated to take positive security actions, but only for the threats they believed existed. Users who believed that hackers are teenaged troublemakers were likely to install firewall and other security software to keep them out, while the users who saw hackers as criminals frequently believed that they were not rich or important enough to be a target, and therefore didn't need to secure their computers.

Neither model is correct, though both are used to make decisions. These users were trying to do something to protect against the threats as they understood them. One of the most interesting results from this study was that even users who had wrong or incorrect mental models sometimes made good security choices. Even though most hackers today are not teenagers trying to impress friends, people who had that mental model worked hard to ensure their computer was protected.

This suggests a promising approach that may help home computer users better secure their systems: induce mental models of security threats that lead to good security behaviors. *Even if the mental models are wrong, they can still lead to good security behaviors and more secure computers.* We should help people develop an understanding of computer security that leads to good security behaviors, even if those understandings might not be technically “correct”. To change people’s mental models, we need to do two things: 1) Identify how people form these mental models, and how we can influence them; and 2) Identify which models are associated with what security behaviors, so we know which models we want home computer users to possess.

This approach does not require informing non-technical users about the details of computer security and how computers work, which most people are not interested in learning. Nor does it require making decisions for the users, which is a difficult solution to implement because it is technically challenging to come up with defaults that work for everyone, and frustrating for users who feel like they have little choice in the matter. Instead, we hypothesize that it is possible to empower home computer users to make their own choices, but in a way that leads to positive security for all.

High-level Research Agenda.

The interviews conducted by Wash [44] suggested that most people formed their mental models of security threats based on reasoning about information provided by stories recounted by their friends and colleagues. This process we described has much in common with existing research about “folk models” or “lay theories” that form through stories shared among people in a community, and through personal experience [26].

Like “folk models,” the models identified by Wash [44] are simple enough that many home computer users can understand them, but powerful enough to affect behavior. Therefore, we believe that the first step is to better understand *how people learn about and form mental models of security threats*, so that we can develop ways to influence these models and inspire better security behaviors. We suspect that most of the information comes from stories from friends and colleagues and other “people like me” [16]. This knowledge will also allow us to be more effective at teaching non-technical home-computer users about security by giving us a way to talk to them that they can understand and incorporate into their behavior.

One of the challenges to this approach is that most of the work in this area, including the previous interview study, provides only self-reported data about security behaviors. Unfortunately, people often do not accurately report their security behaviors, because they want to be seen as more security conscious than they actually are, or because they are mistaken or unaware. For example, the 2010 National Cyber Security Alliance Norton/Symantec Online Safety study

found that 58% of people answered “yes” to indicate that they had a full software security suite, but only 37% actually had one installed [35]. To make widespread progress, we need to *measure the connection between mental models and actual security behaviors*. Not all mental models lead to positive security behaviors; we want to identify the ones that do. And we want to understand how prevalent different models are, which may help us and other security researchers to understand why certain types of vulnerability are more common than others.

Finally, to actually make a difference in home computer security, we need to find a way to distribute this information—to actually help people form the mental models that lead to good security behaviors. Traditional training methods, such as having an expert teach a group of home computer users, will not work here, both because these methods are intractable and expensive, and because previous work suggests that mental models are best transmitted through stories from friends and other “people like me.” A cost-effective method of improving home security therefore is to get home computer users to train each other and spread the good mental models amongst themselves.

1.1 Inspiring Example: Home Thermostats

As an inspirational example, consider some work done by Willet Kempton in the mid 1980s [27]. Kempton is an anthropologist interested in energy conservation, and was studying how people make decisions about home heating. One of his papers focused on the question of thermostat setting: how do people set their thermostats in their homes to keep the house warm in the winter, and does this pattern significantly effect the amount of energy they use?

Kempton discovered that most people had one of two mental models of how thermostats work, and those models played a significant role in how these people made thermostat setting decisions. In the *valve* model, people believed the thermostat operated like a valve on a faucet; when you turn it higher, more heat comes out. People who operated based on this model showed fairly erratic thermostat settings, frequently turning the thermostat up to heat the house faster, then turning it way down once the house was warm. Kempton confirmed this with data from the energy company. Alternatively, people who operate with the *feedback* model believe that the thermostat turns the furnace on and off based on room temperature, but the furnace runs at a single constant output when on. People with this model frequently set it once and allow the thermostat to keep the temperature approximately the same throughout the day.

As Kempton puts it, “Heating engineers are fairly comfortable with the [feedback] theory described here; they consider it simplified but essentially correct.” The feedback theory is closer to the correct model of how thermostats work than the valve theory. However, upon interviewing users and looking at energy use data, Kempton concluded that the valve theory actually works better than the feedback theory. Kempton claimed that thermostat use consistent with the valve theory leads to more comfortable houses. But, more importantly, he found that it also leads to less overall energy use. He traced this to one important decision: should a user turn the thermostat down overnight?

Users with the feedback model believe that it takes more energy to raise the temperature of the house from 55° to 65° than simply maintaining a steady temperature of 65°. Users

operating with the valve theory correctly predict that more fuel is consumed at higher settings than at lower settings, and thus turning the thermostat down overnight saves energy. The valve model prediction is correct, even though the reasoning is wrong; hotter houses lose heat faster to the outside than colder houses. Thus, valve model users are more likely to turn down the thermostat overnight, and thus generally use less energy than feedback model users even though their mental model is less correct.

From this example we draw a number of interesting lessons. First, users make everyday decisions using simplified mental models. The people in Kempton’s study made multiple decisions every day based on simple mental models that do not represent a full understanding of home heating and energy use. Neither model includes an understanding of how the house interacts with the outside air, which is an important part of expert reasoning about home heating. However, both models serve the goals of the user in helping to guide their decision making in setting the thermostats in their homes.

Second, none of the users in Kempton’s study possess a complete and accurate model of home heating. No one he interviewed included interactions with outside air in their mental model. We suspect that very few people include these interactions in their mental model of thermostats even today. The more correct model is complicated, and difficult to use for decision making. Having to reason about a wide variety of factors (energy use, comfort, air movement, outside air interaction, etc.) to make a simple thermostat setting decision is not worth it, and people prefer to use simpler models that make the decisions easier and yield sufficiently good outcomes. In a word, most people *satisfice* [40] when choosing a mental model.

Third, just because a model is closer to the correct (or expert) model does not necessarily mean that it leads to better decisions. Valve model users make decisions that both lead to greater comfort and lower energy use. All models that lay people use are simplifications and will get some decisions *wrong*; it is important to figure out which models lead to better decisions, not which models are closer to correct. Kempton argues that “Technical experts will evaluate folk theory from this perspective [correctness] – not by asking whether it fulfills the needs of the folk. But it is the latter criterion [...] on which sound public policy must be based.” Experts often use correctness as a shortcut, assuming that more correct models lead to more correct decisions. But that isn’t necessarily true.

Finally, simplified mental models that non-experts use often lead to behavior that seems erratic. Valve model users frequently change the setting on their thermostat, micromanaging the heating of the house. But it is important to look at the big picture; that erratic micromanaging uses extra energy, but that energy is made up for by the (correct) lowering of the temperature overnight. It is extremely difficult to find a mental model that is both simple enough that people use it, and always leads to a correct decision. But sometimes changing to a different model can lead to a better overall outcome even though there exist individual decisions that are problematic.

Kempton’s study of home thermostats is inspiring, but it is important to recognize some of the important differences between home thermostat control and information security. Home thermostat users have the benefit of strong feedback; if they really mess up the setting of the home thermostat

they either freeze or roast in their house. This bounds how badly the mental model they use can perform; a model that regularly causes the person to be freezing is unlikely to be held for long. Information security has very little direct feedback. Even if a person holds a model that leads to very poor choices, they may never realize how their actions are causing security problems. Also, home thermostat feedback is fairly timely, and fairly directly connected to user actions. If a person sets their thermostat poorly, they usually find out in a matter of hours. On the other hand, a person who makes poor home computer security decisions may not find out that they have been phished for weeks, and then will have trouble associating that outcome with the specific decision that led to the security problem. This lack of timely feedback means that mental models for home computer security are more difficult, and more variable, than mental models of home thermostats.

2. RELATED WORK

2.1 The Problem: Securing Home Computers

According to the US Census Bureau, there are over 80 million households in the United States that have a computer with Internet access in their home; this represents almost 70% of all households in this country [10]. One of the biggest challenges these computer users face is operating these computers securely. There is a large ecosystem of threats, and a large percentage of them specifically target home computer users. Symantec, a computer security vendor, analyzed the security threats they addressed in 2006 and found that “from the 2249 new threats identified during the first 6 months of 2006, 86% were aimed at home users [42].” In the last ten years or so, this has become an even bigger problem through the rise of botnets that compromise thousands of computers—often home computers—and combine them into a network that can be used to attack third parties. The insecurity of a home computer no longer only affects the owner of that computer; it negatively affects all of the victims of the botnet also [38]. This is what economists call a negative externality [29] and it means that even rational, fully informed home users would choose to provide less security than society would prefer.

Managing the security of a computer system is very difficult, and the vast majority of home computers are administered by people who have little security knowledge or training. The main reasons that novice users cite for improper security relate to a lack of knowledge and understanding: “43% claimed not to understand the threats, 38% claimed they did not know how to use security packages, 35% indicated that they did not know how to secure their computer, and 32% indicated that they did not know about the threats” [15]. Similarly, Dourish et al.[14] found that corporate knowledge workers saw security technologies as a “barrier” (like a locked door) that keeps the bad guys out. Security technologies were expected to keep out all potential bad guys, and technologies that focused on one specific type of bad (like anti-spam technologies without anti-virus capabilities) were seen as “partial” or imperfect. Users also felt futility with security, referencing unknown others (like hackers) who will “always be one step ahead.” Most users, lacking the time and inclination to deal with security, attempted to delegate their security concerns.

Grinter et al. [17] interviewed home network users and

found that homes generally had a single person who assumed the role of “system administrator.” It was his or her job to maintain a network, troubleshoot and fix problems, and help others with network connectivity, and this person “unanimously resented the amount of time” spent in this role. Gross and Rosson [18] studied what security knowledge end users, who were not directly responsible for security but had access to sensitive information, possessed in the context of large organizations. Users’ security knowledge was “neither comprehensive nor sufficient” to maintain proper security, but common security actions such as locking the screen when away were better understood and practiced. All participants were aware of some sensitive information they had access to, and knew to protect it and to be wary of being tricked into revealing it (social engineering). Gross and Rosson also noted that their participants frequently conflated security and functionality failures.

Users rely on others for security because they feel like they don’t have the skills to maintain proper security themselves, so they often try to avoid security decisions. They find ways to delegate the responsibility for security to some external entity which could be technological (like a firewall), social (another person or IT staff), or institutional (like a bank). However, despite this delegation of responsibility, many users still make numerous security-related decisions on a regular basis. The literature does not explain how those decisions get made; rather, it focuses mostly on the anxiety the decisions create.

2.2 Three Common Approaches

Because home computer security is such a major problem today, many people have put forth a wide variety of solutions to this problem and its various sub-problems. We believe that these solutions can be classified into three main categories. First are technical solutions that attempt to take the decisions out of the hands of the user. Second are education approaches that attempt to teach home computer users the details of computer security. And third, a group of academics are now focusing on trying to understand why these users behave the way they do, and how we can support them in their efforts.

2.2.1 The “Stupid User” Approach

The majority of solutions to home computer security involve technologies that try to make it easier to act securely by removing the user from the decision-making process. The argument is usually that unsophisticated users do not have the knowledge and skills to make good security choices, and therefore we should remove the decisions from them for their own good [11]. Or, at the very least, experts should choose secure defaults and make it difficult for users to act insecurely. Many of these technologies have been quite successful. For example, modern anti-virus software has very little interactivity with users; it regularly scans computers for known malware and removes it automatically. Microsoft has automatic updates turned on by default in all recent versions of its Windows operating system; once a month the system automatically downloads and installs all security patches from Microsoft without asking for user permission [34]. Many modern firewalls do not require user configuration to operate securely.

Unfortunately there are limits to the effectiveness of this approach. A number of modern threats are particularly diffi-

cult to defend against with technical solutions. Attackers are increasingly relying on social engineering attacks [3]. Malware must first trick an end user into activating it. Once activated, it must gain access to the user’s computer or information either through some vulnerability, or by authenticating as if it were the user (i.e. through stealing a password). Once it is “in” it does what it was designed to do, and often propagates itself so it can infect others. In addition, there are often legitimate reasons why a user would want to operate insecurely; for example, a system update may break compatibility with some favorite software, and therefore the user may not want to install that update. Or a user may want to visit websites with invalid or unsigned SSL certificates that they know are legitimate. The “stupid human” approach requires a one-size-fits-all solution to security problems, but people use computers for such a variety of different purposes that rarely does one solution work for everyone. When it works, it is often the best solution; but it is limited in what security problems it can adequately address.

2.2.2 The Education Approach

A second approach can be seen as the polar opposite: allow users the freedom to choose, and provide them with appropriate training so they can make good security choices. There has been much effort devoted to training users in organizations to be more secure, and some researchers have investigated the effects of different kinds of training programs and security policies on security outcomes [2, 13]. Microsoft has an extensive online resource for teaching users all about computer security [33]. Many organizations including Microsoft, CERT and US-CERT include lists of advice for being more secure. This approach has many parallels in other domains; for example, many organizations are also working hard to educate consumers about environmentally-friendly activities and goods. However, these approaches are most effective when the desired behavior changes are not very difficult or costly to adopt, and often produce only modest, short-term improvements [41]. Home computer users are rarely interested in learning the details of how security software works in order to make appropriate security decisions. Often, the details are so complicated that users get frustrated and don’t really understand enough.

2.2.3 The “Understand How Users Think” Approach

A more recent approach involves working to understand how computer users think about security, and how they actually make security decisions. Cormac Herley [20] argued that when non-expert users reject security advice, it is often rational to do so. He wrote that security experts provide advice that ignores the costs of the users’ time and effort, and therefore overestimates the net value of security. We agree, though we want dig deeper into understanding how users actually make these security / effort tradeoffs.

Jean Camp [8] proposed using mental models as a framework for communicating complex security risks to the general populace. She created five possible analogies or metaphors for computer security: physical security, medical risks, crime, warfare, and markets. Asghapour et al. [4] built on this by conducting a card sorting experiment in which participants were instructed to match these analogies with a set of computer security related concepts. They found that experts and non-experts show sharp differences in which analogy they felt the concepts were closest to, and hypothesized that

the analogies might function as mental models. However, they did not test this hypothesis. Also, to our knowledge no prior work has examined how mental models of security are formed, or how we can influence them.

The Management Information Systems (MIS) literature has several examples of research projects that approach computer security behaviors from an “adoption” perspective; these researchers suggest that people adopt security behaviors much in the same way as they might adopt a new technology. They seek to understand psychological characteristics that lead people to adopt security behaviors, and how these characteristics interact with the messages about security that home computer users might receive [47, 3, 28, 48].

The studies in this section all stem from a common goal: to understand behavior, one must understand how people think. However, there is very little previous work that takes this chain of logic one step further, to consider what shapes how people think about security. Where does the information come from that users base their decisions on? In an enterprise context, researchers talk about security policy and education campaigns; some studies try to look at how this carries over into home computer security. However, it is much harder to collect data the different kinds of information that home computer users have access to.

2.3 Mental Models and Computer Security

A *mental model* is a “simplified representation of reality that allows people to interact with the world” [25]. Mental models describe how a person reasons and makes inferences about a problem or situation, allow people to make predictions about what might happen, and provide heuristics and guidelines upon which to base behavioral choices. They are not exact replicas of the world down to every last detail; they are representations based on reality as a person experiences it, and used to help people make choices about how to behave. Mental models are functional, rather than complete or accurate. This means that people need and use them all the time; however, they are abstractions that contain inaccuracies when compared with the real world.

We sometimes refer to mental models—the terminology cognitive psychologists use—as *folk models* which is terminology borrowed from anthropology. The “folk” in folk models are the average normal “folks”, people who are not experts, and who have not been formally trained in a particular area. Folk models are “ways of understanding” [26] the world that arise informally, not from explicit efforts to understand something via formal instruction or education; these are the kind of models we refer to in this paper.

Mental models change and develop over time, adapting to new information and new experiences [25]. This is an unconscious process, meaning that you don’t have to think explicitly about your mental model of something to modify it. However, if you haven’t thought about a particular circumstance or situation, or experienced it directly, then it can’t be a part of your mental model.

Mental models have a “chain of causation” [12], a process aspect that helps us figure out how stuff works and what to do next. People are born knowing how to reason about complex situations involving cause and effect relationships, and this skill supports the formation of mental models. People are also able to go beyond the information they receive and generate new hypotheses/explanations based on the information in their mental models [23]. While mental models

represent causal relationships, they are not necessarily procedural [25]. This means that while they help us reason about cause-and-effect, they are not always good at allowing us to see effects that might be multiple steps removed from the immediate ones.

An example of a mental model of home computer security from Wash [44] is what he called the “Big Fish” model: “Hackers are criminals who target ‘big fish’; I’m ‘small potatoes’, so nobody is going to target me.” This mental model leads home computer users to take few steps to protect themselves—they don’t believe they are at risk, so it is not a high priority to protect themselves. This model is causal at a high level, and includes the idea that if one becomes a “big fish” one becomes a target. It is not, however, procedural: the steps it would take to become a “big fish” worthy of the attention of hackers are underspecified, as well as how hackers make choices about who is a good target. The model can be used to guide security choices, namely “I don’t need to do much to protect myself”, but at the same time it is missing information that if it were part of the model might lead to different choices.

Mental models can end up being “incorrect” in a number of ways. Incoming information from past experience or from things one has read or heard about is filtered and organized according to existing mental models [25], and this can prevent new experiences that contradict existing mental models from being incorporated. For example, people whose mental models of computer security include “hackers go after big fish” will have a harder time incorporating the idea that “sometimes people who aren’t big fish also get hacked” into their mental model. In addition, coincidences can reinforce existing mental models, and people can apply mental models to situations just because they are comfortable and familiar with that particular way of interacting with the world [46].

Mental models also represent what is true at the expense of what is false. This allows us to internalize commonalities among the information we encounter rather than having to remember all of the exceptions, which makes our models easier and more efficient to use, but also leads to systematic deduction errors and misconceptions [23, 26]. For example, if one believes “hackers go after big fish”, the idea that “the big fish I know haven’t been hacked” is not likely to be part of the model in the model.

2.4 Stories, Mental Models and Behavior

We want to understand how and why laypeople—non-technical home computer users—choose the security behaviors they do. We believe that people form mental models of threats to security based on information they receive via stories from other people like themselves, from the media, from communications and interactions with experts, and from their own experiences.

Home computer users often do not have a lot of examples around them of security problems that they can use to learn how to react, or security experts to learn from. So how do people learn how to behave in situations they don’t have direct experience with? We learn from other people. Social information sharing is an important way that we learn about the world around us and how to behave in it. Narratives—stories told by other people—are an important component of our ability to learn about the world around us and behave appropriately [7]. Stories people tell about each other, sometimes labeled gossip (a word that has unfortunately negative

connotations), constitute observational learning, and help us avoid others' mistakes [5].

Stories about others reveal useful information about how our culture and society operate; it is easier to make our way through our complex world if we can learn from the experiences of others, and stories are a vehicle for this information [5]. Stories that affirm what is already represented in our mental models are remembered more easily, are given more weight, and are more likely to be passed on [30, 19]. In an exploratory study of gossip, Baumeister [5] found that most instances of gossip people could remember hearing were about people personally known to the participant (85%). People find gossip most interesting when it is about people similar to them [31], and model their behavior after people they perceive to be similar to them [16]. This is evidence that people do indeed learn about how to behave from the experiences of similar others. Stories that arouse emotion, such as stories about bad things that happened to acquaintances, are more likely to be remembered and passed on [37, 19].

Mental models are helpful for making inferences and guiding behavioral choices. For example, Uther and Haley [43] taught a group of participants what they deemed the "correct" model for the web browser 'back' button, and compared participants' performance with a group of people who had not received any training. The group who had received the training had significantly fewer page traverses when navigating a website, indicating that their mental models had been altered by the training. Schobel and Manzey [39] wrote about how engineers in a nuclear power plant use their mental models of the socio-technical system comprised of the technical staff and the computing infrastructure to make predictions about what might go wrong in a crisis, and imagine how they would react. Bostrom [6] conducted three studies eliciting mental models from groups of people about the risk associated with lead paint, smallpox vaccine, or climate change, and then asked them to think through hypothetical scenarios that required them to reason and make inferences about the topic. They found that participants responses to the scenarios differed, depending on which analogies and metaphors were present as part of their mental models.

Medin [32] conducted a study of expert fishermen in the Northwoods of Wisconsin, in which they elicited and compared the mental models of both Native American fishermen and of majority-culture fishermen. Despite both groups being experts, the two groups showed dramatic differences in the way fish were categorized and classified. Majority-culture fishermen grouped fish into standard taxonomic and goal-oriented groupings, while Native American fishermen groups fish mostly by ecological niche. This illustrates how even experts can have dramatically different mental models of the same phenomenon, and any single expert's model is not necessarily correct.

2.5 Folk Models of Home Computer Security

In prior work, Rick Wash interviewed 33 non-expert home computer users in 3 major midwestern cities in order to better understand the mental models home computer users possess and use to make everyday security decisions. He focused on differences between people, and characterized different methods of dealing with security issues rather than trying to find general patterns. The mental models he described may explain differences observed between users in these studies [44].

Wash identified eight different mental models in his data, which he divided into two broad categories based on a distinction that most subjects possessed: 1) models about viruses, spyware, adware, and other forms of malware, which everyone referred to under the umbrella term 'virus'; and 2) models about the attackers, referred to as 'hackers,' and the threat of 'breaking in to' a computer. Each respondent had at least one model from each of the two categories. For example, one respondent believed that viruses were mischievous, and hackers are criminals who target big fish. The models were not necessarily mutually exclusive, and every mental model was shared by multiple respondents. Most respondents made a distinction between 'viruses' and 'hackers.' To them, these are two separate threats that can both cause problems. Some people believed that viruses are created by hackers, but they still usually saw them as distinct threats.

Models of Viruses and other Malware.

All of the respondents used the term 'virus' as a catch-all term for malicious software. Everyone seemed to recognize that viruses are computer programs. Almost all of the respondents classify many different types of malicious software under this term: computer viruses, worms, trojans, adware, spyware, and keyloggers were all mentioned as 'viruses.' The respondents don't make the distinctions that most experts do; they just call any malicious computer program a 'virus.' Wash [44] found four distinct mental models of 'viruses.' One model was very under-specified, labeling viruses as simply 'bad' and expecting that they cause generically bad things to happen. Respondents with this model had trouble using it to make any kind of security-related decisions because the model didn't contain enough information to provide guidance. Two other models (the *Mischief* and *Crime* models) were fairly well-described, including how viruses are created and why, and what the major effects of viruses are. In the *Mischief* model, viruses are created for mischievous purposes and cause visible damage; respondents felt like they would know if they were infected, and waited until they were infected to take action. In the *Crime* model, respondents felt that viruses were created by criminals to steal information, and took many precautions to ensure that they didn't get a virus.

Models of Hackers and Break-ins.

The second major category of mental models describe the attackers, or the people who cause Internet security problems. These attackers are always given the name "hackers," and all of the respondents seemed to have some concept of who these people were and what they did. The term "hacker" was applied to describe anyone who does bad things on the Internet, no matter who they are or how they work. All respondents described the main threat that hackers pose as "breaking in" to their computer. They disagreed about why a hacker would want to "break in" to a computer, and which computers they would target, but all agreed on the terminology for this basic action. To the respondents, breaking in to a computer meant that the hacker could then use the computer as if they were sitting in front of it, and could cause a number of different things to happen to the computer. Many respondents stated that they did not understand how this worked, but they still believed it was possible. Wash [44] described four distinct mental models of hackers. Some

people see hackers as a digital equivalent of graffiti artists, causing mischief but otherwise little harm. Most of the respondents saw hackers as a form of criminal, but differed widely on what type of criminal. One group saw hackers as a form of street thief, opportunistically attacking any computer they could find. These respondents worked hard to avoid places (websites) where hackers might be found, and to protect their computers. Another group saw hackers as a form of cat burglar, carefully choosing which computers to break in to. None of these respondents felt that they were a target, and consequently put forth little effort in securing their computer. Finally, the last group saw hackers as a low-level worker for organized crime, methodically stealing as much identity information as possible.

3. PROPOSED RESEARCH AGENDA

We believe that one of the foundations of information security should be a better understanding of how users make security decisions, and an accurate characterization of the problems that result from these decisions. We believe that a multi-disciplinary approach will help home computer users choose more desirable security behaviors. The main goal of this agenda is to understand how home computer users form the mental models they use to make security decisions, use that knowledge to influence people’s mental models, and hopefully affect real-world security behaviors as a result.

Following the advice of Kempton [27], we do not evaluate models on the basis of correctness in terms of the reality of security threats home computer users face, but rather on the potential benefits that could be attained from behaviors that stem from the models regardless of correctness. In other words, we believe that to make progress in this area, we should not constrain ourselves to teaching “correct” models and ideas about security; rather, we should focus on communicating models that are simple, coherent, easy to understand, and lead to desirable behaviors.

Once we discover where users’ mental models of security come from and have information on how prevalent they are, we as a community will be in a better position to develop interventions that will influence users’ mental models and the behaviors that inevitably follow. Unlike one-time incentives, changing a mental model can change behavior long after the impetus for that change has passed.

3.1 Where do Mental Models Come From?

In previous work, Wash showed that users have distinct mental models of security threats, and that these models affect home computer users’ security-related choices and behaviors [44]. However, he did not directly investigate the potential sources of the information incorporated in users’ mental models, or the impact of different kinds of information on users’ mental models. Anecdotally, those respondents seemed much better able to recount stories about security problems that had happened to them or to someone in their social circle, than they were to provide information from any other sources [44]. This leads us to hypothesize that information learned from and about other people is an important component of home computer users’ mental models of security. This is something that has not been rigorously addressed in the computer security literature; however, researchers developing interventions intended to promote energy conservation are beginning to develop similar approaches intended to inform people about the behavior of

others and thereby promote behavior change in the target participants [9, 36].

It is important to remember that mental models held by laypeople are not necessarily accurate representations of the real world [12, 26]. In fact, it is well-known that in technological contexts users often operate with incorrect mental models [1] — even security experts disagree about the correct way to think about viruses or hackers. To understand the rationale for home users’ behavior, it is important to understand the models upon which their behavior is based, correct or incorrect.

We suspect that a number of factors affect the incorporation of information into users’ mental models. These factors may be important for home computer users’ decision to incorporate new information into their mental models:

- Content: information about threats vs. behaviors [2, 24]
- Source: friends, other people like me, experts, or unfamiliar strangers [5, 16]
- Format: stories vs. policies vs. explicit advice [44]
- Valence: positive vs. negative cases [5]
- Intensity: emotional vs. not [37]
- Style: figurative vs. literal language [21]

In particular, we suspect that: a) Stories about threats have a larger influence on mental models than behavioral advice; and b) Information from friends and colleagues have a stronger influence on mental models than information from strangers or experts. These hypotheses are important because, if true, they suggest that the standard way of educating users — experts providing behavioral advice — is a poor way to teach security.

3.2 Prevalence, Antecedents, and Consequences of Mental Models

An important limitation of the study by Wash [44] is that he was not able to measure the prevalence of each mental model. The qualitative interviews helped understand many different ways of thinking about computer security, but he was not able to answer the question “How common are these models?” This is an important question, because people use mental models to reason about complex situations and decide how to behave. For example, if a common mental model leads people to believe that most security tools and advice are useful, then computer security is mainly a usability issue—we only need to make existing security technologies easier to use. On the other hand, if a large number of people have mental models that lead them to ignore security tools and advice because they think they are not a target of hackers, then better tools cannot solve the problem.

In addition to looking at the prevalence of each mental model, it would be valuable to look at antecedents for these mental models. Are people with more experience with security breaches more likely to have a more realistic model? Or a model that leads to secure behaviors? Are office workers who have received information security training more likely to have accurate models? Looking at common antecedents of each mental model can help to identify existing behaviors and properties that are likely to lead to good security behaviors or bad security behaviors.

Finally, we believe that we need more data about which security behaviors are commonly associated with each mental model. Wash [44] listed a number of behaviors that occurred with each model he found, but those behaviors are

not necessarily common across everyone with that model. By explicitly connecting which behaviors commonly appear with each model, we will better understand what models are the ‘good’ models that we should encourage, and which models are the ‘bad’ models that should be discouraged.

One of the main challenges in connecting models with behavior is measuring those behaviors. Self-report can be misleading, especially for security data. Many people feel social pressure to seem more secure than they really are, and therefore may report more security behaviors than they actually do. Additionally, there are many security-related behaviors that people simply don’t think about, or don’t recognize as security behaviors, and therefore cannot report accurately.

3.3 Sharing Security Stories

In the end, our goal is not just to understand user security behavior, but to improve security. We believe that a promising approach to improving security is to change people’s mental models to a new model that, while not entirely correct, leads to better security behaviors. Our research leads us to believe that one promising way to change people’s mental models of security is to provide a mechanism for regular, non-technical users to share stories with each other. Since, as we hypothesized above, stories from regular people are more powerful than advice from experts, we suspect that sharing stories will actually lead to important mental model changes.

However, simply sharing stories isn’t enough; people need to share the *right* stories. This may be our biggest challenge for this research agenda, but have reason to believe we can overcome it. Ideally, people would only share stories that lead to “better” mental models — where better means mental models that cause people to take appropriate security precautions. This is an opportunity for experts to be involved. If these stories are curated by experts, and stories that lead to positive security behaviors are highlighted, then it may be possible to cause aggregate change in user behavior.

Systems for sharing stories are not without precedent; in the medical world there is a fascinating website called *Patients Like Me* (<http://www.patientslikeme.com>). Patients Like Me is an online information resource where people who have been diagnosed with specific long-term medical diseases such as Parkinson’s disease, Fibromyalgia, and Depression can go and share stories about their disease, and share information about treatments. Users can self-identify as having a disease, write reports about individual symptoms of that disease, and write experience reports about treatments. That information is then aggregated and shared; a user can, for example, search for stories from people who have fatigue due to Parkinson’s disease, or stories about using Wellbutrin to treat depression. Patients Like Me is a system that allows normal people to share stories and experiences with various diseases, and make suggestions to each other for treatment; however, it is still curated by doctors and many doctors participate on the site to help direct people toward positive ideas.

Potential Risks.

One potential risk is that mental models that lead to positive security behaviors might not exist; however, we believe this is unlikely. Wash found a number of incorrect models that respondents reported led to positive security behav-

iors [45]. Additionally, mental models and “lay theories” are surprisingly good at helping us cope with a complex world; people seem good at coming up with models that lead to effective decisions despite not understanding the details of, for example, how a car works or how a web browser works [26].

Another risk is that users of the site might not be willing to share sensitive information about security incidents, or won’t see such stories as credible. We think this is also unlikely; home computer users are already sharing information and stories with each other in the real world. Additionally, users on *Patients Like Me* and other social media sites are willing to share highly personal, highly sensitive information with the other users on the site, and remain credible in doing so. We believe that home computer users will likewise be willing to share stories and experiences about security issues from themselves and others.

4. SUMMARY

People frequently make decisions by consulting a simplified *mental model* of the way the world works. When making decisions about the security of a home computer, people frequently use mental models of security threats to balance security concerns with the time, effort, and money involved in being secure. We propose a new way of thinking about end-users and security: rather than trying to teach people “correct” mental models, we accept the fact that mental models are always simplified and incomplete. Instead, we focus on finding ways to encourage models that lead to valuable security behaviors even if they are “incorrect.” And rather than trying to force mental models on people, we should take advantage of the work done by end-users in forming usable mental models, and encourage those existing mental models to spread.

Acknowledgements

We thank Jeff MacKie-Mason for his continued advice and thoughts in this matter. We also thank the SOUPS and SHB communities for feedback. This material is based upon work supported by the National Science Foundation under grants CNS 1116544 and CNS 1115926.

References

- [1] ADAMS, A., AND SASSE, M. Users are not the enemy. *Communications of the ACM* 42, 12 (1999), 46.
- [2] ALBRECHTSEN, E., AND HOVDEN, J. Improving information security awareness and behaviour through dialogue, participation and collective reflection. An intervention study. *Computers & Security* 29, 4 (June 2010), 432–445.
- [3] ANDERSON, C. Practicing Safe Computing: A Multimethod Empirical Examination of Home Computer User Security Behavioral Intentions. *Citeseer* 34, 3 (2010), 613–643.
- [4] ASGHARPOUR, F., LIU, D., AND CAMP, L. Mental models of computer security risks. In *Workshop on the Economics of Information Security (WEIS)* (2007), Citeseer.

- [5] BAUMEISTER, R. F., ZHANG, L., AND VOHS, K. D. Gossip as Cultural Learning. *Review of General Psychology* 8, 2 (2004), 111–121.
- [6] BOSTROM, A. Lead is like mercury: risk comparisons, analogies and mental models. *Journal of Risk Research* 11, 1 (Jan. 2008), 99–117.
- [7] BRUNER, J. The narrative construction of reality. *Critical inquiry* 18, 1 (1991), 1–21.
- [8] CAMP, L. Mental models of privacy and security. *Technology and Society Magazine, IEEE* 28, 3 (2009), 37–46.
- [9] CARRICO, A. R., AND RIEMER, M. Motivating energy conservation in the workplace: An evaluation of the use of group-level feedback and peer education. *Journal of Environmental Psychology* (Nov. 2010).
- [10] CENSUS, U. Current Population Survey, Computer Use and Ownership supplement, 2009.
- [11] CRANOR, L. F. A Framework for Reasoning About the Human in the Loop. In *Proceedings of the 1st Conference on Usability, Psychology, and Security (UPSec)* (2008).
- [12] D’ANDRADE, R. *The Development of Cognitive Anthropology*. Cambridge University Press, 2005.
- [13] DOHERTY, N. F., ANASTASAKIS, L., AND FULFORD, H. The information security policy unpacked: A critical study of the content of university policies. *International Journal of Information Management* 29, 6 (Dec. 2009), 449–457.
- [14] DOURISH, P., GRINTER, R. E., DELGADO DE LA FLOR, J., AND JOSEPH, M. Security in the wild: user strategies for managing security as an everyday, practical problem. *Personal and Ubiquitous Computing* 8, 6 (Sept. 2004), 391–401.
- [15] FURNELL, S., BRYANT, P., AND PHIPPEN, A. Assessing the security perceptions of personal Internet users. *Computers & Security* 26, 5 (Aug. 2007), 410–417.
- [16] GOLDSTEIN, N. J., CIALDINI, R. B., AND GRISKEVICIUS, V. A Room with a Viewpoint: Using Social Norms to Motivate Environmental Conservation in Hotels. *Journal of Consumer Research* 35, 3 (2008), 472–482.
- [17] GRINTER, R., EDWARDS, W., NEWMAN, M., AND DUCHENEAUT, N. The Work to Make a Home Network Work. In *ECSCW 2005* (2005), Springer, pp. 469–488.
- [18] GROSS, J., AND ROSSON, M. Looking for Trouble: Understanding End-User Security Management. In *Proceedings of the 2007 Symposium on Computer Human Interaction For the Management of information Technology* (2007), pp. 30–31.
- [19] HEATH, C., BELL, C., AND STEINBERG, E. Emotional Selection in Memes: The Case of Urban Legends. *Journal of Personality* 81, 6 (2001), 1028–1041.
- [20] HERLEY, C. So Long, And No Thanks for the Externalities: The Rational Rejection of Security Advice by Users. In *NSPW ’09 Proceedings of the 2009 workshop on New security paradigms workshop* (2009).
- [21] HSU, Y. The effects of metaphors on novice and expert learners’ performance and mental-model development. *Interacting with Computers* 18, 4 (July 2006), 770–792.
- [22] JOHNSON-LAIRD, P., GIROTTO, V., AND LEGRENZI, P. Mental Models: A Gentle Guide for Outsiders, 1998.
- [23] JOHNSON-LAIRD, P. N. Inaugural Article: Mental models and human reasoning. *Proceedings of the National Academy of Sciences* 2010 (Oct. 2010).
- [24] JOHNSTON, A. Fear Appeals and Information Security Behaviors: An Empirical Study. *MIS Quarterly* 34, 3 (2010), 549–566.
- [25] JONES, N. A., ROSS, H., LYNAM, T., PEREZ, P., AND LEITCH, A. Mental Models: An Interdisciplinary Synthesis of Theory and Methods. *Ecology And Society* 16, 1 (2011).
- [26] KEIL, F. C. The Feasibility of Folk Science. *Cognitive science* 34, 5 (May 2010), 826–862.
- [27] KEMPTON, W. Two Theories of Home Heat Control. *Cognitive Science* 10, 1 (1986), 75–90.
- [28] KUMAR, N., MOHAN, K., AND HOLOWCZAK, R. Locking the door but leaving the computer vulnerable: Factors inhibiting home users’ adoption of software firewalls. *Decision Support Systems* 46, 1 (Dec. 2008), 254–264.
- [29] MAS-COLELL, A., WHINSTON, M. D., AND GREEN, J. R. *Microeconomic Theory*. Oxford University Press, 1995, ch. 11 Externalities and Public Goods, pp. 350–382.
- [30] MAZZOCCO, P., GREEN, M., AND BROCK, T. The Effects of a Prior Story-Bank on the Processing of a Related Narrative. *Media Psychology* 10, 1 (2007), 64–90.
- [31] MCANDREW, F. T., AND MILENKOVIC, M. A. Of Tabloids and Family Secrets: The Evolutionary Psychology of Gossip. *Journal of Applied Social Psychology* 32, 5 (May 2002), 1064–1082.
- [32] MEDIN, D. L., ROSS, N. O., ATRAN, S., COX, D., COLEY, J., PROFFITT, J. B., AND BLOK, S. Folkbiology of freshwater fish. *Cognition* 99 (2006), 237–273.
- [33] MICROSOFT. The latest in computer security, 2010. <http://www.microsoft.com/security/default.aspx>.
- [34] MICROSOFT. Security bulletin advance notification, 2010. <http://www.microsoft.com/technet/security/Bulletin/advance.msp>.

- [35] NCSA-Norton-Symantec Online Safety Study, 2010. <http://www.staysafeonline.org/blog/ncsanorton-symantec-online-safety-study-released-today>.
- [36] NIXON, H., AND SAPHORES, J.-D. Information and the decision to recycle: results from a survey of US households. *Journal of Environmental Planning and Management* 52, 2 (Mar. 2009), 257–277.
- [37] PETERS, K., KASHIMA, Y., AND CLARK, A. Talking about others: Emotionality and the dissemination of social information. *European Journal of Social Psychology* 39, 2 (2009), 207–222.
- [38] ROWE, B., AND WOOD, D. How the Public Views Strategies Designed to Reduce the Threat of Botnets. *Trust and Trustworthy Computing* 6101, May 2009 (2010), 337–351.
- [39] SCHÖBEL, M., AND MANZEY, D. Subjective theories of organizing and learning from events. *Safety Science* 49, 1 (Jan. 2011), 47–54.
- [40] SIMON, H. A. Rational choice and the structure of the environment. *Psychological Review* 63, 2 (1956), 129–138.
- [41] STEG, L. Promoting household energy conservation. *Energy Policy* 36, 12 (Dec. 2008), 4449–4453.
- [42] Symantec Internet Security Threat Report: Trends for January 06–June 06, 2006. http://www.symantec.com/specprog/threatreport/ent-whitepaper_symantec_internet_security_threat_report_x_09_2006.en-us.pdf.
- [43] UTHER, M., AND HALEY, H. Back vs. stack: training the correct mental model affects web browsing. *Behaviour & Information Technology* 27, 3 (May 2008), 211–218.
- [44] WASH, R. Folk models of home computer security. In *Proceedings of the Sixth Symposium on Usable Privacy and Security* (2010), ACM, pp. 1–16.
- [45] WASH, R., AND LAMPE, C. The Power of the Ask in Social Media. Tech. rep., Michigan State University, 2010.
- [46] WESTBROOK, L. Mental models: a theoretical overview and preliminary study. *Journal of Information Science* 32, 6 (2006), 563.
- [47] WORKMAN, M., BOMMER, W., AND STRAUB, D. Security lapses and the omission of information security measures: A threat control model and empirical test. *Computers in Human Behavior* 24, 6 (Sept. 2008), 2799–2816.
- [48] ZURKO, M. E. User-centered security: Stepping up to the grand challenge. In *Proceedings of Annual Computer Security Applications Conference (ACSAC)* (December 2005).