

BY RYAN WEST

THE PSYCHOLOGY OF SECURITY

*Why do good users make
bad decisions?*

*"... [the system] must be easy to use and
must neither require stress of mind nor
the knowledge of a long series of rules..."*

AUGUSTE KERCKHOFFS ON THE
DESIGN OF CRYPTOGRAPHIC SYSTEMS
(*La cryptographie militaire*, 1883)

The importance of the user in the success of security mechanisms has been recognized since Auguste Kerckhoffs published his treatise on military cryptography, *La cryptographie militaire*, over a century ago. In the last decade, there has been tremendous increase in awareness and research in user interaction with security mechanisms.

Risk and uncertainty are extremely difficult concepts for people to evaluate. For designers of security systems, it is important to understand how users evaluate and make decisions regarding security. The most elegant and intuitively designed interface does not improve security if users ignore warnings, choose poor settings, or unintentionally subvert corporate policies. The user problem in security systems is not just about user interfaces or system

ILLUSTRATIONS BY SERGE BLOCH





People tend to believe they are less vulnerable to risks than others. People also believe they are less likely to be harmed by consumer products compared to others. It stands to reason that any computer user has the preset belief that they are at less risk of a computer vulnerability than others.

interaction. Fundamentally, it is about how people think of risk that guides their behavior. There are basic principles of human behavior that govern how users think about security in everyday situations and shed light on why they undermine security by accident.

This article offers a brief introduction to research on risk, uncertainty, and human decision making and how they relate to users making security decisions, and provides a few key concepts and possibilities in how they may be used to improve users' security behavior.

Non-acceptance of security tools is recognized as a major problem facing the information security world [5]. Research in the usability of security mechanisms has exploded over the last decade and an excellent trove of research papers is cataloged by the *HCI Sec Bibliography* hosted at www.gaudior.net/alma/biblio.html. Among the studies listed there is a mountain of evidence that mechanisms for encryption, authorization, and authentication can be difficult for people to understand or use [1, 9] and that people often fail to recognize security risks or the information provided to cue them [3, 4]. Accordingly, researchers have promoted the need for user-centered design throughout the development process and warn that usability testing security systems only at the end of the process does not guarantee a usable or acceptable system [7, 11, 12].

However, there is more to this than interaction with technology. Human decision making has been a topic of study in social sciences from economics to psychology for over a century. The net sum of that research suggests that individuals are often less than optimal decision makers when it comes to reasoning

about risk. However, we have predictable and exploitable characteristics in our decision-making process. Understanding these principles and how users come to make decisions about security may suggest places where we can improve the outcome of the decisions.

Users do not think they are at risk. First of all, people tend to believe they are less vulnerable to risks than others. Most people believe they are better than average drivers and that they will live beyond average life expectancy [6]. People also believe they are less likely to be harmed by consumer products compared to others. It stands to reason that any computer user has the preset belief that they are at less risk of a computer vulnerability than others. It should come as no surprise that, in 2004, a survey from AOL and the National Cyber Security Alliance reported that roughly 72% of home users did not have a properly configured firewall and that only one-third had antivirus virus signatures updated within the past week.¹

Even as security measures improve, users will remain at risk. There is evidence that individuals maintain an acceptable degree of risk that is self-leveling, known as *risk homeostasis*.² Applied to security, it suggests that as users increase their security measures, they are likely to increase risky behavior. For example, the user who has just installed a personal firewall may be more likely to leave his machine online all the time.

Users aren't stupid, they're unmotivated. In social

¹America Online and the National Cyber Security Alliance. AOL/NCSA Online Safety Study, 2004; www.staysafeonline.info/news/safety_study_v04.pdf.

²G.J.S. Wilde. *Target Risk 2: A New Psychology of Safety and Health*. PDE Publications, Toronto, Ontario, 2001.

From Windows Explorer: UI #1

1. Right click on folder in public share (invokes UI #2)
2. Click on Properties in context menu (invokes UI #3)
3. Click on Sharing tab (invokes UI #4)
4. Click Share... (invokes UI #5)
5. Enter the User or Group name to share with
6. Click Add (automatically sets permission level to "Reader" which sets ACEs for Read, Read & Execute, and List Folder Contents)
7. Click Share (invokes UI #6)
8. Click Done (returns to UI #3)
9. Click Close (returns to UI #1)

cognition, the term is *cognitive miser*. Humans have a limited capacity for information processing and routinely multitask. As a result, few tasks or decisions receive our full attention at any given time. To conserve mental resources, we generally tend to favor quick decisions based on learned rules and heuristics. While this type of decision making is not perfect, it is highly efficient. It is efficient in the sense it is quick, it minimizes effort, and the outcome is good enough most of the time. This partially accounts for why users do not reliably read all the text relevant in a display or consider all the consequences of their actions.

Safety is an abstract concept. When evaluating alternatives in making a decision, outcomes that are abstract in nature tend to be less persuasive than outcomes that are concrete [2]. This is key to understanding how users perceive security and make decisions. Often the pro-security choice has no visible outcome and there is no visible threat. The reward for being more secure is that nothing bad happens. Safety in this situation is an abstract concept. This, by its nature, is difficult for people to evaluate as a gain when mentally comparing cost, benefits, and risks.

Compare the abstract reward (safety) garnered from being more secure against a concrete reward like viewing an attachment in instant messaging or Web content that requires a browser add-on and the outcome does not favor security. This is especially true when a user does not know what his or her level of risk is or believes they are at less risk than others to start. Returning to the principle of the cognitive miser, the user is also more likely to make a quick decision without considering all of the risks, consequences, and options.

Feedback and learning from security-related decisions. The learning situation created by many common security and risk decisions does not help either. In a usual learning situation, behavior is shaped by positive reinforcement when we do something "right." We do something good, we are rewarded. In the case of security, when the user does something good, the reinforcement is that bad things are less likely to happen. There is seldom an immediate

reward or instant gratification, which can be a powerful reinforcer in shaping behavior.

In another common learning situation, behavior is shaped by negative reinforcement when we do something "wrong." We do something bad, we suffer the consequences. In the case of security, when the user does something bad, the negative reinforcement may not be immediately evident. It may be delayed by days, weeks, or months if it comes at all. Cause and effect is learned best when the effect is immediate and the anti-security choice often has no immediate consequences. This

Table 1. Nine steps and six UIs are required to set file permissions on a public share in Windows Vista. It takes four steps just to find the settings.

makes learning consequences difficult except in the case of spectacular disasters.

Evaluating the security/cost trade-off. While the gains of security are generally abstract and the negative consequences are stochastic, the cost is real and immediate. Security is integrated into systems in such a way that it usually comes with a price paid in time, effort, and convenience—all valuable commodities to users.

For example, in the simplest case—restricting access to a public share in Microsoft's Windows Vista to a group of users—requires about nine separate steps and six distinct user interfaces (see Table 1). While each step seems small, they add up a real cost to users. In deciding what to do, users weigh the cost of the effort against the perceived value of the gain (safety/security) and the perceived chance that nothing bad would happen either way.

Making trade-offs between risk, losses, and gains. Given that security gains are often intangible, the costs known, and the negative consequences involve probabilities, we can look at several known factors at play when people evaluate risks, costs, and benefits.

Users are more likely to gamble for a loss than accept a guaranteed loss. First of all, people react to risk differently depending on whether they think they are primarily gaining something or losing something. Tversky and Kahneman [8] showed that people are more likely to avoid risk when alternatives are presented as gains and take risks when alternatives are presented as losses. For example, consider the following scenario where a person has to decide between two options presented as gains:

Scenario 1:

- A) Gain \$5 at no risk

B) Gain \$10 if a coin toss lands heads up

When Tversky and Kahneman used a similar scenario, 72% of those surveyed chose the sure bet offered by option A because there was less risk and the outcome was guaranteed. Now consider a similar scenario presented as a choice between two losses:

Scenario 2:

- A) Lose \$5 guaranteed
- B) Lose \$10 if a coin toss lands heads up

When Tversky and Kahneman framed their scenario as a choice between losses, 64% of the respondents chose option B. People tended to focus on the chance to not lose anything offered in B compared to the sure loss guaranteed by option A.

When evaluating a security decision, the negative consequences are potentially greater of course, but the probability is generally less and often unknown. The principle holds true. When there is a potential loss in a poor security decision compared to the guaranteed loss of making the pro-security decision, the user may be inclined to take the risk. For example, consider the choice between two losses in a common security decision involving the download and installation of a digital certificate and ActiveX control from an unknown




Figure 1. Losses carry more value compared to gains when both are perceived as equal. For non-zero values, if value of loss (X) = value of gain (Y), then motivation of loss (A) > motivation of gain (B) (Adapted from Tversky and Kahneman [8]).

source. In this scenario, the primary goal is to view the Web page content:

Scenario 3:

- A) Do not install digital certificate and ActiveX control from unknown source and do not view the content of the Web page (fail on primary goal), guaranteed.
- B) Install digital certificate and ActiveX control from unknown source, view the Web page (accomplish primary goal), and take a chance that something bad happens.

Like Scenario 2, some users will chance that nothing bad will happen in order to achieve their primary goal than accept the task failure guaranteed by option A. Furthermore, if there are no immediate and obvious negative consequences incurred by option B, the user learns it is an acceptable decision and is more likely to repeat it in the future. The everyday security decisions end users make, like opening file attachments, are often presented in the form of losses as in Scenario 3.



People do not perceive gains and loss equally. This suggests that while a system designer may consider the cost of security effort small, the loss could be perceived as worse than the greater gain in safety. Put simply, the user must perceive a greater magnitude of gain than of loss.

Security as a secondary task.

People tend to focus more on the losses that will affect their immediate goal than the gains when making decisions under time pressure [12]. Users are often called on by the system to make a security decision while they are in the middle of an activity. In these cases, the user is often motivated to get on with the primary task as quickly as possible and, therefore, less likely to make a decision that further interrupts that task. In cases where users are prompted to install software updates, scan a file for viruses before opening, and so forth, users are less likely to comply when in the middle of another task, especially if in a hurry.

Losses perceived disproportionately to gains. People do not perceive gains and losses equally. Tversky and Kahneman [8] showed that when individuals perceive a gain and a loss to have the same value, the loss is more motivating in the decision (see Figure 2). In short, this means that a loss of \$100 is more adverse than a gain of \$100 is attractive to a decision maker.

This suggests that while a system designer may consider the cost of security effort small, the loss could be perceived as worse than the greater gain in safety. Put simply, the user must perceive a greater magnitude of gain than of loss.

IMPROVING SECURITY COMPLIANCE AND DECISION MAKING

Using the principles at work in security decision making, there are several avenues that may improve user security behavior.

Reward pro-security behavior. There must be a tangible reward for making good security decisions. Some suggest that corporate IT organizations would be encouraged to adopt stronger security practices if insurance companies offered lower premiums to those who protect themselves by certain measures [5]. Like-




Figure 2. Can you spot the security message? Message dialogs often look similar enough that no message stands out as more important than others.

wise, end users must be motivated to take pro-security actions. Increasing the immediate and tangible reward for secure actions may increase compliance. One form of reward is to see that the security mechanisms are working and that the action the user chose is, in fact, making them safer. This makes safety a visible gain when evaluating gains and losses in a security decision.

A good example of this is when an antivirus or antispyware product finds and removes malicious code. In these cases, the security application often issues a notification that it has found and mitigated a threat. This is an effective way for a security system to prove its value to the user by showing there was a risk and that the system protected them. By returning to the access control scenario for file sharing, it would be possible to report attempts at unauthorized access to the file owner.

Improve the awareness of risk. As discussed earlier, people often believe they are at less risk compared to others. One way to increase security compliance is to increase user awareness of the risks they face. This could be achieved through user training and education in general but should also be built into systems to support specific events.

One classically deficient area in the security of systems is messages and alerts. Security messages often resemble other messages dialogs (Figure 2). As a result, security messages may not stand out in importance and users often learn to disregard them.

To avoid the response bias problems faced by most message dialogs, security messages should be instantly distinguishable from other message dialogs. Security messages should look and sound very different (illustrated in Figure 3). This helps mitigate the blasé attitude with which users attend to the information. Once the message dialog has the user's attention, they are more likely to read




Figure 3. Can you spot the security message? (Part 2) Well-designed security messages have distinct visual and auditory properties that make them stand apart from all other message dialogs and indicate the criticality of the message.

and consider the choices given to them.

Catch corporate security policy violators. Increasing the awareness of risk could also mean increasing the likelihood that a corporate user is caught violating security policy. Having a corporate security policy that is not monitored or enforced is tantamount to having laws but no police. If the security systems have good auditing capabilities and are watched by event monitoring systems, users who make poor security decisions could be “caught” in a way. This would serve as an immediate negative consequence by itself. Like automated systems at traffic lights that snap pictures and issue violations to drivers that run red lights, users who make poor security decisions could receive automated email notifications of their actions and the corporate policy or safe computing practice. In general, the best deterrent to breaking the rules is not the severity of consequences but the likelihood of being caught.

Reduce the cost of implementing security. Obviously, if users need to take additional steps to increase their level of security, they will be less likely to do so. As the cost of implementing security increases, the overall value of the decision decreases. To accomplish a task, users often seek the path of least resistance that satisfies the primary goal. It should be common knowledge that in making the secure choice the easiest for the user to implement, one takes advantage of normal user behavior and gains compliance.

Another way to reduce the cost of security is, of course, to employ secure default settings. Most users never change the default settings of their applications. In this way, one increases the cost to make non-secure decisions in terms of time and effort. While good default settings can increase security, system designers must be careful that users do not find an easier way to slip around them. For example, users who are directed by their IT departments to use strong passwords across multiple systems are more likely to write them down [1].

CONCLUSION

Core to security on an everyday basis is the compliance of the end user, but how do we get them to make good decisions when they are often the weakest link in the chain? Users must be less motivated to choose anti-security options and more motivated to choose pro-security options. Obviously, no one would suggest training end users with USB devices that deliver an electric shock or food pellet reward based on their actions. But, generally speaking, we can increase compliance if we work with the psychological principles that drive behavior.

The ideal security user experience for most users

would be none at all. The vast majority would be content to use computers to enrich their lives while taking for granted a perfectly secure and reliable infrastructure that makes it all possible. Security only becomes a priority for many when they have problems with it. However, now, and in the foreseeable future, users are in the control loop. We must design systems with an understanding that, at some point, must make a decision regarding security. The question is, what will they decide? **C**

REFERENCES

1. Adams, A. and Sasse, A.S. Users are not the enemy. *Commun. ACM* 42, (1999) 40–46.
2. Borgida, E., and Nisbett, R.E. The differential impact of abstract vs. concrete information on decisions. *J. Applied Social Psychology* 7 (1977) 258–271.
3. Dhamija, R., Tygar, J.D., and Hearst, M. Why phishing works. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Montreal, Quebec, Canada, Apr. 22–27, 2006). R. Grinter, T. Rodden, P. Aoki, E. Cutrell, R. Jeffries, and G. Olson, Eds. ACM, New York, 581–590.
4. Downs, J.S., Holbrook, M., and Cranor, L.F. Behavioral response to phishing risk. In *Proceedings of the Anti-Phishing Working Groups 2nd Annual Ecrime Researchers Summit* (Pittsburgh, PA, Oct. 4–5, 2007). ACM, New York, 37–44.
5. Greenwald, S.J., Olthoff, K.G., Raskin, V., and Ruch, W. The user non-acceptance paradigm: INFOSEC's dirty little secret. *New Security Paradigms Workshop*, 35–43. ACM, New York.
6. Slovic, P., Fischhoff, B., and Lichtenstein, S. Facts versus fears: Understanding perceived risks. *Judgment under Uncertainty: Heuristics and Biases*. D. Kahneman, P. Slovic, and A. Tversky, eds. Cambridge University Press, New York, 1986, 463–489.
7. Smetters, D.K. and Grinter, R.E. Moving from the design of usable security technologies to the design of useful secure applications. *New Security Paradigms Workshop*. ACM, New York, 2002, 82–89.
8. Tversky, A. and Kahneman, D. Rational choice and the framing of decisions. *J. Business* 59 (1986), 251–278.
9. Whitten, A. and Tygar J.D. Why Johnny can't encrypt: A usability evaluation of PGP 5.0. In *Proceedings of the 8th USENIX Security Symposium* (1999). USENIX Association, Berkeley, CA, 169–184.
10. Wright, P. The harassed decision maker: Timer pressure, distractions, and the use of evidence. *J. Applied Psychology* 59, (1974), 555–561.
11. Yee, K.P. User interaction design for secure systems. *Proceedings of the 4th International Conference on Information and Communications Security*. Springer-Verlag, London, 2002.
12. Zurko, M.E. and Simon, R.T. User-centered security. *New Security Paradigms Workshop*. ACM, New York, 27–33.

RYAN WEST (ryan.west@acm.org) has conducted academic research in risk and decision making and applied research in areas ranging from medical mistakes to computer security. He currently works as a design researcher at Dell, Inc., Austin, TX.

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.

© 2008 ACM 0001-0782/08/0400 \$5.00

DOI: 10.1145/1330311.1330320

