Incentive Design for Home Computer Security

Poster Submission

People are the weakest link in security (Anderson, 1993). People write passwords on sticky notes on the screen. People don't patch their home systems and become botnet zombies. People choose whether to label a patch critical or just recommended. Our motivating insight is that these actions generally reflect *motivated behavior* in response to the configuration of incentives confronting individuals.

Since behavior is motivated by the goals and preferences of the individual, this behavior can be altered by designing appropriate incentives. By carefully structuring the benefits received from using a technology, we can induce users to make choices that most benefit the system. Along with some colleagues, we are developing a methodology for incentive-centered design of technology systems. We are working to provide guidelines and examples of how to carefully structure benefits to induce appropriate user choices.

We are applying these technology design ideas to a major open problem in computer security: botnets. Botnets are large collections of computers (called *zombies*) that are under the control of a single attacker. Botnets are behind a number of large security problems including spam email, distributed denial of service attacks, and multiple types of fraud and extortion (Ratliff, 2006). A significant part of the problem concerns security vulnerabilities inherent in the design of operating systems, network protocols and middleware. We do not address this well-studied issue. Instead, we focus on the problem that many zombies result from home computers that are poorly administered; that is, they are left more vulnerable than necessary given the current state of protective software. Home computer users frequently lack the skills necessary to properly secure their computer, and to properly clean the computer once it has been compromised. By providing appropriate incentives, it may be possible to induce these home users to make better choices in securing their computers.

An individual's use of software is largely driven by his or her perception of the direct benefits and costs of use (including the costs of learning the technology). The problems of non-use and mis-use are especially great for information security technologies for at least two reasons. First, many of the benefits of maintaining a secure system accrue not to the user, but to others. A home computer user rarely suffers from the insecurity he causes; it is the victims of the botnet that benefit from increased security. Ratliff (2006) describes how botnets can be used for extorting protection money from online businesses. Second, due to the nature of security systems, users are often not well-informed about the benefits to themselves. Most security systems are not directly productive; they exist to prevent productivity losses. As such, there is little feedback to users as to their own benefits (e.g. which losses were avoided) from their security choices. On the other hand, costs of recommended security behavior are usually more obvious, and thus receive more weight in user decisions. For example, CERT recommends[1] turning off Java and JavaScript, which

[1] `http://www.cert.org/tech_tips/home_networks.html`

will cripple many popular websites such as Google GMail, MSN Games, and most so-called Web 2.0 services.

Integrating Ideas from Economics and Psychology

The botnet problem is not entirely unique, of course. It shares properties with a well studied problem in economics — the voluntary provision of public goods. A public good is any good that is both *non-excludable* and *non-rivaled.* A good is *non-excludable* if no one (in a well-defined group) can be prevented from benefiting from the good. A *non-rivaled* good can be used by many people at once, where one person using the good does not prevent others from also using the good. Public goods usually have a problem with free riding — people using the good without contributing to its creation. This problem occurs when the public good is provided through voluntary contributions. Individual people will rationally choose to contribute less than optimally. Andreoni (2006) and Chen (2003) provide good summaries of public goods research.

Thanks to botnets, home computer security is an instance of a public good. Increasing the security of a home computer increases the difficulty of forming a botnet, benefiting all of society. Conversely, a person can free ride by permitting his or her computer to remain insecure. Currently home computer security is voluntarily provided, and evidence indicates that home users are insufficiently securing their computers, leading to large numbers of easily-compromised zombies. This is a well-known observation, first observed by Camp and Wolfram (2000), and the literature is summarized by Anderson and Moore (2006) and Wash and MacKie-Mason (2007).

Psychology also has interesting theories to add to this problem. It identifies the problem as one of social loafing, which is the reduction in motivation when individuals work collectively. It is a problem of perception — individuals only social loaf when he or she perceives a group outcome rather than separate individual outcomes.

A meta-analysis (Karau and Williams, 1993) of the social loafing literature proposed a useful model that integrates the results of many experiments into a coherent theory. Basically, people are less motivated to work toward group goals when they cannot see as strong of a connection between their effort and the final outcome that they value. Therefore, social loafing can be reduced by increasing the perceived value and importance of the individual contributions. Ling et al. (2005) tested a number of useful design principles derived from this theory in the context of contributions to an online community.

Designing a Social Firewall System

We are using these ideas to design a new personal intrusion detection system, also known as a personal firewall. Personal firewalls monitor the local computer for events such as network accesses or applications being run. These systems then consult a policy to determine whether this event should be permitted to happen, or denied with an error. A good policy can limit the system to only be able to do what the user wants it to do, and nothing more. Most personal firewall systems have the ability to leave part (or most) of the policy incomplete, and prompt users interactively as needed. However, when prompted few users have the knowledge or experience necessary to make a 'good' choice. Our system tries to help users with this decision by sharing information from other users faced with a similar decision.

A simple mechanism for sharing information between users would simply share and aggregate policy decisions. Whenever an application attempts to access the Internet, the user will be presented with a dialog box asking them to "Permit" or "Deny" the access. This binary decision could then be sent to a central server to be aggregated with other people's decisions. Future users, faced with the same decision, can see behavior of others ("80% of users have chosen 'Permit' ") to help with their choice. This might work based on a 'wisdom of crowds' (Surowiecki, 2004) idea that while individual users are imperfect, the aggregate can make the correct choice. But is this what will happen?

We model the user's decision in an economic model to attempt to predict user decisions. An appropriate model from economics is called 'information cascades' (Bikhchandani et al., 1998). We assume users prefer to make the 'right' choice between the two options, but there is uncertainty as to which choice is 'right.' The user also has some private information, a signal (like a gut feeling), that is accurate p percent of the time, where $p > 0.50$. This means that on average, the users' gut feelings are correct, and if properly aggregated, the group could do better than if everyone chose individually. However, an important result was proved by Bikhchandani et al. (1992): there is a non-trivial probability that all users will choose incorrectly, despite receiving informative signals. This comes because users can only observe the results of other user's discrete choices and cannot observe signals. If a relatively small number of users who make the first few choices have an erroneous signal, then all the subsequent users will rationally ignore their own signal and follow the previous users. Intuitively, if your gut says that a certain action should be 'permit'ed but everyone before you has chosen 'deny,' then it is rational to question your gut and also choose 'deny.' However, by doing so you don't provide any information about your 'permit' signal to future users. Future users encounter the same situation, observing everyone else choosing 'deny,' and will choose 'deny' for similar reasons. Anderson and Holt (1997) have validated this effect in the lab with human subjects.

We are exploring additional ideas for user contributions. It may be useful to have users report if they are compromised by a virus, worm, or hacker. Such occasional outcome information can help avoid the information cascade described in the previous paragraph. Additionally, it may be possible to collect comments from users describing why a given choice was made. These comments could be filtered, ordered, and voted upon much like reviews on `amazon.com` to provide useful information to other users. The difficulty in all these situations is motivating appropriate contributions by users. We hope to use design principles from the literature mentioned above from economics and psychology to motivate useful contributions.

References

Lisa Anderson and Charles Holt. Information cascades in the laboratory. *American Economic Review*, 87(5):847–862, December 1997. URL `http://links.jstor.org/sici?sici=0002-8282(199712)87%3A5%3C847%3AICITL%3E2.0.CO%3B2-9`.

Ross Anderson. Why cryptosystems fail. In *CCS '93: Proceedings of the 1st ACM conference on Computer and communications security*, pages 215–227. ACM Press, 1993. ISBN 0897916298. doi: 10.1145/168588.168615. URL `http://portal.acm.org/citation.cfm?id=168615`.

Ross Anderson and Tyler Moore. The economics of information security. *Science*, 314(5799):

610–613, October 2006. doi: 10.1126/science.1130992. URL http://www.sciencemag.org/cgi/content/abstract/314/5799/610.

James Andreoni. Philanthropy. In S-C. Kolm and J. Mercier Ythier, editors, *Handbook of Giving, Reciprocity and Altruism*, pages 1201–1269. North Holland, Amsterdam, 2006. URL http://econ.ucsd.edu/~jandreon/WorkingPapers/Philanthropy.pdf.

Sushil Bikhchandani, David Hirshleifer, and Ivo Welch. A theory of fads, fashion, and cultural change as information cascades. *Journal of Political Economy*, 100(5):992–1026, October 1992. URL http://www.jstor.org/view/00223808/di980598/98p00557/0.

Sushil Bikhchandani, David Hirshleifer, and Ivo Welch. Learning from the behavior of others: Conformity, fads, and informational cascades. *Journal of Economic Perspectives*, 12(3):151–170, Summer 1998. URL http://www.jstor.org/view/08953309/di014715/01p0058j/0.

L Jean Camp and Catherine Wolfram. Pricing security. In *Proceedings of the Information Survivability Workshop*, 2000. URL http://www.springerlink.com/index/m44317165u727779.pdf.

Yan Chen. Incentive-compatible mechanisms for pure public goods: A survey of experimental research. In Charles Plott and Vernon Smith, editors, *The Handbook of Experimental Economics Results*, volume 1. Elsevier Science Publishing Company, April 2003. URL http://www.si.umich.edu/~yanchen/papers/chenpub.doc.

Steven Karau and Kipling Williams. Social loafing: A meta-analytic review and theoretical integration. *Journal of Personality and Social Psychology*, 65(4):681–706, 1993.

K. Ling, G. Beenen, P. Ludford, X. Wang, K. Chang, D. Cosley, D. Frankowski, L. Terveen, A. M. Rashid, P. Resnick, and R Kraut. Using social psychology to motivate contributions to online communities. *Journal of Computer-Mediated Communication*, 10(4), 2005. URL http://jcmc.indiana.edu/vol10/issue4/ling.html.

Evan Ratliff. The zombie hunters. *The New Yorker*, October 10 2006. URL http://www.newyorker.com/fact/content/articles/051010fa_fact.

James Surowiecki. *Wisdom of Crowds*. Little Brown, 2004.

Rick Wash and Jeffrey K. MacKie-Mason. Security when people matter: Structuring incentives for user behavior. In *International Conference on Electronic Commerce*, Minneapolis, MN, 2007. ACM. URL http://www-personal.si.umich.edu/~rwash/pubs/icec702w-wash.pdf.