Network Security Games: Combining Game Theory, Behavioral Economics, and Network Measurements

Nicolas Christin

Carnegie Mellon University
Information Networking Institute and CyLab
Pittsburgh, PA 15123
nicolasc@cmu.edu

Computer and information networks are a prime example of an environment where negative externalities abound, particularly when it comes to implementing security defenses. A typical example is that of denial-of-service prevention: ingress filtering, where attack traffic gets discarded by routers close to the perpetrators, is in principle an excellent remedy, as it prevents harmful traffic not only from reaching the victims, but also from burdening the network situated between attacker and target. However, with ingress filtering, the entities (at the ingress) that have to invest in additional filtering are not the ones (at the egress) who mostly benefit from the investment, and, may not have any incentive to participate in the scheme. As this example illustrates, it is important to understand the incentives of the different participants to a network, so that we can design schemes or intervention mechanisms to re-align them with a desirable outcome.

Game theory offers a solid bedrock for formally assessing the incentives of non-cooperative participants. In this talk, I will start by discussing a framework for network security games [4, 5] that we devised to help model how rational, individual, end-users would respond to security threats in large-scale networks. We decouple security decisions between self-insurance (which does not present any externalities) and self-protection (which does present externalities). Assuming fully rational players, acting with perfect information, and with the ability to perfectly execute their security decisions, we can derive results showing how much of a negative impact externalities can have on security decision-making. I will also introduce extensions of this work which deal with more limited information cases [6].

However, humans are not acting perfectly rationally when it comes to security decision-making. Prospect theory tells us that humans tend to be risk-averse when it comes to gains; and risk-seeking when it comes to losses [7]. In other words, people tend to "gamble" more than they should when it comes to security risks. I will further show, through an experiment related to our framework [3] that in addition to these biases, users have very limited "computational" ability; in particular, they seem unable to strategize over more than one decision variable at a time. I will present complementary experimental results [1] that suggest that Peltzman effects [11] also apply in computer security. Much like drivers wearing seat belts or helmets tend to drive faster, people tend to behave more insecurely online when they believe they have adopted secure precautions, such as installing an anti-virus scanner. As a result, I will postulate that game-theoretic modeling either needs to be complemented by behavioral analysis (for

To appear as a keynote in the Second Conference on Decision and Game Theory for Security (GameSec'11)
College Park, MD. November 2011.
The original publication is available on www.springerlink.com

individual users) or is better suited to describing institutional users (e.g., corporations, governments, ISPs...).

In the second part of this presentation, I will make the case that to provide improved resilience to attacks, we must be simultaneously mindful of the capabilities of the attackers, as well as their own economic incentives. Indeed, since the early- to mid-2000's, attackers have become mostly profit-driven [9]. By primarily conditioning their actions on their best financial interest, attackers are more and more behaving rationally in the economic sense of the term, and are considerably more predictable than attackers driven by less mundane ideals. Trying to disrupt the economic incentives that drive attackers to commit their forfeits appears to be a defensive strategy worth investigating, as a complement to the technical approaches that have been proposed.

I will contend that modeling attacker behavior is easier than modeling defender behavior. First, attackers show much stronger economic rationality than defenders: the success of the attack directly conditions their profits, while for defenders, security precautions are often viewed as sunk costs. Second, attackers' actions are often publicly observable: attacks such as phishing, malware distribution or search-engine manipulation leave a visible footprint. I will present a couple of recent measurement studies we conducted [2, 8, 10] in an effort to acquire more information on attacker behavior, and will show that a priori disparate attacks all present *concentration points*. Specifically, very often, the number of actual perpetrators behind entire class of attacks (e.g., search engine manipulation) are small. This in turn helps us inform security games where we want to model attackers as players, rather than exogenous entities.

Finally, I will conclude by presenting a roadmap for future research integrating network measurements and formal, game theoretic, modeling.

References

1. N. Christin, S. Egelman, T. Vidas, and J. Grossklags. It's all about the Benjamins: Incentivizing users to ignore security advice. In *Proceedings of IFCA Financial Cryptography'11*, Saint Lucia, March 2011.
2. N. Christin, S. Yanagihara, and K. Kamataki. Dissecting one click frauds. In *Proc. ACM CCS'10*, Chicago, IL, October 2010.
3. J. Grossklags, N. Christin, and J. Chuang. Predicted and observed behavior in the weakest-link security game. In *Proceedings of the 2008 USENIX Workshop on Usability, Privacy and Security (UPSEC'08)*, San Francisco, CA, April 2008.
4. J. Grossklags, N. Christin, and J. Chuang. Secure or insure? A game-theoretic analysis of information security games. In *Proceedings of the 2008 World Wide Web Conference (WWW'08)*, pages 209–218, Beijing, China, April 2008.
5. J. Grossklags, N. Christin, and J. Chuang. Security and insurance management in networks with heterogeneous agents. In *Proceedings of the 9th ACM Conference on Electronic Commerce (EC'08)*, pages 160–169, Chicago, IL, July 2008.
6. B. Johnson, J. Grossklags, N. Christin, and J. Chuang. Are security experts useful? Bayesian Nash equilibria for network security games with limited information. In *Proceedings of the 15th European Symposium on Research in Computer Security (ESORICS 2010).*, pages 588–606, Athens, Greece, September 2010.
7. D. Kahneman and A. Tversky. Prospect theory: An analysis of decision under risk. *Econometrica*, XLVII:263–291, 1979.

8. N. Leontiadis, T. Moore, and N. Christin. Measuring and analyzing search-redirection attacks in the illicit online prescription drug trade. In *Proceedings of USENIX Security 2011*, San Francisco, CA, August 2011.
9. T. Moore, R. Clayton, and R. Anderson. The economics of online crime. *Journal of Economic Perspectives*, 23(3):3–20, Summer 2009.
10. T. Moore, N. Leontiadis, and N. Christin. Fashion crimes: Trending-term exploitation on the web. In *Proceedings of ACM CCS 2011*, Chicago, IL, October 2011.
11. S. Peltzman. The effects of automobile safety regulation. *Journal of Political Economy*, 83(4):677–726, August 1975.