An Attacker-Defender Game in System Exploitation

Extended Abstract

Kay-Yut Chen
Information Systems and Operations Management
College of Business
University of Texas
Arlington TX 76019
kychen@uta.edu

Jingguo Wang
Information Systems and Operations Management
College of Business
University of Texas
Arlington TX 76019
jwang@uta.edu

Laura Marusich
U.S. Army Research Laboratory South Field Element
University of Texas
Arlington TX 76019
laura.r.marusich.ctr@mail.mil

Jonathan Bakdash
1. U.S. Army Research Laboratory South Field Element
University of Texas
Dallas TX 75080
2. Department of Psychology, Counseling, and Special Education,
Texas A&M
Commerce 75429
jonathan.z.bakdash.civ@mail.mil

Keywords: Information security, game theory, behavioral game, vulnerability

An Attacker-Defender Game in System Exploitation

Recent outbreaks of ransomware and repeated massive data breaches, recently of Equifax, highlight the urgency of understanding how organizations effectively patch and protect their systems in defending against cyberattacks. Game theoretical approaches have been widely applied to study organizations' optimal security spending (Cavusoglu et al. 2008; Wang et al. 2008). Such approaches take account of strategic behaviors of hackers and defenders (Anderson and Moore 2006). For instance, incorporating hackers' strategic responses to a firm's security investment, Cavusoglu et al. (2008) compare game theory and decision theory approaches for the investment decisions. What drive vendors and organizations to patch vulnerable systems has been also investigated (Arora et al. 2010; August and Tunca 2006; August and Tunca 2008; August and Tunca 2011; Kim et al. 2011) . Most studies in this aspect focus on the design of vendors' and organizations' incentives to facilitate their patching (August et al. 2014; August and Tunca 2008; Choudhary and Zhang 2015), and a few papers examine the impact of the vulnerability market (Kannan and Telang 2005) and the disclosure of software vulnerability on vendors' patch development (Arora et al. 2010). These studies are largely based on the assumption that hackers' efforts are exogenous and independent of vendors' and organizations' patching behavior. August and Tunca (2011) consider zero-day attacks, but also assumes the probabilities of zero-day attacks are not influenced by patching decision. Studies on patching with strategic hackers is still to be explored (August et al. 2014). Hackers' strategic reaction to patching may lead to different implications (Kannan and Telang 2005).

These studies provided significant insights to security investment and vulnerability management. However, all models assume that individuals make rational choices given available information. Research in behavioral economics have well documented that the rationality of individuals are bounded (Simon 1947), and decision makers suffers a range of cognitive biases when making choices under uncertainty (Tversky and Kahneman 1974). How defenders and hackers interact and behave in the episodes of cyberattacks, in light of bounded rationality and cognitive biases, is an unaddressed question.

We model a two time-period attacker-defender cyber security game. The goal is to be able to develop insights through game theoretic analysis, and behavioral human-based experiments, via multiple rounds of attacker-defender interactions. We show, under certain conditions, the optimal strategy for the defender is to eliminate some but not all potential exploits. Moreover, we also find that attackers deviate from rational strategies. Specifically, we consider a case with a single hacker to attack and a single defender to protect a system or a computer. The basic unit for attacking and defending is a "vulnerability" or "exploit" (we use the term "vulnerability" and "exploit" interchangeably). We assume the system has a total of vulnerabilities = N. We model TWO types of vulnerabilities. The first type is referred to as "known vulnerabilities". N_k= number of total known vulnerabilities. In practice, there is a database of the known vulnerabilities (such as National Vulnerability Database), and N_k is known. The second type is referred to as "zero-day vulnerabilities". The number of total zero-day vulnerabilities, N_o, is unknown to both attackers and defender.

The attacker may use zero-day vulnerabilities at a cost higher than that of the known vulnerabilities. There are m_k known vulnerabilities on the system, which is common knowledge. The attacker decides k_k, the number of known exploits to acquire/use, each costing c_A. We further assume there are m_o zero-day vulnerabilities on the system, but again, this number is

not known to the attack and defender. However, there is a common belief, shared by the attacker and defender, over both the total number of zero-day vulnerabilities and the specific number on the computer. Let $g(N_o\ m_o)$ be the joint distribution of N_o and m_o. The attacker also chooses k_o the number of zero-day exploits to use. Each zero-day exploit costs c_o where $c_o > c_A$, that the zero-day exploits cost more than the known ones.

The defender chooses s, the number of known vulnerabilities to eliminate each costing c_D. Note that in this basic setting all vulnerabilities/exploits are assumed to cost equally, and the specific ones (to be used by an attacker or eliminate by a defender) are randomly chosen. The system will be successfully hacked if either the attacker uses a known exploit that matches a known vulnerability present on the computer, or use a zero-day exploit that matches a zero-day exploit on the computer.

While the model can be extended to cover a variety of attacker-defender interactions, we focus on two specific scenarios. In the first scenario, the defender moves first, and chooses the number of exploits to eliminate (s). Note that he can choose not to eliminate any. Subsequently, the attacker chooses k_k and k_o, the number of known and 0-day exploits to use. In the second scenario, the model is extended to two decision rounds (i.e. the defender moves again after the first-round result is known). There is a positive probability that the defender can detect a successful 0-day attack, and the particular exploit will be eliminated from the second round.

We show, with numerical analysis, given the right parameters, there is a Nash equilibrium where the defender eliminates some but not all of the exploits, and the attacker only uses 0-day exploits. (i.e. $0<s< m_k$, $k_k=0$ and $k_o>0$). We further show that a similar Nash equilibrium can be constructed for the case with detection. In that case, the number of 0-day exploits (k_o) will be lower because it is worthwhile for the attacker to "save some" for the next round of attack.

We have conducted two behavioral experiments to verify if Nash equilibrium is a good predictor of participants' behavior in these scenarios. In the baseline treatment, the median of the defender's decision of number of exploits to eliminate is 3, consistent with the Nash prediction (Wilcoxon p-value = 0.18). However, we observe substantial noise in the decisions which suggests bounded rationality and cognitive biases in the decision space. The decisions range from 0 to 6, covering all possible choices, with a standard deviation of 1.45. For the attacker, the median number of known exploits is 1, and it is significant higher than 0, the Nash equilibrium prediction (Wilcoxon p-value basically 0). We also observe a long-tailed distribution with a maximum $k_k=12$. Hence, the use of known exploits cannot be completely eliminated as predicted by rational theory. The median of the number of 0-day exploits used is 4 but it is not significantly different from the Nash prediction of 5 (Wilcoxon p-value = 0.44). Again, we observe a high level of decision noise (standard deviation = 2.38). We conclude that while there is substantial amount of decision noise, the subjects do response to the reward and cost structure. The Nash equilibrium seems to be able to predict the medians of the defender's decision (number of exploits to eliminate) and the attacker's use of 0-day exploits. However, the use of known exploits is depressed but not eliminated in the scenario, contrary to what rational theory predicts.

The detection treatment is calibrated so that it has the same equilibrium prediction as the baseline case, except that the use of 0-day exploits, in the first decision round, is reduced (from 5 to 4) because the attacker should "save" a 0-day exploits for the next round in light of the detection threat. However, we found that individuals are much more aggressive in the detection treatment. The defender chooses a higher number of known exploits to eliminate (median = 4, p-value < 0.01). The attacker chooses both a higher number of known exploits (median = 2, p-

value < 0.01), and a higher number of 0-day exploits (median = 5, p-value < 0.01). We shall discuss more findings, and their interpretations, in our presentation.

Our study is among the first to relying on behavioral economics experiments to understand the strategic interactions between the hacker and the defender, and intend to show how their behavior may deviate from the predictions based on the traditional game theoretical framework. This study aims to provide better practical implications for security management.

References

Anderson, R., and Moore, T. 2006. "The Economics of Information Security," *Science* (314:5799), pp. 610-613.

Arora, A., Krishnan, R., Telang, R., and Yang, Y. 2010. "An Empirical Analysis of Software Vendors' Patch Release Behavior: Impact of Vulnerability Disclosure," *Information Systems Research* (21:1), pp. 115--132.

August, T., August, R., and Shin, H. 2014. "Designing User Incentives for Cybersecurity," *Commun. ACM* (57:11), pp. 43--46.

August, T., and Tunca, T. I. 2006. "Network Software Security and User Incentives," *Management Science* (52:11), pp. 1703--1720.

August, T., and Tunca, T. I. 2008. "Let the Pirates Patch? An Economic Analysis of Software Security Patch Restrictions," *Information Systems Research* (19:1), pp. 48--70.

August, T., and Tunca, T. I. 2011. "Who Should Be Responsible for Software Security? A Comparative Analysis of Liability Policies in Network Environments," *Management Science* (57:5), pp. 934--959.

Cavusoglu, H., Raghunathan, S., and Yue, W. T. 2008. "Decision-Theoretic and Game-Theoretic Approaches to It Security Investment," *Journal of Management Information Systems* (25:2), pp. 281–304.

Choudhary, V., and Zhang, Z. 2015. "Research Note—Patching the Cloud: The Impact of Saas on Patching Strategy and the Timing of Software Release," *Information Systems Research* (26:4), pp. 845--858.

Kannan, K., and Telang, R. 2005. "Market for Software Vulnerabilities? Think Again," *Management Science* (51:5), pp. 726--740.

Kim, B. C., Chen, P.-Y., and Mukhopadhyay, T. 2011. "The Effect of Liability and Patch Release on Software Security: The Monopoly Case," *Production and Operations Management* (20:4), pp. 603--617.

Simon, H. A. 1947. *Administrative Behavior: A Study of Decision-Making Processes in Administrative Organization* New York, USA: Macmillan Publishers.

Tversky, A., and Kahneman, D. 1974. "Judgment under Uncertainty: Heuristics and Biases," *Science* (185:4157), pp. 1124-1131.

Wang, J., Chaudhury, A., and Rao, H. R. 2008. "A Value-at-Risk Approach to Information Security Investment," *Information Systems Research* (19:1), pp. 106–120.