Applying Norms and Sanctions to Promote Cybersecurity Hygiene

Extended Abstract

Shubham Goyal
Amazon
Seattle, Washington, USA
gsshubha@ncsu.edu

Nirav Ajmeri
NC State University
Raleigh, North Carolina, USA
najmeri@ncsu.edu

Munindar P. Singh
NC State University
Raleigh, North Carolina, USA
mpsingh@ncsu.edu

ABSTRACT

Cybersecurity breaches cause enormous harm to the safety, privacy, and prosperity of individuals and organizations. Many security breaches occur due to people not following security regulations such as applying software patches, updating software applications, and so on. We term these regulations as *cybersecurity hygiene*. This paper investigates different sanctioning mechanisms with respect to the success in establishing these regulations for cybersecurity hygiene. Our findings have implications for workforce training to promote cybersecurity.

KEYWORDS

Cybersecurity; Normative systems; Sanctions; Agent societies

ACM Reference Format:
Shubham Goyal, Nirav Ajmeri, and Munindar P. Singh. 2019. Applying Norms and Sanctions to Promote Cybersecurity Hygiene. In *Proc. of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019), Montreal, Canada, May 13–17, 2019,* IFAAMAS, 3 pages.

1 INTRODUCTION

Organizations have created security regulations to promote cybersecurity hygiene but compliance with regulation is rarely adequate and security breaches continue to occur. Singh [7] points out the importance of cybersecurity as an application domain for multiagent systems (MAS), especially in light of modeling human behavior. Normative MAS provide an appropriate way of thinking about challenges such as Advanced Persistent Threat (APT) [3] because the actions or nonactions (e.g., carelessness) of one user can affect the outcomes for other users. In essence, the science of security must include considerations of *cybersecurity hygiene*, the theoretical foundations of which can be based on norms. Sanctions [4] provide a recognized means to promote establishment of norms but have not been studied in connection with cybersecurity.

We motivate three main varieties of sanctioning mechanisms, specifically, group, individual, and peer sanctions. *Individual* sanction is where the person who failed to comply with the regulation is sanctioned. *Group* sanction is where every member of the group is sanctioned when a subset of the group has failed to comply with the regulation. *Peer* sanction is when a member of the group, sanctions another member of the group for not following the regulations. Individual and group sanctions are applied by an administrator, a designated party who has the responsibility of monitoring and sanctioning, and peer sanctions are applied by users.

We empirically investigate the effects of these sanctioning mechanisms in promoting compliance with cybersecurity regulations as well as the detrimental effect of sanctions on the ability of users to complete their work. We do so by developing a game that emulates the decision making of workers in a research lab.

Contributions. To this end, we investigate how sanctions can promote cybersecurity hygiene. Specifically, we investigate two research questions in reference to the above-mentioned types of sanctions: group, individual, and peer.

- How effectively does a sanction type lead to improved cybersecurity hygiene?
- How detrimental is a sanction type to user productivity?

Approach and findings in brief. We develop a game to simulate a real-life work setting, such as a corporate office in which workers complete assigned tasks while using computers. Each player assumes the role of an office worker. Each player is challenged to complete assigned tasks (captured as points earned) along with maintaining the security of his or her computer. Failure to complete the security tasks may attract sanctions, causing loss in points earned or loss in opportunity to earn points.

We conduct several experiments in this setting. We find that workers complete more tasks and are sanctioned less often under individual sanctions than under group sanctions.

Individual and group sanctions are applied by an administrator, a designated party who has the responsibility of monitoring and sanctioning, and peer sanctions are applied by the users.

2 THE GAME MODEL

A *norm* characterizes sound or *normal* interactions among the participants of a social group, reflecting their mutual expectations from the system [6]. The administrator *expects* each worker to be security compliant all the time in addition to completing his or her project tasks. An agent can learn about the norms by experiencing *sanctions* or observing sanctions being applied on others [1, 2, 5].

To investigate the effect of the sanctioning to humans' security-related behavior, we model a multiagent system comprising agents who play the worker and manager roles.

The secure multiagent system model $O = \{A, C, M, E\}$ contains four components: A is a set of workers who perform their project tasks, security-related tasks and can sanction other workers. C represents a PC that is owned by the worker A. M is the unique

Proc. of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019), N. Agmon, M. E. Taylor, E. Elkind, M. Veloso (eds.), May 13–17, 2019, Montreal, Canada.

manager in charge of the system, and environment E is everything besides the previous three entities, including the attackers.

Workers complete their productivity task while maintaining the security of the PC. Workers have to choose between completing their project tasks and security tasks. In the absence of a sanctioning mechanism, they lack the motivation to complete security tasks. The *manager* is the central administrator responsible for maintaining the security of the system. The manager observes the security of all PCs and sanctions workers in the system who are not security compliant. We do not consider positive sanctions but employ negative sanctions as a consequence of not being security compliant. Outside of the system, there are potential *attackers*. They attack the system with the goal of compromising the PCs in the system.

Each PC in the system is associated with a worker and is in one of the three states: safe, vulnerable, and unusable. Initially, every PC is in the safe state. Failure to complete a security task on time by the worker owning the PC moves the PC to the vulnerable state. An external attack on a vulnerable PC makes it unusable. The worker owning the PC must complete the security task to transition an unusable or vulnerable PC back to safe state.

Each worker has a state—compliant or noncompliant. A worker is compliant only when his or her PC is in safe state.

3 GAME DESIGN

To investigate how our model fits into real life situation, we designed a web-based game as shown in Figure 1, following the model.

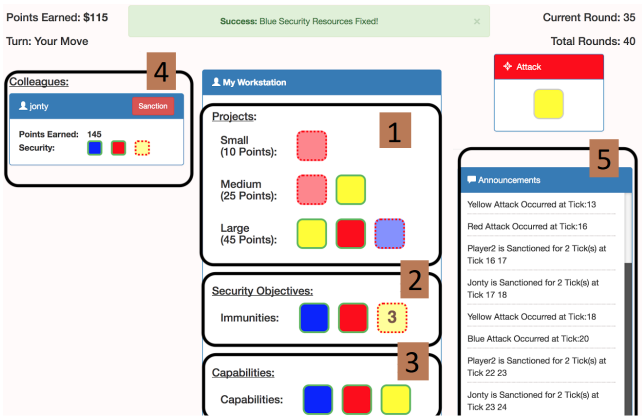

Figure 1: Game screenshot showing its five parts.

The game has three tasks: T_{blue}, T_{red} and T_{yellow}. The tasks can be completed by clicking the tiles in Part 1. Each task has a corresponding immunity and capability, as shown in Part 2 and Part 3, respectively. The game is divided into 40 rounds. Each player takes a turn in every round. There can be an attack at the beginning of a round. After an attack all the players lose the corresponding immunity. If the immunity is already lost, players lose the corresponding capability. Losing a capability means that player won't be able to complete the corresponding task.

To gain back the immunity, the player has to complete the immunity task by clicking on tiles in Part 2. Completing an immunity task and a regular task (Part 1 tasks) both takes a single round, but no point is awarded for completing immunity task.

There is a built in manager in game engine. At the beginning of every round the manager checks the immunity of all players and can sanction them if a player does not have immunity. If a player gets sanctioned, he loses two rounds in the game for each incomplete immunity.

The tasks in Part 1 are equivalent to project tasks in game model, immunity tasks are equivalent to security tasks and manager emulates the central administrator responsible for maintaining security of the system.

4 EXPERIMENTS AND RESULTS

We conducted a study on Amazon Mechanical Turk where we asked participants to play our game. Thirty participants participated in the study playing 107 games. The study was approved by our university's Institutional Review Board. We encouraged and rewarded players to give their best in the game by promising them a bonus based on the score in the game. Participants completed multiple surveys as a part of the study to note their feedback on different sanctioning mechanisms in the game.

Each participant played two games with group sanction and two games with individual sanction in a span of 60 minutes. All the game parameters other than the sanctioning method were kept constant throughout the study in all the games. We recorded every move made by a player and evaluated the data. We test significance via the two-tailed paired t-test.

After each game we asked the players, on a Likert scale of 1 (*not at all influential*) to 5 (*very influential*), how effective was the sanctioning mechanism. 77 percent of participants identified sanctions as a strong factor (4–5) in influencing their decisions.

Players were more compliant—completed more security tasks and were sanctioned less often—under individual than under group sanctions. Player completed more productivity tasks under individual than under group sanctions, indicating that individual sanctions impose a lower cost for achieving compliance.

Under both kinds of sanctions, players once sanctioned took almost the same time to get back to the *normal* state, indicating equivalent resilience for both sanctioning technique. This result was not statistically significant.

Peer sanctions were more prevalent under group than under individual sanctions. This fact can be explored in further studies to create a self-sustained system that maintains security without an external administrator.

5 CONCLUSIONS

We investigate the effectiveness of group, individual, and peer sanctions in promoting cybersecurity hygiene and improving productivity. We establish that individual sanctions are more effective than group sanctions in enforcing compliance with cybersecurity regulations. Peer sanctions can provide the basis for how a community of workers can self-regulate itself and promote cybersecurity hygiene.

6 ACKNOWLEDGMENTS

We thank Jon Doyle for valuable discussions, Hongying Du and Bennett Narron for contributions to an earlier version of the game, and the US Department of Defense for support through the Science of Security Lablet at NC State University.

REFERENCES

[1] Nirav Ajmeri, Hui Guo, Pradeep K. Murukannaiah, and Munindar P. Singh. 2018. Robust Norm Emergence by Revealing and Reasoning about Context: Socially Intelligent Agents for Enhancing Privacy. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI)*. IJCAI, Stockholm, 28–34.

[2] Giulia Andrighetto, Jordi Brandts, Rosaria Conte, Jordi Sabater-Mir, Hector Solaz, and Daniel Villatoro. 2013. Punish and Voice: Punishment Enhances Cooperation when Combined with Norm-Signalling. *PLoS ONE* 8, 6 (06 2013), 1–8.

[3] Ping Chen, Lieven Desmet, and Christophe Huygens. 2014. A Study on Advanced Persistent Threats. In *Proceedings of the 15th IFIP TC 6/TC 11 International Conference on Communications and Multimedia Security (CMS) (Lecture Notes in Computer Science)*, Vol. 8735. Springer, Aveiro, Portugal, 63–72.

[4] Luis G. Nardin, Tina Balke-Visser, Nirav Ajmeri, Anup K. Kalia, Jaime S. Sichman, and Munindar P. Singh. 2016. Classifying Sanctions and Designing a Conceptual Sanctioning Process for Socio-Technical Systems. *The Knowledge Engineering Review* 31, 2 (March 2016), 142–166.

[5] Bastin Tony Roy Savarimuthu and Stephen Cranefield. 2011. Norm Creation, Spreading and Emergence: A Survey of Simulation Models of Norms in Multiagent Systems. *Multiagent Grid Systems* 7, 1 (Jan. 2011), 21–54. http://dl.acm.org/citation.cfm?id=2019196.2019199

[6] Munindar P. Singh. 2013. Norms as a Basis for Governing Sociotechnical Systems. *ACM Transactions on Intelligent Systems and Technology (TIST)* 5, 1 (Dec. 2013), 21:1–21:23.

[7] Munindar P. Singh. 2015. Cybersecurity as an Application Domain for Multiagent Systems. In *Proceedings of the 14th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*. IFAAMAS, Istanbul, 1207–1212. Blue Sky Ideas Track.