Role of Intrusion-Detection Systems in Cyber-Attack Detection

Varun Dutt[1], Frederic Moisan[2], and Cleotilde Gonzalez[2]

[1] Applied Cognitive Science Laboratory
Indian Institute of Technology, Mandi, Kamand, India-175005
[2] Dynamic Decision Making Laboratory
Carnegie Mellon University, Pittsburgh, USA-15213
varun@iitmandi.ac.in, fmoisan@gmail.com, coty@cmu.edu

Abstract. Currently, little is known about how defenders' reliance on decision-support technology influences their decisions. Here, we designed a cyber-security game, where "hackers" decide whether to attack a computer network and "analysts" decide whether to defend the network based upon recommendations from IDS. We present results from an experiment with 200 participants randomly paired and assigned to one of four between-subjects conditions that varied in the IDS's availability (absent/present) and its accuracy (when present, it is 10%, 50%, or 90% accurate). Results revealed that proportion of attack and defend actions were similar and close to their Nash proportions when IDS was absent and when it was 50% accurate; but, these proportions were smaller and different from their Nash proportions when the IDS was inaccurate (10% accurate) or very accurate (90% accurate). Our results suggest that the presence of decision-support technology is likely to make defenders over rely on this technology.

Keywords: Behavioral cybersecurity · Intrusion detection system (IDS) · IDS accuracy · IDS availability · Instance-based learning theory

1 Introduction

Cyber-attacks, i.e., attempts by hackers to steal data and damage networks and systems, are increasing at an alarming rate [10, 16]. In fact, a recent survey by Cyber Ark suggests that currently, nations are at a greater risk from cyber-attacks compared to physical attacks [8]. According to Barack Obama, the cyber threat is one of the most serious economic and national security challenges [22]. In order to safeguard computer systems against cyber-attacks, the role of specialized human decision makers (i.e., Analysts) is indispensable [9, 12]. However, with the sudden and highly changing conditions in which analysts work, science has not yet caught-up with understanding of the basic cognitive and psychological processes that may influence analysts' work, and ultimately the safety of our information in the network. In this paper, we aim at a

better understanding of some of the fundamental processes in attack detection: the reliance on variable and inaccurate decision support.

Analysts protect a computer network by identifying, as early and accurately as possible, threats and non-threats during cyber-attacks [9]. However, analysts cannot directly observe hackers' actions on the network which are often like finding a needle in a haystack [1, 12, 20]. The existence of multiple and diverse sensors result in a large amount of network activity data. Analysts often need to rely upon cybersecurity tools, e.g., Intrusion Detection Systems (IDS) to organize and structure network activity; and, to help make information relevant, meaningful, and useful for an analyst to use [23]. However IDS possess inaccuracies; false-alarms (reporting an attack when there is none) and misses (not reporting an attack when there is one) are common, and it is up to the analyst to rely or not on the recommendation the IDS provides [13]. Hackers also know that IDS are not fault-free [5]. Therefore, it is important to not only understand how defenders react to inaccurate recommendations from IDS, but how attackers would also behave in the presence of inaccurate recommendation to the analyst.

We address these questions by relying on cybersecurity games, inspired on Behavioral Game Theory (BGT) [7] and on Instance-Based Learning Theory (IBLT) [9, 11], a theory of decision making in dynamic environments. Current research includes cybersecurity games [1] to document Nash equilibria for attacker-defender interactions while considering inaccuracies of the IDS [20]. This is a very promising area of research that can help us understand how behavioral interactions between hackers and analysts deviate from the mathematically determined optimal Nash behavior [4, 14, 15]. However, it is important to understand the source of deviation from optimal behavior, and BGT and IBLT can help in understanding actual human limitations based on preferences, motivations, memory and other psychological aspects of behavior.

In this paper, we use a cybersecurity game involving two players: a "hacker" and an "analyst" who are presented with imperfect and possibly unreliable decision support ("IDS"). We investigate the interaction between two human players, hackers and analysts, and document the deviations of this interaction from optimal behavior. Next, we explain the cybersecurity game that we designed for this research. Then, we motivate our hypotheses on the interaction between human hackers and analysts, while the human players play the cybersecurity game. We then present results from an experiment in which we manipulate the presence and the accuracy of the information provided by the IDS. Finally, we discuss the implications of these results for development and use of IDSs for cybersecurity.

2 An imperfect cybersecurity game with and without recommendations

Consider a firm and an anonymous individual (a potential hacker) that can penetrate the firm's network infrastructure to illegally access valuable information. This firm has invested in cybersecurity and an analyst, who is in charge of protecting the firm's assets, constantly protects the network. This situation can be translated into a simple strategic game where a hacker may either attack the firm's network (a) or do nothing (na), while the analyst may either defend the network (d) or do nothing (nd). Figure 1 shows an instantiation of a security game with these features. It is clear from Figure 1

that there exists no Nash equilibrium in pure strategies (in each of the four possible outcomes, one player is better off deviating). As a result, the only equilibrium solution in this game is in mixed strategies (i.e., selecting each action with some probability), which specifies the following: the hacker attacks with 0.2 ($\frac{1}{5}$) probability, while the analyst defends with 0.66 ($\frac{2}{3}$) probability. The corresponding expected payoffs are 0.0 for the hacker and -3.0 for the analyst.

		Analyst: Defend (d)	Analyst: Not Defend (nd)
Hacker	Attack (a)	-5, 5	10, -15
	Not attack (na)	0, -5	0, 0

Fig. 1. Payoff matrix in a security game played repeatedly between a Hacker and an Analyst.

We now extend the definition of the above security game by introducing an IDS that can alert the analyst regarding the decision made by the hacker (thus, the analyst does not see the actions of the hacker directly; rather, she gets messages from the IDS based upon hacker's decisions). The resulting game theoretic representation of the cybersecurity game that is played sequentially and involves a recommendation system is depicted in Figure 2. The hacker first makes a choice, followed by the IDS that reports the existence/absence of an attack to the analyst. In our security game, we define *pa* as the probability of the IDS to accurately predict the hacker's choice (a wrong prediction therefore occurs with probability *1-pa*). The report from the IDS is determined through probability *pa* (e.g., if the hacker attacks by choosing a, IDS reports an attack with probability *pa* and non-attack with probability *1-pa*). After receiving the IDS recommendation, the analyst makes a choice. The dotted lines in Figure 2 represent the analyst's information sets, each of which regroups states (nodes) that are undistinguishable by the analyst (e.g., after seeing an attack alert from the IDS, the analyst cannot know whether it is a true positive or a false positive).

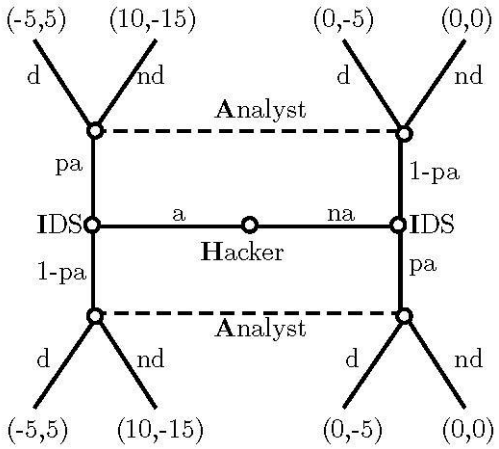

Fig. 2. A cybersecurity game with IDS recommendations

We found the Nash equilibria in the cybersecurity game described above for three cases of IDS accuracy: 10%, 50%, and 90%. The extensive form of the Cybersecurity game along with the Nash equilibria were generated by the Gambit software (McKelvey, McLennan, and Turocy, 2010). When the IDS is 50% accurate (pa = 0.5), then, according to the Nash equilibrium, the hacker's probability of attack actions equal 20%; and the analyst's probability of defend actions, regardless of whether the IDS says "attack" or "not attack", equals 67%. Similarly, when the IDS is 10% accurate, the probability of attack is 03% and the probability of defend is 09%, and when the IDS is 90% accurate, the probability of attack is 03% and the probability of defend is 09%. As per the Nash equilibrium predictions, the expected payoffs in a single trial for the hacker and the analyst were the following in different games: when IDS is 50% accurate, hacker's payoff = 0.0 and analyst's payoff = -3.0; when IDS is 10% accurate, hacker's payoff = 0.0 and analyst's payoff = -0.4; and, when IDS is 90% accurate, hacker's payoff = 0.0; analyst's payoff = -0.4. Although the above game theoretic analysis is determined by the static properties of the game, it is worth indicating that these theoretical implications still hold in the case of a dynamic setting (iterative play). In fact, both players following the above strategies in every trial still correspond to the Nash equilibrium in the large game made of N iterations (no one is better off deviating from it in any single trial of the finite game)[2].

3 Hypotheses

According to mixed-strategy equilibrium [7], when the hacker attacks according to the Nash prediction, the human analyst is indifferent between any of her actions. Similarly, when the analyst defends according to the Nash prediction, the human hacker is

[2] In the above cybersecurity game, playing the unique Nash equilibrium in mixed strategies of the stage game in every round leads to the unique subgame perfect Nash equilibrium in the repeated game.

indifferent between any of her actions. However, in a dynamic setting as in the current study, the above static predictions may not suffice to characterize human behavior. In addition, we may consider learning and dynamic decision making theories to hypothesize the human behavior [9] .

IBLT has been used for understanding and accounting for the interactions between hackers and analysts in simulated cybersecurity games [4, 9, 14]. According to IBLT, a human player would try to maximize her utility by selecting actions that are contingent upon other player's actions [24]. For example, in the baseline security game (Figure 1), the analyst would try to maximize utility by taking defend actions, which would increase the overall proportion of defend actions. In contrast, as an analyst increases defend actions, one expects a hacker to reduce attack actions (in order to minimize the disutility of facing a defend action when she attacks). Overall, according to IBLT, one expects to find that the proportion of attack and defend actions should agree with the mixed strategy Nash equilibria in the baseline security game, i.e., a higher proportion of defend actions and a smaller proportion of attack actions.

Next, in games where IDS is present, according to IBLT, the human proportion of attack and defend actions should be a function of the accuracy of the IDS. For example, when the IDS is 50% accurate, the analyst would not be able to use IDS informatively to maximize her payoffs. That is because, when she uses the IDS, it is likely to be inaccurate half the time and accurate only half the time. Due to an uninformative IDS and the goal to increase payoff, the analyst would likely increase her defend actions. That is because by defending there is a positive reward (+5) for the analyst, while there is a punishment (-15) for not defending. As per predictions from IBLT, we expect to find the proportion of attack and defend actions in the 50% accuracy case to be similar to the baseline condition.

When the IDS is 90% or 10% accurate, it is expected that the analyst would be able to rely on the IDS to make good decisions. If IDS is 90% accurate, then one expects to find it accurate on 90 out of 100 repeated trials on average and this number is likely sufficient to convince the analyst follow the IDS's predictions most of the time. If the analyst follows IDS's predictions and they are accurate, then she may not need to defend as much as she would when the IDS is only 50% accurate and uninformative. Thus, we expect to find lesser proportion of defend actions when the IDS is 90% accurate and, as IDS is accurate, the hacker should attack less (in order to avoid the disutility of facing a defend action when he attacks). Finally, when the IDS is only 10% accurate, the proportion of defend actions should be in between those when it is 90% accurate and when it is 50% accurate. That is because the analyst would tend to learn over repeated trials that most of the time IDS predictions are opposite to the ground truth and thus he needs to do opposite to what IDS says.

4 Methods

We conducted an online experiment involving human participants performing as hackers and analysts in a cybersecurity scenario using the cybersecurity game presented above. We manipulated the presence of the IDS recommendation as well as the IDS's accuracy (Figure 2), involving two human players, one performing as a hacker and the other performing as an analyst.

4.1 Experimental Design

Participants were randomly paired together and assigned to one of two between-subjects conditions over 50-trials: IDS-absent and IDS-present. In the IDS-absent condition, participants playing as hackers and analysts simultaneously decided their actions in each trial. For example, in each trial the hacker decided to either attack or not-attack the network and the analyst decided to defend or not defend the network. After both players made decisions, they were provided feedback that consisted of the following information: the actions taken by both players, their resulting payoffs, and a player's own cumulative payoff (opponent's cumulative payoff was not shown). The feedback was followed by a new trial where both participants made decisions again.

In the IDS-present condition, participants were randomly paired together and assigned to one of three between-subjects conditions over 50-trials: IDS 90% accurate (pa = 0.9), IDS 50% accurate or uninformative (pa = 0.5), and IDS 10% accurate (pa = 0.1). Figure 4 provides a snapshot of a trial in one of the IDS-present conditions. As shown in Figure 4, in the three IDS-present conditions, the hacker made a decision first, whether to attack or not-attack the network. The hacker's decision was followed by a report of the IDS to the analyst about whether the network event in the current trial was a cyber-threat or not (cyber-threats are generated by hackers and indicate an attack on the network). Based upon IDS's recommendation, the analyst decided to defend or not-defend the network. Once the analyst had made her decision, both players were provided the following feedback: actions taken and payoffs obtained for both players, IDS recommendation in the last trial, and a player's own cumulative payoff (opponent's cumulative payoff was not shown).

Overall, this design resulted in a total of 4 between-subjects conditions: IDS-absent and IDS-present with IDS 90% accurate, 50% accurate, and 10% accurate. The probability value *pa* as well as the end-point of the game were unknown to participants and remained constant across all trials of the game. In order to test our hypotheses, we analyzed the proportion of attack and defend actions averaged over all participants and trials.

Overall, this design resulted in a total of 4 between-subjects conditions: IDS-absent and IDS-present with IDS 90% accurate, 50% accurate, and 10% accurate. The probability value *pa* as well as the end-point of the game were unknown to participants and remained constant across all trials of the game. In order to test our hypotheses, we analyzed the proportion of attack and defend actions averaged over all participants and trials and the proportion of attack and defend actions per trial averaged over all participants. Also, we analyzed the hacker and analyst's expected payoffs per trial averaged over all participants and trials as well as the proportion of times the analyst matched her actions to IDS's recommendations.

4.2 Participants

A total of 200 American respondents were recruited through Amazon Mechanical Turk to voluntarily participate in an online cyber-security study. Among the 200 respondents, the online system randomly paired individuals across the four between-subjects conditions and across one of the two roles, hacker or analyst, in a condition: IDS-absent (N = 20 pairs) and IDS-present (10% accurate: N = 25 pairs; 50% accu-

rate: N = 29 pairs; and, 90% accurate: N = 26 pairs). Forty-five percent of participants were males. Age ranged from 19 years to 63 years (Mean = 33; SD = 10). About 35.5% of participants self reported to possess a 4-year undergraduate college degree; 54.5% reported to either have high-school degrees, 2-year college degrees, or some college experience; and, 10% reported to either have a graduate or a professional degree. Hacker participants started with $3 in their account and analyst participants started with $5 in their account. These unbalanced starting payoff were kept based upon a pilot study, where it was found that starting with these payoffs both players win about the same amount from the game at the end of their experiment. As can be seen in Figure 1, like in the real-world, the payoff matrix would give advantage to the hacker over the analyst. Upon completion of the experiment, each person was paid based on the points obtained from the game. For this purpose, points in the game were converted to real money in the ratio: 1 point in the game = 1 cent in real money. The average time spent to complete the task was 29 minutes, and the average amount of total payment was $3.06 with both hackers and analysts getting close to this amount across different conditions.

4.3 Procedure

Participants were instructed to repeatedly play a game with their opponent where each trial would determine either a gain or a loss of points from their initial capital. Participants were informed through instructions that the immediate outcome of each trial depended on their own choice between two options available, as well as the choice made by their opponent. Participants were not provided with any information regarding the payoff matrix prior to starting the game and they only learnt the payoffs by the feedback provided after each trial of the game about the actions and payoffs obtained by both players. As shown in Figure 4, in the conditions involving the IDS, the recommendation made by IDS to the analyst was also reported along with its accuracy (whether it correctly predicted the hacker's move) after each trial.

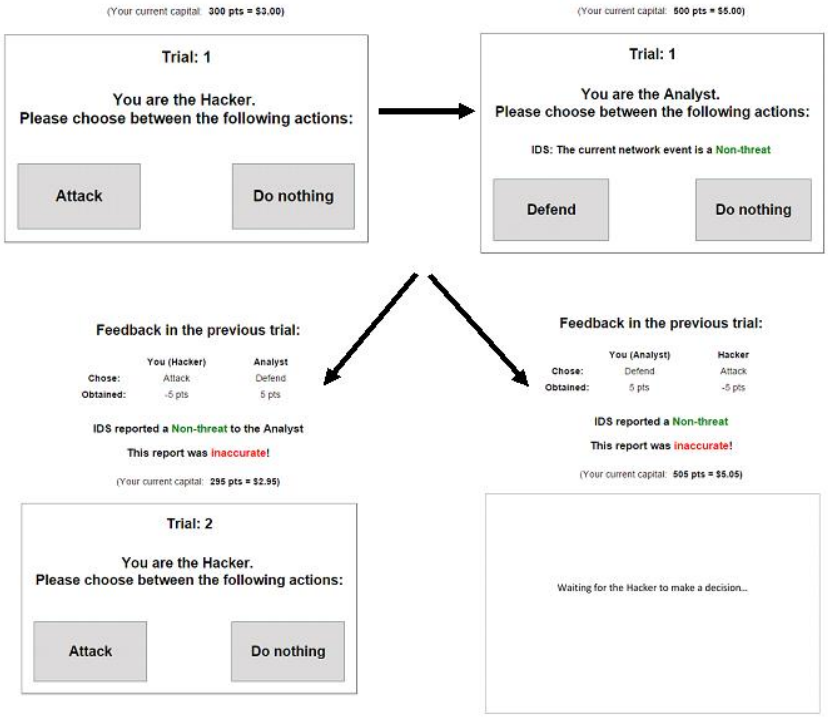

Fig. 3. The dynamics of a trial in one of the IDS-present conditions in the Cybersecurity game. The hacker acts first and her action is followed by an IDS recommendation to the analyst (green color was used for non-threats and red for threat events). The IDS recommendation was followed by an action from the analyst. Finally, both players were shown feedback of the previous trial and the hacker was asked to act first.

5 Results

We conducted analyses on the overall average behavior as well as on the behavior over trials for each of the roles to which the participants were assigned, the hacker or the analyst. When appropriate, we also compared human behavior to the Nash equilibria solutions and predictions from IBLT described above.

5.1 Proportion of Attack and Defend Actions

Fig. 4 shows the average proportion of attack and defend actions in the different conditions (the dotted lines show the Nash proportions). In the absence of the IDS, the overall proportion of actions indicated agreement with the optimal Nash solutions. The proportion of defend actions (0.63) was no different than the Nash defend proportion (0.67) ($t(19) = -0.70$, *ns*); and, as expected the proportion of attack actions (0.25)

was no different than the Nash attack proportion (0.20) ($t(19) = 1.66$, *ns*). Thus, as expected, on average, in the absence of the IDS, human behavior agreed with the Nash equilibrium solutions.

In the presence of the IDS, however, the proportion of actions seemed to directly depend on how accurate the IDS was, particularly for analyst's behavior. When the IDS's accuracy was 50%, humans in the hacker and analyst roles practically ignored the IDS recommendations and behaved in a similar way as when the IDS was absent (Hacker: Absent vs Present for 50% accuracy: $t(47) = -1.41$, *ns*; Analyst: Absent vs Present for 50% accuracy: $t(47) = 0.51$, *ns*). Furthermore, when the IDS was 50% accurate, analysts behaved in agreement with their Nash proportion of actions ($t(28) = -1.68$, *ns*). However, hackers deviated from their Nash proportion of actions: The proportion of attack actions (0.32) was significantly higher than the Nash proportion (0.20) ($t(28) = 3.49$, $p < .01$). Overall, the behavior of hackers and analysts in the presence of an uninformative (i.e., 50% accurate) IDS is similar to the condition in which the IDS is absent and these results are in agreement with our expectations based upon IBLT.

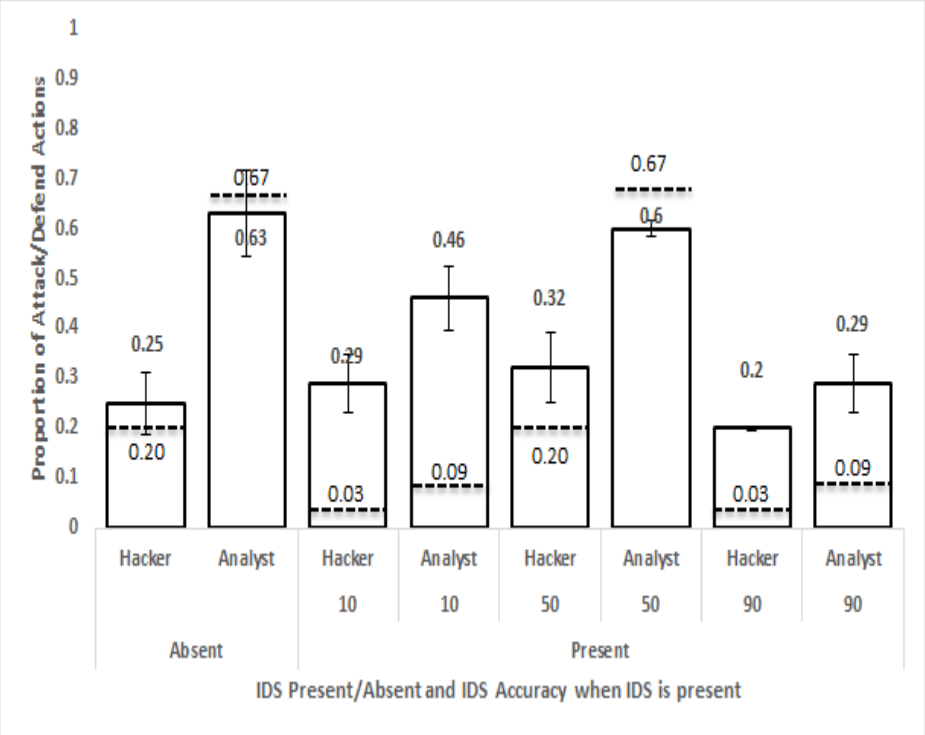

Fig. 4. The average proportion of attack or defend actions as a result of IDS being absent or present as well as for different IDS accuracies (10%, 50%, and 90%) when IDS is present.

When the IDS was informative ($pa \neq 0.5$), the proportion of defend actions was lower (0.46 for 10% accurate IDS and 0.29 for 90% accurate IDS) compared to the

condition in which the IDS was absent (0.63 for IDS absent; Absent vs Present for 10% accuracy: $t(43) = 3.16$, $p < .01$; Absent vs Present for 90% accuracy: $t(44) = 5.66$, $p < .001$). Furthermore, the proportion of defend actions was higher than the Nash optimal (10% accuracy: $0.46 > 0.09$, $t(24) = 11.13$, $p < .001$; 90% accuracy: $0.29 > 0.09$, $t(25) = 5.21$, $p < .001$). Next, the proportion of defend actions for 90% level of accuracy was significantly lower (0.29) than those for the 10% level of IDS accuracy (0.46) (10% accuracy: $t(49) = -3.21$, $p < .01$).

For an informative IDS, hackers tend to attack at about the same rate (0.29 for 10% accurate and 0.20 for 90% accurate) as when the IDS was absent (0.25 for IDS absent; Absent vs Present for 10% accuracy: $t(43) = -0.87$, *ns*; Absent vs Present for 90% accuracy: $t(44) = -1.22$, *ns*); however, this proportion of attack action was significantly higher than the Nash optimal (10% accuracy: $0.29 > 0.03$; $t(24) = 8.83$, $p < .001$; 90% accuracy: $0.20 > 0.03$; $t(25) = 5.81$, $p < .001$). Furthermore, the proportion of attack actions was significantly higher when IDS accuracy was 10% compared to 90% ($0.29 > 0.20$; $t(49) = -2.17$, $p < .01$). Overall, these results are in agreement with our expectations in the case of an informative IDS (i.e., when it is 10% and 90% accurate).

6 Discussion and Conclusions

Using an online experiment with human participants interacting in a cybersecurity game, we showed how the presence and accuracy of decision support influences the proportion of attack and defend actions. We found that in the absence of IDS, both hackers and analysts performed optimally, as per Nash solutions. However, the presence of IDS biases participants' decisions. As per IBL models of cybersecurity [9] participants in the analyst role take more defend actions to maximize utility. Hackers also try to reduce their proportion of attack actions in order to reduce the disutility of facing a defend action while attacking the network, converging to Nash proportions in the absence of IDS.

Interestingly, when the IDS was present but uninformative (i.e., 50% accurate), analysts practically ignored the IDS recommendations and performed again optimally as in the case where the IDS was absent, while hackers attacked more than optimal. A high level of defend actions assures the analyst to maximize her payoff. That is, because defending an attack resulted in a positive reward (+5) and not defending an attack led to a large punishment (-15), the overall learning effect is to maintain the level of defend actions. As a result of high level of defend actions, the hacker reduced the attack actions (in order to minimize the disutility by facing defend actions from the analyst). This suggests that informing an attacker of uninformative IDS may benefit cybersecurity.

In the presence of an informative IDS (accuracy = 10% or 90%), participants performing as analysts reduced their proportion of defend actions compared to when the IDS was absent, but unfortunately not enough to behave optimally as predicted by the Nash equilibrium. The hackers kept a similar level of attack actions, which was also higher than optimal. Therefore, the presence of informative IDS might help in cybersecurity by keeping the analyst calmer and become more effective by reducing the number of defend actions.

There are a number of implications of our results from the real word, in particular on the use of IDS technology for cyber-attack detection. First, if IDS technology is going to be uninformative (or its accuracy is close to 50%), then it is as good for organizations for not investing in this costly technology. However, if the IDS technology is very accurate or very inaccurate, then, based upon repeated feedback, this technology will become advantageous for real-world analysts to make good decisions to catch cyber-attacks from occurring. Also, making hackers know about an accurate IDS over repeated attacks, provides a good deterrence from attacking such networks. That is because we found that the feedback about IDS's accuracy after a trial in our experiment caused hackers to reduce their proportion of attack actions.

7 Acknowledgements

This research was partially supported by the Department of Science & Technology (Cognitive Science Research Initiative, Award number: SR/CSRI/28/2013(G)) award to Varun Dutt. The authors thank Zahid Maqbool, Akash Rao, and Palvi Aggarwal, Applied Cognitive Science Laboratory, Indian Institute of Technology Mandi, for their help with editing this manuscript.

8 References

1. Alpcan, T., & Başar, T. (2010). *Network security: A decision and game-theoretic approach*. Cambridge University Press.
2. Anderson, J. R. A. L., & Lebiere, C. (1998). The atomic components of thought Lawrence Erlbaum. *Mathway, NJ*.
3. Anderson, J. R., & Lebiere, C. (2003). The Newell test for a theory of cognition. *Behavioral and brain Sciences*, *26*(05), 587-601.
4. Arora, A., & Dutt, V. (2013). Cyber Security: Evaluating the Effects of Attack Strategy and Base Rate through InstanceBased Learning. In *12th International Conference on Cognitive Modeling. Ottawa, Canada*.
5. Bhatt, C., Koshti, A., Agrawal, H., Malek, Z., & Trivedi, B. (2011). Architecture for intrusion detection system with fault tolerance using mobile agent. *International Journal of Network Security & Its Applications*, *3*(5), 167.
6. Bloem, M., Alpcan, T., & Başar, T. (2006, December). Intrusion response as a resource allocation problem. In *Decision and Control, 2006 45th IEEE Conference on* (pp. 6283-6288). IEEE.
7. Camerer, C. (2003). *Behavioral game theory: Experiments in strategic interaction*. Princeton University Press.
8. Cyber Ark. (2014). CyberArk 2014 Global Advanced Threat Landscape. Retrieved from http://www.dit.co.jp/ditplus/report/201408/pdf/survey-cyberarks-2014-global-advanced-threat-landscape%20survey-07-28-14.pdf.

9. Dutt, V., Ahn, Y. S., & Gonzalez, C. (2013). Cyber situation awareness modeling detection of cyber attacks with instance-based learning theory.*Human Factors: The Journal of the Human Factors and Ergonomics Society*,*55*(3), 605-618.
10. Forbes. (2015). Cybersecurity Trends To Watch In 2015. Retrieved from http://www.forbes.com/sites/riskmap/2015/02/05/cybersecurity-trends-to-watch-in-2015/#612d8b685da3
11. Gonzalez, C., Lerch, J. F., & Lebiere, C. (2003). Instance-based learning in dynamic decision making. *Cognitive Science*, *27*(4), 591-635.
12. Jajodia, S., Liu, P., Swarup, V., & Wang, C. (2010). *Cyber situational awareness* (Vol. 14). New York, NY: Springer.
13. Laszka, A., Abbas, W., Sastry, S. S., Vorobeychik, Y., & Koutsoukos, X. Optimal Thresholds for Intrusion Detection Systems.
14. Dutt, V., & Kaur, A. (2013). Cyber security: testing the effects of attack strategy, similarity, and experience on cyber attack detection. *International Journal of Trust Management in Computing and Communications*, *1*(3-4), 261-273.
15. Maqbool, Z., Makhijani, N., Pammi, C., & Dutt, V. (2016). Effects of motivation: rewarding analysts for good cyber-attack detection may not be the best strategy. Manuscript submitted for publication
16. McAfee. (2016). In the Dark: Crucial Industries Confront Cyberattacks. Retrieved from http://www.mcafee.com/us/resources/reports/rp-critical-infrastructure-protection.pdf
17. McKelvey, R. D., McLennan, A. M., & Turocy, T. L. (2006). Gambit: Software tools for game theory.
18. Rabadia, P., & Valli, C. analysis into developing an accurate and efficient intrusion detection approach.
19. Rajasekaran, K., & Nirmala, K. (2012). Classification and Importance of Intrusion Detection System. *International Journal of Computer Science and Information Security*, *10*(8), 44.
20. Roy, S., Ellis, C., Shiva, S., Dasgupta, D., Shandilya, V., & Wu, Q. (2010, January). A survey of game theory as applied to network security. In *System Sciences (HICSS), 2010 43rd Hawaii International Conference on* (pp. 1-10). IEEE.
21. Spathoulas, G. P., & Katsikas, S. K. (2010). Reducing false positives in intrusion detection systems. *computers & security*, *29*(1), 35-44.
22. White House. (2011). Remarks by the President on securing our nation's cyber infrastructure. Retrieved from http://www.whitehouse.gov/the_press_office/Remarks-by-the-President-on-Securing-Our-Nations-Cyber-Infrastructure/
23. Gonzalez, C., Ben-Asher, N., Oltramari, A., & Lebiere, C. (in press). Cognitive models of cyber situation awareness and decision making. In C. Wang, A. Kott, & R. Erbacher (Eds.), *Cyber defense and situational awareness*.
24. Gonzalez, C., Ben-Asher, N., Martin, J. M. and Dutt, V. (2015), A Cognitive Model of Dynamic Cooperation With Varied Interdependency Information. Cognitive Science, 39: 457–495. doi: 10.1111/cogs.12170