Effects of Motivation: Rewarding Hackers for Undetected Attacks Cause Analysts to Perform Poorly

Zahid Maqbool, **Nidhi Makhijani**, Indian Institute of Technology, Mandi, India, **V. S. Chandrasekhar Pammi**, University of Allahabad, India, and **Varun Dutt**, Indian Institute of Technology, Mandi, India

Objective: The aim of this study was to determine how monetary motivations influence decision making of humans performing as security analysts and hackers in a cybersecurity game.

Background: Cyberattacks are increasing at an alarming rate. As cyberattacks often cause damage to existing cyber infrastructures, it is important to understand how monetary rewards may influence decision making of hackers and analysts in the cyber world. Currently, only limited attention has been given to this area.

Method: In an experiment, participants were randomly assigned to three between-subjects conditions (n = 26 for each condition): equal payoff, where the magnitude of monetary rewards for hackers and defenders was the same; rewarding hacker, where the magnitude of monetary reward for hacker's successful attack was 10 times the reward for analyst's successful defense; and rewarding analyst, where the magnitude of monetary reward for analyst's successful defense was 10 times the reward for hacker's successful attack. In all conditions, half of the participants were human hackers playing against Nash analysts and half were human analysts playing against Nash hackers.

Results: Results revealed that monetary rewards for human hackers and analysts caused a decrease in attack and defend actions compared with the baseline. Furthermore, rewarding human hackers for undetected attacks made analysts deviate significantly from their optimal behavior.

Conclusions: If hackers are rewarded for their undetected attack actions, then this causes analysts to deviate from optimal defend proportions. Thus, analysts need to be trained not become overenthusiastic in defending networks.

Application: Applications of our results are to networks where the influence of monetary rewards may cause information theft and system damage.

Keywords: motivation, cybersecurity, decision making, behavioral game theory, instance-based learning theory

Address correspondence to Dr. Varun Dutt, Applied Cognitive Science Laboratory, Indian Institute of Technology, Mandi, Himachal Pradesh 175005, India; e-mail: varun@iitmandi.ac.in.

HUMAN FACTORS
Vol. 59, No. 3, May 2017, pp. 420–431
DOI: 10.1177/0018720816681888

INTRODUCTION

Cyberattacks, that is, attempts by hackers to damage computer systems, are increasing at an alarming rate (Tobias, 2015). A survey by Cyber Ark (2014) suggests that nations are at a greater risk from cyberattacks compared with physical attacks. According to U.S. president Barack Obama, cyber threat is one of the most serious economic and national security challenges for a nation to address (White House, 2009). Given the increasing trend of cyber warfare, it is necessary to safeguard networks against cyber threats. The job of safeguarding networks against cyber threats is that of security analysts (referred to as "analysts" hereafter), people who are tasked with the responsibility of defending networks against hackers. However, currently little is known on how analysts may make decisions in situations involving cyberattacks on networks (Dutt, Ahn, & Gonzalez, 2011).

We need to investigate how certain human factors (e.g., monetary motivations) may influence decision making of hackers and analysts in the cyber world (Arora & Dutt, 2013; Dutt, Ahn, & Gonzalez, 2013; Kaur & Dutt, 2013). For example, an analyst would likely have different benefits and costs for successfully detecting or failing to detect cyberattacks, respectively. On one hand, an analyst could earn large benefits in terms of career promotions, financial incentives, and recognition for successfully detecting cyberattacks. However, the same analyst is likely to also face large costs due to his or her committing false alarms and misses in detecting cyberattacks. Such costs could be in terms of demotions, losing one's job, and paying for damages incurred. Likewise, the hacker could experience benefits and costs for launching cyberattacks. If the hacker is caught attacking a network, then he or she is likely to face fines and jail. However, if he or she manages to initiate an undetected

attack, then benefits could include financial gains, fame, and notoriety. An analyst's false alarms would cost nothing to a hacker in cases where the hacker decides not to attack the network.

Although research that investigates the role of benefits and costs in influencing actions of human hackers and analysts is much needed, to date, very little attention has been given to this area in the cyber world (Roy et al., 2010). Mostly, using game theory, research has investigated the best response strategies in security games between hackers and analysts (Alpcan & Başar, 2010). However, these approaches have restricted their analyses to finding mathematical Nash equilibriums without testing whether such equilibriums actually hold in the real world. Without a clear understanding of the dynamics of human behavior of both analysts and hackers, any tools created to predict hacker actions and to support analyst actions are expected to be futile. Hence, it is important to study how monetary motivations affect the behavior of hackers and analysts in the cyber world.

One way to study the role of motivations in cybersecurity is by using behavioral game theory (Camerer, 2003; Dutt et al., 2013), which conceptualizes the interaction between hackers and analysts using simple 2 (players) × 2 (actions per player) repeated games. In these games, players make repeated decisions from a set of decision actions available to them and receive payoffs for their actions. The payoff obtained by both players is a function of both players' actions. In this paper, we use behavioral game theory to study the role of motivations in influencing decisions in a cybersecurity game. Here, hacker's action is abstracted as a choice between an "attack" action and a "not-attack" action. An attack action means attacking a network using cyber threats, whereas a not-attack action means not attacking the system. Similarly, an analyst's action is abstracted as a choice between a "defend" and a "not-defend" action in the network. Motivations for hackers and analysts are conceptualized in terms of payoffs for both players (Alpcan & Başar, 2010; Roy et al., 2010).

These payoffs result due to a combination of attack and defend actions from hackers and analysts individually, and they also determine the mathematical Nash (optimal) solutions in the game (Camerer, 2003). In this paper, human participants performing as hackers or analysts play against their Nash counterparts, that is, play against computer programs that perform according to predefined Nash strategies computed based upon payoffs. By analyzing these human–Nash games, we are able to investigate deviations from optimal proportions that humans show against their optimal Nash counterparts. Furthermore, we investigate how these deviations are influenced by different benefits and costs to hackers and analysts.

Literature has revealed instance-based learning theory (IBLT; Dutt et al., 2013; Dutt & Gonzalez, 2012; Gonzalez, Lerch, & Lebiere, 2003) to be an accurate account of decision making when humans perform as hackers and analysts (Arora & Dutt, 2013; Dutt et al., 2013; Kaur & Dutt, 2013). IBLT uses a cognitive information representation in memory called an instance. An instance consists of three parts: a situation (a set of attributes that define the decision situation), a decision or action, and an outcome or utility of the decision made in the situation. The different parts of an instance are built through a general decision process with three steps: (1) characterizing a situation by its attributes, (2) making a decision based on expected utility, and (3) updating the utility when feedback is available. Instances are accumulated over time in memory. They are retrieved, blended, and used repeatedly in future decisions. Their strength in memory, called activation, is reinforced according to statistical procedures proposed as part of the ACT-R cognitive architecture (Anderson & Lebiere, 1998; Lebiere, 1999). In such situations, hackers and analysts, playing a game against each other, possess cognitive limitations and rely on recency and frequency of available information to make decisions. Therefore, application of IBLT to experiential decisions of hackers and analysts will help explain how these decisions are affected by motivational factors.

Furthermore, research shows that prospect theory (PT; Kahneman & Tversky, 1979; Tversky & Kahneman, 1992) provides a robust account of decisions in situations involving gains and losses. According to PT, people value gains and losses differently, and losses have more emotional impact than an equivalent amount of gains.

		Analyst	
		Defend (*d*)	Not defend (*nd*)
Hacker	Attack (*a*)	A(a, d), -D(a, d)	-A(a, nd), D(a, nd)
	Not Attack (*na*)	0, D(na, d)	0, 0

Figure 1. Actions and payoffs in the cybersecurity game between hackers and analysts. The payoffs represent costs to players, and negative costs are benefits. In each cell, the first payoff value corresponds to the hacker and the second value corresponds to the analyst.

For example, the amount of utility gained from receiving $50 is less in magnitude to the utility when one gains $100 and then loses $50 (both situations being mathematically identical). Given that motivational factors involve gaining and losing payoffs, PT is likely to help explain experiential decisions made by hackers and analysts. In the next section, we detail the mechanics of a game that we used to investigate the role of motivations in influencing decision making in simulated cyberattack situations.

THE CYBERSECURITY GAME

The cybersecurity game (see Figure 1) is conceptualized as a 2 × 2 game (Alpcan & Başar, 2010), that is, between two players, where each player has two actions to repeatedly choose between (for the hacker, the actions include attack [*a*] and not attack [*na*]; whereas for the analyst, the actions include defend [*d*] and not defend [*nd*]). One of the players is randomly assigned to be the hacker, and the other player is assigned to be the analyst. As shown in Figure 1, there will be benefits for the hacker (–A[*a, nd*]) due to an undetected attack and costs (A[*a, d*]) due to a detected attack. Similarly, there will be benefits for the analyst (–D[*a, d*]) due to a successful defend action against an attack and costs (D[*na, d*] and D[*a, nd*]) due to unsuccessful defend and not-defend actions against the hacker's not-attack and attack actions, respectively. Here A denotes the payoff obtained by the hacker and D refers to the payoff obtained by the analyst. Analysts are penalized twice compared with hackers in such a game, because in the real world, analysts suffer costs when there is a successful attack on the network and when analysts spend time and resources defending against hackers who do not attack. The objective for both the hacker and analyst is to maximize payoffs when the game is played repeatedly over several rounds.

A pure-strategy Nash equilibrium is an action profile with the property that no single player can obtain a higher payoff by deviating unilaterally from this profile. In some situations, there may not exist a pure-strategy Nash equilibrium, and players may want to randomize their choices over available actions as they play repeatedly. The latter case of random mixing of choices is called a mixed-strategy Nash equilibrium. As participants performing as hackers and analysts might want to choose any of the two actions available to them, there exists a mixed-strategy Nash equilibrium in this game. Let $0 \leq p \leq 1$ and $1 - p$ be the proportions for the hacker choosing *a* and *na* actions, respectively. Also, let $0 \leq q \leq 1$ and $1 - q$ be the proportions where the analyst chooses *d* and *nd* actions. Then, the mixed-strategy Nash equilibrium is given as

$$p = \frac{\mathrm{D}(na, d)}{\mathrm{D}(na, d) + \mathrm{D}(a, d) + \mathrm{D}(a, nd)} \quad (1)$$

$$q = \frac{\mathrm{A}(a, nd)}{\mathrm{A}(a, nd) + \mathrm{A}(a, d)} \quad (2)$$

Therefore, payoffs determine the Nash expectations for attack and defend actions. In this paper, we observe deviations from the optimal *p* and *q* proportions for the human players.

Effects of Motivations in Cybersecurity Game

Using the cybersecurity game, we manipulate payoffs for hackers and analysts in three different conditions: equal payoff, rewarding hacker, and rewarding analyst (see Figures 2a, 2b, and 2c). In equal-payoff condition (baseline), analysts are awarded +5 points for an *a-d* action combination, and hackers are awarded an equal number of points for an *a-nd* action combination. In rewarding-hacker condition, the hacker is awarded +50 points for an *a-nd* action combination, which is 10 times the award in equal-payoff condition. Thus, hackers are motivated to wage (undetected) attacks on the network. Furthermore, in rewarding-analyst condition, the analyst is awarded +50 points for an *a-d*

Equal - Payoff:

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

NE: p = 33%, q = 33%

Rewarding - Hacker:

		Analyst: Defend (*d*)	Analyst: Not defend (*nd*)
Hacker	Attack (*a*)	-10, 5	50, -5
	Not Attack (*na*)	0, -5	0, 0

NE: p = 33%, q = 83%

Rewarding - Analyst:

		Analyst: Defend (*d*)	Analyst: Not defend (*nd*)
Hacker	Attack (*a*)	-10, 50	5, -5
	Not Attack (*na*)	0, -5	0, 0

NE: p = 8%, q = 33%

Figure 2. Monetary payoffs and proportion of attack (*p*) and defend (*q*) actions for hackers and analysts in three different conditions: (a) equal payoff, (b) rewarding hacker, and (c) rewarding analyst. The *p* and *q* percentages define the Nash equilibriums in different conditions.

action combination, which is again 10 times the award in equal-payoff condition. Therefore, in rewarding-analyst condition, analysts are motivated to correctly detect cyberattacks waged on the network. In this paper, human players, whether hackers or analysts, are specifically asked to maximize their payoffs while they play against their Nash opponents. In this situation, the proportion of actions for human players can be hypothesized according to IBLT (Dutt et al., 2013). Furthermore, research has shown that PT provides an accurate account of situations involving gains and losses (Kahneman & Tversky, 1979; Tversky & Kahneman, 1992). The theory predicts that losses have more emotional impact than an equivalent amount of gains. Thus, according to PT, hackers and analysts would tend to perform those actions where the negative emotional impact (due to a loss) is less.

According to IBLT, a human player would try to maximize his or her outcomes from his or her actions in the game. Thus, as per IBLT, human players would tend to reinforce those actions that result in rewards compared with the baseline

payoffs. As hackers are rewarded for their attack actions in the rewarding-hacker condition, these rewards would likely increase the proportion of attack actions compared with the baseline condition. However, in the rewarding-hacker condition, the Nash analysts also perform a large number of defend actions (~83%). As predicted by PT, human hackers would minimize the loss caused due to an analyst's defend actions. As losses loom larger than gains (due to PT), overall, we expect a decrease in the proportion of attack actions in the rewarding-hacker condition.

Furthermore, as per IBLT, we expect a greater proportion of defend actions for human analysts in the rewarding-analyst condition compared with the equal-payoff condition. However, a Nash hacker attacks the network very little in the rewarding-analyst condition. Thus, an excessive proportion of defend actions would make the analyst lose points. As losses are felt more painful compared with gains, according to PT, human analysts would tend to be cautious while defending the network. Overall, this dynamic would lead human analysts to decrease their proportion of defend actions in the rewarding-analyst condition compared with the baseline condition.

To test these expectations, we detail an experiment where human hackers and analysts played against their Nash counterparts. We validate our expectations of the proportions of attack and defend actions, whether in agreement or disagreement with the optimal proportions, based upon predictions from IBLT and PT.

EXPERIMENT

Experimental Design

One hundred and fifty-six participants were randomly assigned to one of three between-subjects conditions: equal payoff (n = 52), rewarding hacker (n = 52), and rewarding analyst (n = 52) (see Figure 2 for matrices used in these conditions). In each condition, 26 participants performed as hackers against Nash analysts, and 26 participants performed as analysts against Nash hackers. Each condition was 50 rounds long and involved an interaction between hackers and analysts in real time. Nash players, whether hackers or analysts, played according to the computed Nash proportion in the respective condition. In each round, a uniformly distributed random number R ($\sim U[0, 1]$) was generated and compared with the Nash proportion. If $R \leq$ Nash proportion, then the corresponding Nash action was implemented. For example, in the equal-payoff condition, if $R \leq 0.33$ (Nash proportion), then the Nash hacker played an attack action; otherwise, he played a not-attack action. For testing our expectations, we compared the average proportion of attack and defend actions from human players across different conditions (average was computed over all human participants and rounds in a particular condition). Also, we compared the average proportion of attack and defend actions across 50 rounds of play in a condition (average was computed over all human participants in a particular round).

Participants

The cybersecurity game was deployed on the Prolific crowd-sourcing platform (Prolific, 2016). Sixty-nine percent of participants were males. Ages ranged from 18 years to 54 years (M = 28 years, SD = 8 years). Information consents were obtained from each participant, and Indian Institute of Technology Mandi approved the study. Participants were from different education levels: 20% possessed high school degrees, 10% possessed senior secondary degrees, 52% possessed college degrees, 14% possessed master's degrees, 2% possessed PhD degrees, and 2% possessed other professional degrees. Furthermore, participants possessed different educational backgrounds: 52% were from science, technology, engineering, and mathematics (STEM) backgrounds, and the remaining were from humanities and social sciences backgrounds. Participants reported pursuing different occupations: 34% were students, 16% worked in the information technology (IT) industry, 18% worked in engineering, 17% worked in business and consultancy, and 15% worked in manufacturing. Participants were asked to maximize their payoffs and were compensated a flat participation fee of US$1.50 at the end of their study. No participant took more than 20 min to finish the study.

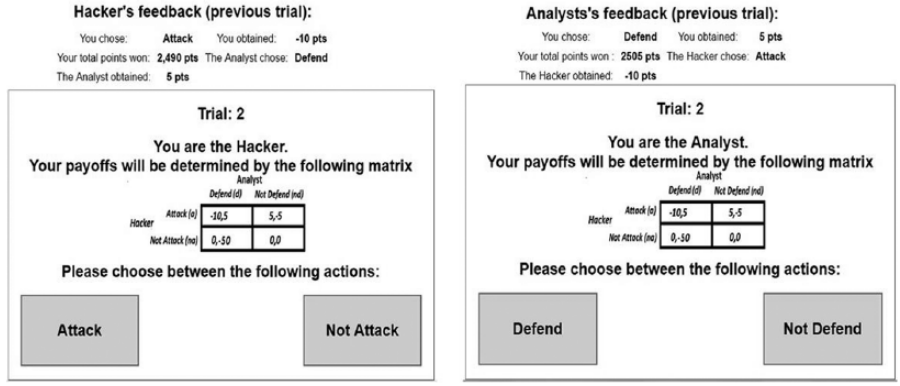

Figure 3. The user interfaces for hackers and analysts in the cybersecurity game. The figure depicts the set of actions available and feedback provided to both players.

Procedure

A total of 156 participants were given instructions about their goal in the cybersecurity game, and they possessed complete information about their own and their opponent's actions and payoffs in all conditions (the payoff matrix was given). Participants could gain or lose points as the game continued, and the game's endpoint was not disclosed at any point during the experiment. As part of the instructions, payoff matrices as well as the set of actions possible were explained to both players (see the appendix). In a round, both participants decided their actions simultaneously and then received feedback (see Figure 3) about each other's actions and payoffs. After feedback, participants were asked to make the next trial's decision. Once the study ended, participants were thanked and given their participation fee.

RESULTS

Proportion of Attack and Defend Actions Across Conditions

Figure 4 shows the average proportion of attack and defend actions for human hackers and analysts in the three conditions. The average proportion of attack and defend actions varied across conditions: attack actions, $F(2, 75) = 5.67$, $p < .01$, $w^2 = .13$; defend actions, $F(2, 75) = 27.05$, $p < .001$, $w^2 = .42$. Post hoc REGWQ tests revealed that the proportion of attack actions was lower in the rewarding-hacker condition compared with both equal-payoff ($p < .05$) and rewarding-analyst ($p < .05$) conditions.

There was no difference between the proportion of attack actions in the equal-payoff and rewarding-analyst conditions ($p > .05$). Thus, as per our expectations from IBLT and PT, overall, human hackers decreased their proportion of attacks when rewarded for successfully attacking the system.

Furthermore, post hoc REGWQ tests revealed that the proportion of defend actions was lower in the rewarding-analyst condition compared with the equal-payoff condition ($p < .001$); however, there was no difference in the proportion of defend actions between the rewarding-hacker and equal-payoff conditions ($p > .05$). In summary, as per expectations from IBLT and PT, analysts decreased defend actions when rewarded for successfully defending the system.

Optimal Proportion of Attack and Defend Actions Across Conditions

We performed a simulation of 2,000 analyst-hacker pairs each in different conditions by assuming one player playing the mixed-strategy Nash equilibrium and the other player playing only one of his or her possible actions across all 50 rounds. For example, if the hacker played according to his or her mixed-strategy equilibrium in the equal-payoff condition, then we simulated 2,000 analysts performing only defend action or only not-defend action across 50 rounds. After running such simulations, we recorded the cumulative payoff of the human player for each set of his or her actions. The action that maximized the cumulative payoff was termed the optimal action. Figure 5 shows

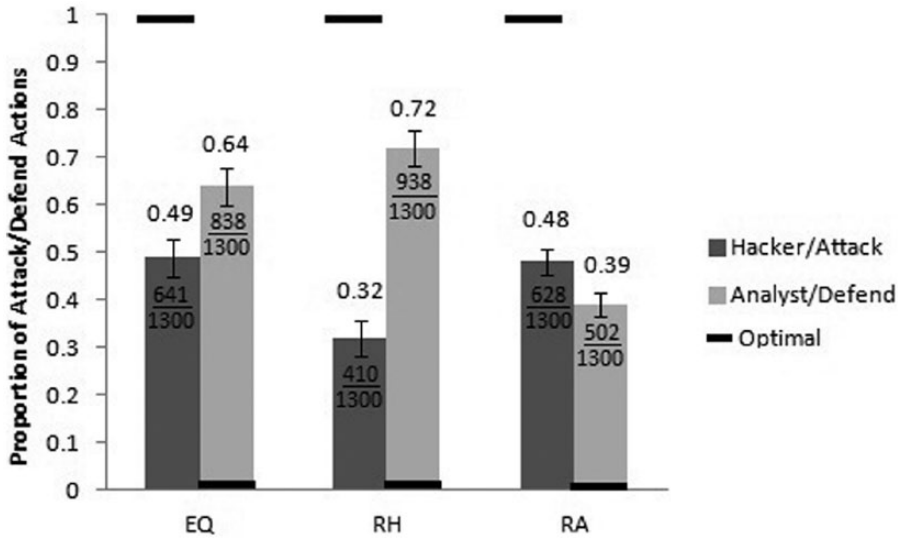

Figure 4. Average proportion of attack and defend actions from human hackers and analysts across the three conditions: equal payoff (EQ), rewarding hacker (RH), and rewarding analyst (RA). The values on the bars depict frequencies against the computed proportions (e.g., there were 641 participant-trial attack actions out of a total of 1,300 participant-trial data points in the EQ condition). The horizontal bars show the corresponding optimal proportions. The error bars show 95% confidence intervals around the mean value.

the cumulative payoffs across different conditions for both players. In all three conditions, the attack action and the not-defend action were those that maximized hacker's and analyst's payoffs, respectively.

These optimal actions for hackers and analysts have been shown in Figure 4 as horizontal levels at 1.0 or 0.0. As seen in Figure 4, across all conditions, the proportions of attack and defend actions deviated significantly from their optimal levels—equal payoff: analyst, $t(25) = 17.84$, $p < .001$, $r = .96$; hacker, $t(25) = -14.25$, $p < .001$, $r = .94$; rewarding hacker: analyst, $t(25) = 25.81$, $p < .001$, $r = .98$; hacker, $t(25) = -17.00$, $p < .001$, $r = .96$; rewarding analyst: analyst, $t(25) = 10.56$, $p < .001$, $r = .90$; hacker, $t(25) = -10.58$, $p < .001$, $r = .90$. In summary, human actions deviated from their optimal behaviors in all conditions.

Proportion of Attack and Defend Actions Across Rounds

As the proportion of attack and defend actions deviated significantly from their optimal proportions (Figure 4), we analyzed the trend of these actions across rounds. Figures 6a and 6b show the average proportion of attack actions across rounds in the rewarding-hacker and rewarding-analyst conditions compared with that in the equal-payoff condition, respectively (optimal proportions are shown as horizontal lines). Although the attack proportions appeared

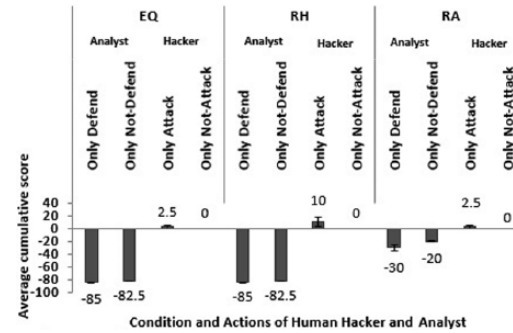

Figure 5. The average cumulative scores for analyst's and hacker's actions in different conditions: equal payoff (EQ), rewarding hacker (RH), and rewarding analyst (RA).

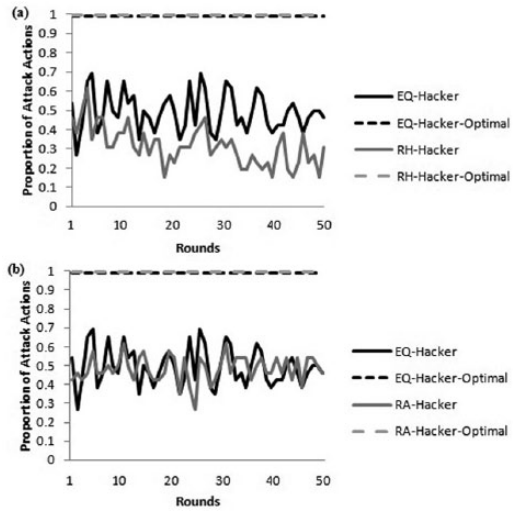

Figure 6. Average proportion of attack actions across rounds in equal-payoff (EQ) and rewarding-hacker (RH) conditions (a) and EQ and rewarding-analyst (RA) conditions (b). Horizontal lines show the corresponding optimal proportions.

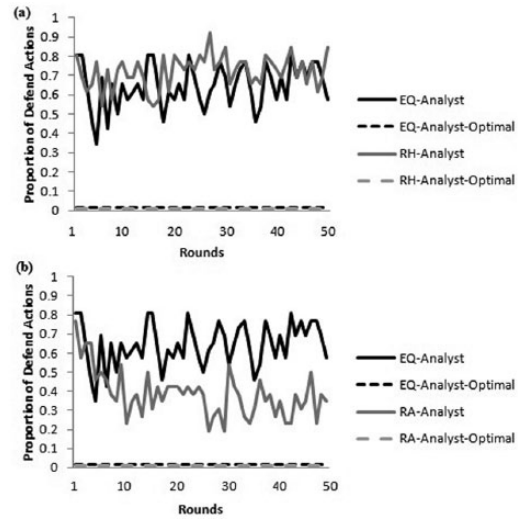

Figure 7. Average proportion of defend actions across rounds in equal-payoff (EQ) and rewarding-hacker (RH) conditions (a) and EQ and rewarding-analyst (RA) conditions (b). Horizontal lines show the corresponding optimal proportions.

decreasing in the rewarding-hacker condition compared with the equal-payoff condition, overall, for both rewarding-hacker and rewarding-analyst conditions, the proportions of attack actions were similar to that in the equal-payoff condition: rewarding hacker ≈ equal payoff, $F(49, 2450) = 1.16$, $p = .20$, $r = .42$; rewarding analyst ≈ equal payoff, $F(49, 2450) = 0.64$, $p = .98$, $r = .02$. These results show that rewarding hackers caused them to decrease their attack actions away from their optimal levels.

Furthermore, we analyzed the trend of proportion of defend actions across rounds. Figures 7a and 7b show the average defend proportions across rounds in rewarding-hacker and rewarding-analyst conditions compared with that in the equal-payoff condition, respectively. In the rewarding-hacker condition, the defend proportions did not change with increasing rounds, and their trend was significantly different from the increasing trend in the equal-payoff condition, $F(49, 2450) = 1.42$, $p < .05$, $r = .23$. In the rewarding-analyst condition, the defend proportions appeared decreasing, and their trend was significantly different from the trend in the equal-payoff condition, $F(49, 2450) = 1.83$, $p < .001$, $r = .58$. Overall, these results imply that rewarding analysts caused them to decrease their defend actions; however, this decrease fell short of reaching optimal levels.

Correlations Between Attack and Defend Proportions and Demographic Variables

First, we pooled data across all three conditions and then correlated proportion of attack and defend actions with demographic variables (age, gender, and education levels). The correlations between attack and defend proportions and demographic variables were insignificant in both the pooled data as well as in the equal-payoff condition. In the rewarding-hacker condition, however, male participants performed more attack actions compared with female participants, $r(24) = .56$, $p = .003$. Furthermore, in the rewarding-analyst condition, higher education levels and males resulted in smaller defend proportions: education level, $r(24) = –.43$, $p = .03$; males, $r(24) = –.41$, $p = .04$.

DISCUSSION AND CONCLUSIONS

Cyberattacks on computer networks is a pressing problem that needs to be addressed

early on (Loukas, 2015). Our results show that motivational factors (monetary payoffs) do influence attack and defend actions in security games involving human players. In particular, motivating players via payoffs for one action (attack or defense) causes a decrease in the proportion of that action compared with the baseline. This decrease may drive actions away from their optimal proportions. Overall, our results could be explained based upon PT (Kahneman & Tversky, 1979; Tversky & Kahneman, 1992) and IBLT (Dutt et al., 2013).

First, attack and defend proportions decreased when hackers and analysts were rewarded for successfully attacking and defending the network, respectively. When a player is rewarded for an action, then, according to IBLT, the expectation for the rewarded action becomes higher than the expectation for other actions (Dutt et al., 2013). According to Dutt et al. (2013), a decision maker would likely perform more of an action that is perceived to provide a higher expectation. However, excessive reliance on rewarded actions will also make it difficult for people to maximize their overall payoffs due to losses incurred from such actions. According to PT, losses influence people's decisions more than gains. Thus, the fear of encountering losses from rewarded actions will make players careful about such actions and decrease their proportions.

The aforementioned explanation is also supported by the findings in Gonzalez, Ben-Asher, Martin, and Dutt (2015), where cooperation increased in a prisoner's-dilemma game over repeated rounds when players possessed full information about their own and their opponent's actions and payoffs. Possessing full information about actions and payoffs allowed players to take into account this information to maximize their rewards. In the present study, possessing information about actions and outcomes of opponents caused players to perform a lesser proportion of attack and defend actions in order to minimize losses from such actions.

Furthermore, we found that the attack and defend proportions decreased over rounds and were different from their optimal levels. In our experiment, players were specifically asked to maximize their rewards and were rewarded more for certain actions. For getting rewards, players decreased their proportion of rewarded actions to minimize losses on account of excessive reliance on such actions. For the hacker, it meant that his attack proportions moved away from the optimal proportion, and for the analyst, it meant that his defend proportions moved toward the optimal proportion. However, due to cognitive limitations on memory and recall for human players (Dutt et al., 2013), the movement of the rewarded actions, whether toward or away from their optimal levels, was gradual over time.

Our results have implications for the real world. First, due to the movement of defend proportions toward their optimal levels, it is advised that organizations reward experienced analysts for their successful defend actions. Such rewards will likely cause analysts to perform close to their optimal level and hackers, with certain hacking experience, to perform much less than their optimal level. In such a situation, it is likely that hackers get caught by analysts most of the time. However, we also found that rewards for hackers caused analysts to fear the worst and increase the defend proportions. Thus, organizations should take steps to ensure that analysts do not become overenthusiastic in defending networks, especially when they perceive hackers getting high rewards.

Also, we found that level of education and gender influenced optimal actions in certain rewarding conditions. Based upon our results, male analysts, when rewarded, would perform more optimally compared with female analysts. Furthermore, people with higher levels of education would likely make better analysts. Finally, male hackers would likely attack networks more than their female counterparts.

Although our study has implications for the real world, those implications should be considered with respect to certain limitations. First, although we used an anonymous crowd-sourcing platform where about a majority of our participants were from STEM backgrounds, being either students or working in the IT or engineering industry, it is likely that our participants may not be real hackers and analysts. Yet, if one considers real hackers and analysts to possess the same cognitive machinery and biases as non-hackers and non-analysts, then, when they are put in game-theoretic scenarios as ours, one

should likely expect decision making according to IBLT and PT. Second, we expect our results to generalize to humans in scenarios similar to ours, but we conducted a laboratory experiment involving abstract cyberattack situations, and thus our conclusions may be seen in this spirit. Third, we created motivations via monetary rewards. In the real world, however, hackers and analysts may have motivations that go beyond monetary rewards. For example, hackers may want to attack networks in order to gain fame or notoriety; similarly, analysts may prize nonmonetary recognitions, like promotions, rather than financial gains from protecting networks. We believe that nonmonetary motivations are important, and these should be studied by researchers using our proposed game-theoretic paradigm in future.

In the real world, hackers may be multiple and may not have full knowledge of networks they are trying to compromise; also, analysts may not possess full information about their own network, as this information is generated in real time and distributed across several computers. Although we could only speculate currently, we believe that in such a limited information environment, analysts may not be able to exploit higher rewards and thus may perform differently from their optimal defend proportions. Ideas concerning limited information availability and multiple hackers attacking a single network are the ones that we plan to execute as immediate next steps in this research program.

APPENDIX

Instructions of the Experiment

Welcome! This study consists of a single task, where you will be randomly matched with an anonymous individual from a pool of online participants. Both you and the other player will have the opportunity to earn money according to the decisions you make in this task. However, there is no payment unless you fully complete the task. Please read the following scenario:

Shoppers.com is a company that sells a number of items online including clothes, books, and electronic devices. The company has a large number of customers that perform online transactions using the company's website that is installed on the company's Webserver (See Figure 1).

Figure 1. The cyber infrastructure at Shoppers.com. The Hacker tries to steal information and damage files on the Webserver. The Analyst tries to protect the Webserver from cyberattacks from the Hacker.

A ***Webserver*** is a computer that enables customers on the Internet to buy the products. Data about these customers, including financial information, preferences for products, and other details is being collected by Shoppers.com on the Webserver. Thus, Webserver stores every transaction that customers make on Shoppers.com as well as their personal information.

Cyberattacks are now one of the main threats to the security of networks. Computer hackers (i.e., ***the Hacker in Figure 1***) are people knowledgeable about computers and they use this knowledge to steal data and private information in networks and damage files that are important to organizations. In this scenario, the **Hacker** tries to steal information from the Webserver in Figure 1 and also damage files on the Webserver that are important to Shoppers.com. Shoppers.com has hired a security analyst (i.e., ***Analyst in Figure 1***) in order to help the company protect its customers' data against the cyberattacks waged by the **Hacker**.

Today, you will take the role of *the* ***Hacker*** and you are being matched with another participant who will take the role of the ***Analyst.***

In this task, you and your opponent will make multiple decisions sequentially in multiple trials. Your payoffs in each trial will depend on what both you and your opponent do. The Hacker may choose between the following two actions: Attack or Not Attack. The Attack action means that the Hacker attacks Shoppers.com's Webserver to steal information and damage files. The Not Attack

		Analyst	
		Defend	**Not Defend**
Hacker	**Attack**	*-10, 5*	*5, -5*
	Not Attack	*0, -5*	0, 0

Figure 2. The payoffs for the Hacker and Analyst in the scenario. If the Hacker takes an Attack action and the Analyst takes a Defend action, then the Hacker gets −10 points and the Analyst gets +5 points. If the Hacker takes an Attack action and the Analyst takes a Not Defend action, then the Hacker gets +5 points and the Analyst gets −5 points. Similarly if the Hacker takes a Not Attack action and the Analyst takes a Defend action, then the Hacker gets 0 points and the Analyst gets −5 points. Finally, if the Hacker takes a Not Attack action and the Analyst takes a Not Defend action then, both the Hacker and Analyst get 0 points each.

action means that the Hacker does not attack Shoppers.com. The Analyst may choose between the following two actions: Defend or Not Defend. The Defend action means that the Analyst protects Shoppers.com's Webserver against the Hacker's actions. The Not Defend action means that the Analyst decides not to defend Shoppers.com's Webserver. After the Analyst and the Hacker have taken an action, both the Hacker and Analyst receive the following feedback: the actions taken by each other and the payoffs that result from those actions. Figure 2 lists the payoffs for different combinations of actions taken by the Hacker and the Analyst. For example, if the Hacker takes an Attack action and the Analyst takes a Defend action, then the Hacker gets –10 points and the Analyst gets +5 points. If the Hacker takes a Not Attack action and the Analyst takes a Defend action the Hacker gets 0 points and the Analyst gets –5 points. A similar combination could be derived for the other two combinations of Hacker and Analyst actions. After this feedback is received the Hacker and Analyst start a new trial and take one of the two actions again.

You will start with 1000 points in your account. According to your actions and your opponent's actions, you may win or lose points. **Your goal in this task is to win as many points as possible.** At the end of the game, we will select 10 best scoring Hacker participants and 10 best scoring Analyst participants based upon the total points won in the game. These Hacker and Analyst participants will enter a lucky draw each. In the lucky draw, 1 Hacker participant and 1 Analyst participant will be randomly selected and given Amazon.com's gift vouchers worth GBP 10 each as prizes. The winners will be intimated via an email message.

Good luck!

ACKNOWLEDGMENTS

This research was supported by the Department of Science and Technology, Government of India award ("Building a Secure and Trustworthy Cyberspace: A Behavioral Game-Theoretic Approach," Award No. SR/CSRI/28/2013[G]) to Varun Dutt and V. S. Chandrasekhar Pammi. Also, we are grateful to Indian Institute of Technology Mandi for providing the necessary computational resources for this project.

KEY POINTS

- Due to most corporate operations becoming online, the threat of cyberattacks is growing; therefore it is important to study the impact of motivational factors (monetary rewards) on hackers' and analysts' decisions.
- The hacker's and analyst's performance is evaluated by the proportion of attack and defend actions over 50 rounds of interaction in a cybersecurity game. Nash proportions of attack and defend actions are a function of monetary rewards set for these two roles.
- Based upon predictions of instance-based learning theory and prospect theory, decisions of hackers and analysts should be driven by those actions where punishments due to bad decisions are minimized.
- Results indicate that hackers and analysts decrease their proportion of attack and defend actions; however, when hackers are rewarded, then analysts, fearing attack, defend far in excess of what they should optimally do.
- Cyber organizations should reward analysts for their successful defend actions and reduce the perception of rewards for hackers in analysts' minds.

REFERENCES

Alpcan, T., & Başar, T. (2010). *Network security: A decision and game-theoretic approach*. Cambridge, UK: Cambridge University Press.

Anderson, J. R., & Lebiere, C. J. (1998). *The atomic components of thought*. New York, NY: Psychology Press.

Arora, A., & Dutt, V. (2013, July). *Cyber security: Evaluating the effects of attack strategy and base rate through instance based learning*. Paper presented at the 12th International Conference on Cognitive Modeling, Ottawa, Canada.

Camerer, C. (2003). *Behavioral game theory: Experiments in strategic interaction*. Princeton, NJ: Princeton University Press.

Cyber Ark. (2014). *2014 Global Advanced Threat Landscape Survey*. Retrieved from http://www.cyberark.com/resource/2014-global-advanced-threat-landscape-survey

Dutt, V., Ahn, Y. S., & Gonzalez, C. (2011). Cyber situation awareness: Modeling the security analyst in a cyber-attack scenario through instance-based learning. In *Data and applications security and privacy XXV* (pp. 280–292). Berlin, Germany: Springer.

Dutt, V., Ahn, Y. S., & Gonzalez, C. (2013). Cyber situation awareness modeling detection of cyber attacks with instance-based learning theory. *Human Factors*, *55*, 605–618.

Dutt, V., & Gonzalez, C. (2012). Making instance-based learning theory usable and understandable: The instance-based learning tool. *Computers in Human Behaviour*, *28*(4), 1227–1240.

Gonzalez, C., Ben-Asher, N., Martin, J. M., & Dutt, V. (2015). A cognitive model of dynamic cooperation with varied interdependency information. *Cognitive Science*, *39*, 457–495.

Gonzalez, C., Lerch, J. F., & Lebiere, C. (2003). Instance-based learning in dynamic decision making. *Cognitive Science*, *27*, 591–635.

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, *47*, 263–291.

Kaur, A., & Dutt, V. (2013, July). *Cyber situation awareness: Modeling the effects of similarity and scenarios on cyber attack detection*. Paper presented at the 12th International Conference on Cognitive Modeling, Ottawa, Canada.

Lebiere, C. (1999). The dynamics of cognition: An ACT-R model of cognitive arithmetic. *Kognitionswissenschaft*, *8*, 5–19.

Loukas, G. (2015). *Cyber-physical attacks*. Retrieved from http://www.professionalsecurity.co.uk/reviews/cyber-physical-attacks

Prolific. (2016). *Prolific finds the right participants for your surveys and tasks*. Retrieved from https://prolific.ac/

Roy, S., Ellis, C., Shiva, S., Dasgupta, D., Shandilya, V., & Wu, Q. (2010). A survey of game theory as applied to network security. In *2010 43rd Hawaii International Conference on System Sciences (HICSS)* (pp. 1–10). New York, NY: IEEE.

Tobias, M. W. (2015, 9 September). Your cybersecurity: Don't count on the government. *Forbes*. Retrieved from http://www.forbes.com/sites/marcwebertobias/2014/05/12/your-cybersecurity-dont-count-on-the-government/#3e46ddfe79dc

Tversky, A., & Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, *5*, 297–323.

White House, Office of the Press Secretary. (2009). *Remarks by the president on securing our nation's cyber infrastructure*. Retrieved from https://www.whitehouse.gov/the-press-office/remarks-president-securing-our-nations-cyber-infrastructure

Zahid Maqbool is a doctoral student at Applied Cognitive Science Laboratory, School of Computing and Electrical Engineering, Indian Institute of Technology Mandi. His interest is in understanding issues in cybersecurity via behavioral game theory.

Nidhi Makhijani is a BTech (computer science engineering) student in the School of Computing and Electrical Engineering, Indian Institute of Technology Mandi. Her interest is in understanding issues in cybersecurity via programming and simulation.

V. S. Chandrasekhar Pammi is a professor at the Centre of Behavioral and Cognitive Sciences, University of Allahabad, India. His current research includes cognitive and computational neuroscience aspects of decision making, sequential skill learning, cross-modal integration, and spatial navigation in built environments.

Varun Dutt is an assistant professor and principal investigator at the Applied Cognitive Science Laboratory, School of Computing and Electrical Engineering, Indian Institute of Technology Mandi. His current research interests include cybersecurity, cognitive science, computational cognitive modeling, judgment and decision making, and artificial intelligence.

Date received: April 17, 2016
Date accepted: October 25, 2016