Cybersecurity: Influence of patching vulnerabilities on the decision-making of hackers and analysts

Zahid Maqbool
Applied Cognitive Science Laboratory
Indian Institute of Technology, Mandi,
India – 175005
Zahid_maqbool@students.iitmandi.ac.in

V.S. Chandrasekhar Pammi
Centre of Behavioral and Cognitive Sciences
University of Allahabad, India – 211002
cpammi@cbcs.ac.in

Varun Dutt
Applied Cognitive Science Laboratory
Indian Institute of Technology, Mandi,
India – 175005
varun@iitmandi.ac.in

***Abstract*—Patching of vulnerabilities on computer systems by analysts enables us to protect these systems from cyber-attacks. However, even after patching, the computer systems may still be vulnerable to cyber-attacks as the patching process may not be full proof. Currently, little is known about how hacker's attack actions would be influenced by the varying effectiveness of the patching process. The primary objective of this study was to investigate the influence of the patching process on the attack-and-defend decisions of hackers and analysts. In this study, we used a 2-player zero-sum stochastic Markov security game in a lab-based experiment involving participants performing as hackers and analysts. In the experiment, participants were randomly assigned to two between-subjects patching conditions: effective (N = 50) and less-effective (N = 50). In effective patching, the probability of the network to be in a non-vulnerable state was 90% after patching by the analyst; whereas, in less-effective patching, the probability of the network to be in the non-vulnerable state was 50% after patching by the analyst. Results revealed that the proportion of attack and defend actions were similar between effective and less-effective conditions. Furthermore, although the proportion of defend actions were similar between vulnerable and non-vulnerable states, the proportion of attack actions were smaller in the non-vulnerable state compared to the vulnerable state. A majority of time, both players deviated significantly from their Nash equilibria in different conditions and states. We highlight the implications of our results for patching and attack actions in computer networks.**

Keywords— ***Analyst, Attack, Defend, Patching, Cyber Security, Hacker, Markov security games, Nash equilibrium*.**

I. INTRODUCTION

With the explosive growth of the Internet and its extensive use in all sectors, protecting computer networks, programs, and data from cyber-attacks and unauthorized access has become a challenge [1]. Hackers, people who attack computer networks, are always trying to find vulnerabilities and to use them to exploit computer systems [13]. Organizations need to identify and fix these vulnerabilities as their presence pose a serious threat to the normal working of computer systems. Analysts, people who protect computer systems, may patch vulnerabilities on these systems to protect these systems from cyber-attacks [14]. Security patches may enable additional functionality or address security flaws within a program [4]. This patching may be effective sometimes; however, patches may also lead to new vulnerabilities in computer systems [15]. The primary objective of this research is to investigate how the effectiveness of the patching process influences the decisions of hackers and analysts. We investigate our objective by using lab-based experiments involving cyber-security games [5, 6, 7, 16, 17], where the expected outcomes in these games are derived using behavioral game theory [12] and other theories of cognition [5, 6, 7, 8].

The influence of the effectiveness of the patching process may be studied using a Markov security games [11]. In the Markov security game, human players performing as hackers and analysts may take attack/not-attack and defend (patch)/not-defend (not-patch) actions and obtain payoffs as a result of their actions. This interaction between hackers and analysts via each other's actions may be repeated over several rounds. As per the Markov assumption, the last patching action of the analyst may influence the vulnerability of the network to cyber-attacks in the current state. In most cases, patching of vulnerabilities may improve the security of a network (i.e., patching may be effective and it may make the network non-vulnerable to cyber-attacks); however, in some cases patching may also lead to more serious types of bugs and vulnerabilities (i.e., patching may be less-effective and it may make the network vulnerable to cyber-attacks). Thus, in Markov security game, a patching action from analyst in the past may lead the network to be either non-vulnerable or vulnerable to cyber-attacks, depending upon whether the patching was effective or less-effective.

Cui et al. [3] have analyzed the role of information availability in Markov security games related to the risk assessment of network information systems. According to Cui et al. [3], in the absence of repairs to vulnerabilities, cyber-attacks could induce damages that became larger as the attacks spread in the network. In contrast, damages to the network became smaller when analysts were able to timely repair the vulnerabilities present in the network. These findings are consistent with the idea of network behaving according to the defenders' last actions in Markov security games. Cui et al. [3] derived predictions about the Nash equilibria using mathematical simulation techniques; however, research that empirically evaluates the Nash predictions in games involving human players is still lacking in literature. In this paper, using Markov security games in lab-based experiments, we attend to this literature gap and study how the effectiveness of the patching process influences analyst's and hacker's decision actions.

In literature on cognition, Instance-based Learning Theory (IBLT) [5, 6, 7, 8], a theory of decisions from experience in dynamic environments, has been shown to be successful in accounting for decisions of participants performing as hackers and analysts [6, 7, 8, 9]. In this paper, we used IBLT to derive our expectations for the decisions of participants performing as hackers and analysts in Markov security games, where the patching process is either effective or less-effective.

According to IBLT, hackers and analysts possess cognitive limitations on memory and recall and these players rely on recency and frequency of available information to make decisions. Therefore, the application of IBLT to the analyst's and hacker's experiential decisions in Markov security games will help us explain how these decisions are impacted by the effectiveness of patching process as well as how these decisions deviate from their Nash proportions on account of limitations upon memory and recall.

In what follows, we first introduce the Markov security game and the Nash proportions for attack and defend (or patch) proportions in this game. Next, we state our expectations based upon IBLT and test these expectations in a lab-based experiment involving Markov security games. Finally, we discuss our experimental results and the implications of these results for the patching and attack actions in computer networks.

II. THE MARKOV SECURITY GAME

The Markov security game (see Figure 1) [2, 3] is conceptualized as a repeated 2 x 2 zero-sum game. The game is played between two players, a hacker and an analyst. The objective for both players in the game is to maximize individual payoffs when the game is played repeatedly over several rounds (where the end-point is unknown to participants). Each player has two different 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). Attack actions correspond to attacking a computer network; whereas, defend actions correspond to patching computers on the network. When the game is used in the laboratory, one human player is randomly assigned to be the hacker and the other human player is assigned to be the analyst.

As shown in Figure 1A, there are two possible network states for a given set of actions available to participants performing as hackers and analysts, vulnerable (*v*) and not-vulnerable (*nv*). In the *v* state, the probability of hacker being able to penetrate into the network is very high while as in non-vulnerable state the probability of hacker being able to penetrate into the network is low. The transition between the *v* and *nv* states is determined by the analyst's last action (*d* or *nd*). If the analyst chooses to patch computers on the network (i.e., initiate a *d* action) in a round *t*, then this patching action likely increases the network's probability of being in the *nv* state in round *t+1*. In contrast, if the analyst does not initiate a patching action (i.e., the analyst decides to take a *nd* action) in a round *t*, then this lack of patching likely increases the network's probability of being in the *v* state in round *t+1*.

The transitions from state *v* to state *nv* or from state *nv* to state *v* depend on the effectiveness of the patching process: effective or less-effective. If the patching is effective, then the probability of transiting from the state *nv* to the state *v* is small (~ 0.1) and the probability of transiting from state *v* to state n*v* is large (~ 0.8). In contrast, if the patching is less-effective, then the probability of transiting from state *nv* to state *v* and from state *v* to state *nv* are equal (~ 0.5). The probability value of each state in a round *t* is determined by the following Markov process:

$$\text{Prob (t)} = \text{M (.)} * \text{Prob (t -1)} \qquad (1)$$

Where, M (.) refers to state-transition matrix (see Figure 1A for the state-transition matrices corresponding to different analyst actions). Prob (t) and Prob (t - 1) refers to probabilities of being in states *v* and *nv* in round t and t - 1 respectively. The probability of being in states *v* or *nv* at the start of the game is made equally likely (= 0.5):

$$\text{Prob (1)} = \begin{bmatrix} 0.5 \\ 0.5 \end{bmatrix} \qquad (2)$$

Where, the 0.5 value in the first and second rows correspond to the probability of being in the *v* state and the *nv* state, respectively.

As shown in Figure 1B, there are separate sets of payoffs associated with each state *v* and *nv*, due to the combination of hacker's and analyst's individual actions. These payoffs form a zero-sum game, where the hacker's (analyst's) gains are greater when the state is *v* (*nv*). For example, in the matrix for the state *v*, an a - d action results in a reward of 5 points for the analyst and a penalty of 5 points for the hacker (the hacker is caught by the analyst attacking the network). For a - nd action, analysts get -10 points and hackers get +10 points. Similarly, one could derive the payoff for other action combinations in Figure 1B. Upon comparing the matrices in states *v* and *nv*, one would find higher (lower) penalties and lesser (greater) benefits for the hacker (analysts) in the state *v* (*nv*).

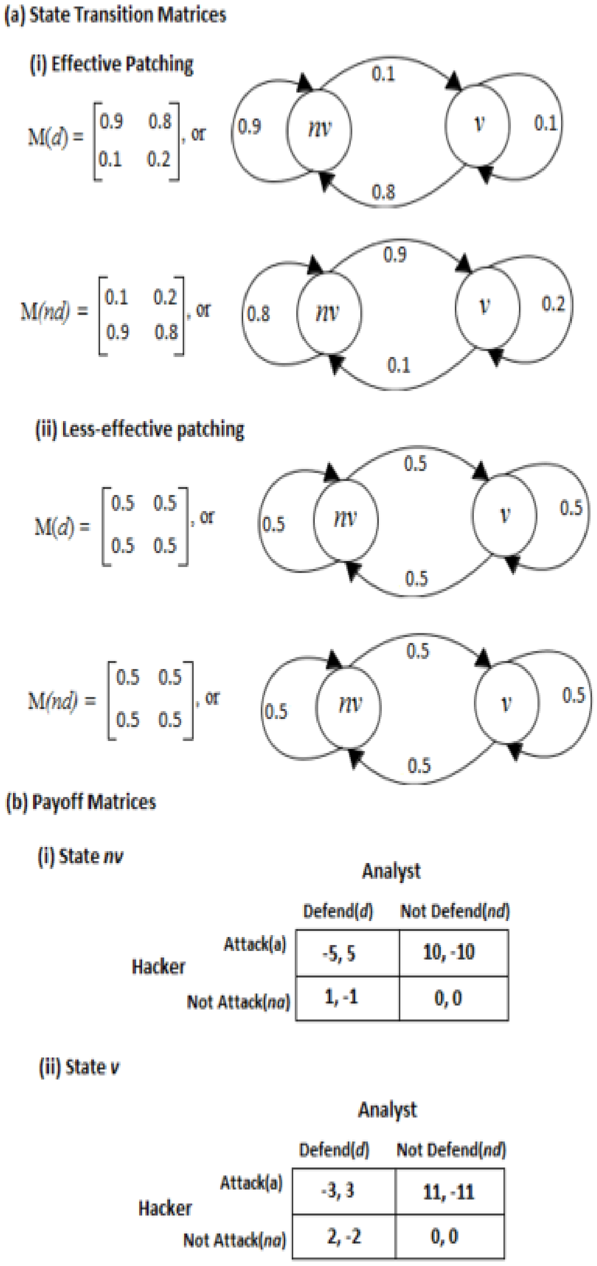

Figure 1. Payoff and transition matrices in a Markov security game. (a) The state transition probability matrices showing transitions between non-vulnerable (*nv*) and vulnerable (*v*) states for both effective patching and less-effective patching conditions. (b) The payoff matrices corresponding to *nv* and *v* states. In each cell, the first payoff is for the hacker and the second payoff is for the analyst.

Using the games in Figure 1B, we computed the mixed strategy Nash equilibria in the Markov security game for the *v* and *nv* states, respectively. Let p represents the proportion of attack actions and 1- p represents the proportion of not-attack actions. Similarly let q represents the proportion of defend actions and 1- q represent the proportion of not-defend actions For Nash equilibria, the hackers and analyst would be indifferent between the payoffs resulting from their two actions. Thus, we get the following Nash proportions:

For the state *v*:

$$3*p - 2*(1 - p) = -11*p + 0 \text{ and } -3*q + 11*(1 - q) = 2*q + 0 \qquad (3)$$

$$\Rightarrow p = 1/8\ (= 0.125) \text{ and } q = 11/16\ (= 0.687)$$

For the state *nv*:

$$5*p - 1*(1 - p) = -10*p + 0 \text{ and } -5*q + 10*(1 - q) = 1*q + 0 \qquad (4)$$

$$\Rightarrow p = 1/16\ (= 0.062) \text{ and } q = 5/8\ (= 0.625)$$

We used these Nash action proportions to compare against human action proportions in a lab-based experiment (reported ahead in this paper).

II. EXPECTATIONS IN THE MARKOV SECURITY GAME

Although the rewards and penalties in the effective and less-effective patching cases are different in the Makov security game (see Figure 1B), the payoffs in these matrices are oriented in the same direction. Thus, across both matrices, the payoff for hackers and analysts are similar and these payoffs possess the same valance (positive or negative). According to IBLT, people maximize their perceived payoff across actions, which is determined by the blended values computed for different actions [5, 6, 7, 8]. As participants performing as hackers and analysts would face similar payoffs across the effective and less-effective patching cases, they would likely possess similar perceived payoffs in both cases. Thus, we expect similar proportion of attack and defend actions across both the effective and less-effective patching conditions. Furthermore, according to IBLT, overall, we expect human decisions to deviate from their Nash proportions. That is because human participants would possess cognitive limitations on memory and recall processes and human beings would tend to rely upon recency and frequency of outcomes to make their repeated decisions. The reliance upon recency and frequency processes would likely not allow participants to form optimal expectations for their actions, where these optimal expectations are determined by the Nash proportions for different actions. We test these expectations in a lab-based experiment next.

II. EXPERIMENT

In this section, we report a lab-based experiment involving people performing as hackers and analysts in the Markov security game (Figure 1). We investigate how the effectiveness of the patching in the game influences the decisions of human players.

A. *Experimental Design*

One hundred participants were randomly assigned to one of two between-subjects conditions: effective patching (N = 50) and less-effective patching (N = 50). In each condition, 25 participants performed as hackers; whereas, 25 participants performed as analysts. The game in each condition was 50-rounds long and involved an interaction between participants performing as hackers and analysts in real time.

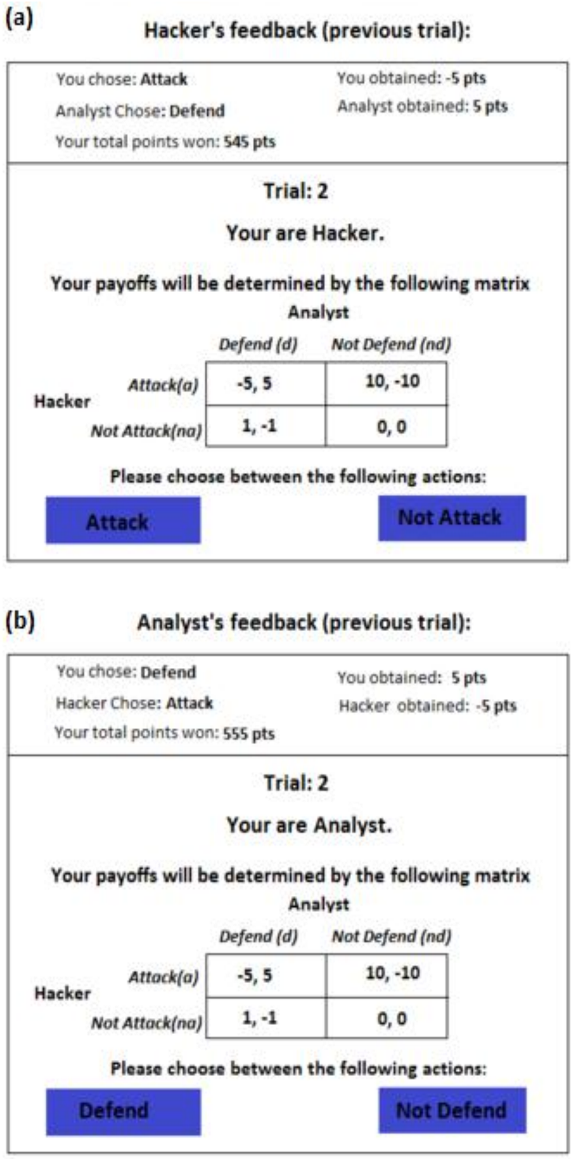

Hacker's feedback (previous trial):

You chose: Attack
You obtained: -5 pts
Analyst Chose: Defend
Analyst obtained: 5 pts
Your total points won: 545 pts

Trial: 2

Your are Hacker.

Your payoffs will be determined by the following matrix

		Analyst	
		Defend (d)	*Not Defend (nd)*
Hacker	*Attack(a)*	-5, 5	10, -10
	Not Attack(na)	1, -1	0, 0

Please choose between the following actions:

Analyst's feedback (previous trial):

You chose: Defend
You obtained: 5 pts
Hacker Chose: Attack
Hacker obtained: -5 pts
Your total points won: 555 pts

Trial: 2

Your are Analyst.

Your payoffs will be determined by the following matrix

		Analyst	
		Defend (d)	*Not Defend (nd)*
Hacker	*Attack(a)*	-5, 5	10, -10
	Not Attack(na)	1, -1	0, 0

Please choose between the following actions:

Figure 2. The graphical user interface shown to participants acting in the roles of hackers (A) and analysts (B) across different patching conditions.

Figure 2 shows the graphical user interface shown to participants acting in the roles of hackers (A) and analysts (B) across different conditions. As shown in Figure 2, in any round, participants were given the following feedback from the last round: actions chosen by them and their opponents, current payoffs obtained by them and their opponents, and total payoffs obtained by them since the start of the game. Both hackers and analysts were also provided with the payoff matrices resulting in different game states (*v* or *nv*) and they were asked to choose between attack/not-attack and defend /not-defend actions in each round. These payoff matrices may change for both players from round-to-round on account of the network being in the *v* or *nv* state (the two payoff matrices are shown in Figure 1B).

For testing our expectations, we compared the proportion of attack and defend actions from human players across different conditions and states. Furthermore, we compared the proportion of attack and defend actions from human players across 5-blocks of 10-rounds each in different conditions and states. Also, we compared the human action proportions with the corresponding Nash action proportions (computed in equations 3 and 4). We used mixed-factorial ANOVAs for testing our expectations. Also, we performed *t*-tests to compare human and Nash proportions in different states and conditions. We used an alpha level of 0.05 and power level of 0.8 for all comparisons.

B. Participants

Seventy-nine percent of participants were males. Ages ranged from 18 years to 30 years (Mean = 21.2 years and standard seviation = 1.92 years). Participants were from different education levels: 74% participants were undergraduates and 26% were graduate students. All participants were from Science, Technology, Engineering, and Mathematics (STEM) backgrounds. Fourty-two percent participants were pursing degrees in computer-science and engineering, 38% participants were pursing degrees in electrical engineering, 18% participants were pursing degrees in mechanical engineering, and 2% participants were pursing degrees in basic sciences. Participants were asked to maximize their payoffs in a cyber-security game and were compensated a flat participation fee of INR 30 (~ USD 0.5). In addition, participants could earn up to INR 20 based on their performance in the game. For calculating the performance incentive, a participant's final score in the game was converted to real money in the following ratio: 55 points in the game = INR 1.0. No participant took more than 20 minutes to finish the study.

C. Procedure

Participant recruitment for the experiment was done through an email advertisement and participation was voluntary. Participants gave their written consent before starting their study and the study was approved by the ethics committee at the Indian Institute of Technology Mandi. Participants were given instructions about the goal in the cyber-security game (to maximize their total payoff) and they were told about the working of the game. As part of the instructions, payoff matrices as well as the set of actions possible were explained to participants. Questions in the instructions, if any, were answered before participants could start their study. Particpants possessed complete information about their own and their opponent's actions and payoffs in all conditions (the payoff matrices were given to both players). Participants could gain or lose points as the game continued. In a round, both participants decided their actions simultaneously and then received feedback 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.

IV. RESULTS

A. *Proportion of attack and defend actions across conditions*

We calculated the proportion of attack and defend actions in each patching condition (see Figure 3). As shown in Figure 3, for the hacker, there was no significant difference in the proportion of attack actions between the less-effective patching and the effective patching conditions (0.31 ~ 0.36; *F*(1, 49) = 0.32, *p* = .57, *ηp²* = .007[1]). Furthermore, for the analyst, there was again no significant difference in the proportion of defend actions between the effective patching and the less-effective patching conditions (0.67 ~ 0.69; *F*(1, 49) = 0.13, *p* = .71, *ηp²* = .003). Overall, these results are as per our expectations.

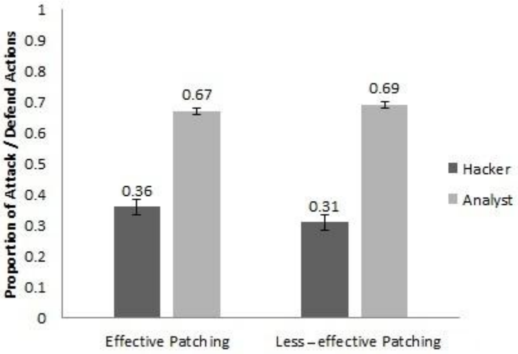

Figure 3. Proportion of attack and defend actions across the two conditions.

B. *Proportion of attack and defend actions across states*

Next, we calculated the proportion of attack and defend actions across different *v* and *nv* states. As shown in Figure 4, the proportion of attack actions were significantly higher in the *v* state compared to the *nv* state (0.38 > 0.28; *F*(1, 98) = 5.70, *p* < .05, *ηp²* =.08). However, there was no significant difference between the proportion of defend actions between the *v* and *nv* states (0.65 ~ 0.71; *F*(1, 98) = 0.74, *p* = .39, *ηp²* =.002). Thus, the defend action proportions agree with our expectations and the attack action proportions do not agree with our expectations. Next, we compared the proportion of attack and defend actions with their Nash proportions in the two states. For hackers, the proportion of attack actions were significantly different from their Nash proportions across both the *v* and *nv* states (state *v*: *t*(49) = 10.34, *p* < .05, *r* = .82; state *nv*: *t*(49) = 7.563, *p* < .05, *r* = .73). For analysts, the proportion of defend actions were not significantly different from their Nash proportions in the *v* state (*t*(49) = -0.481, *p* = .63, *r* = .068) however, the proportion of defend actions were significantly different from their Nash proportion in the *nv* state (*t*(49) = 3.040, *p* < .05, *r* = .40). Thus, overall, these results agree with our expectations.

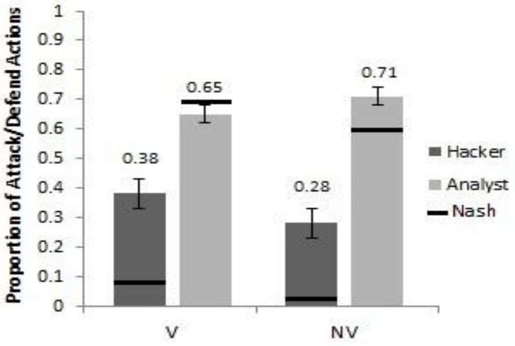

Figure 4. Proportion of attack / defend actions across the states.

C. *Proportion of attack and defend actions across patching conditions and states*

Furthermore, we also analyzed the proportion of attack and defend actions across the two patching conditions and the two network states (i.e., the interaction effect). Figure 5 shows the proportion of attack and defend actions across the patching conditions and states. For hackers, the interaction between conditions and states did not influence the proportion of attack actions (*F*(1, 98) = 0.72, *p* = .39, *ηp²* =.007). Similarly, for analysts, the interaction between conditions and states did not influence the proportion of defend actions (*F*(1, 98) = 0.69, *p* = .41, *ηp²* =.006). Thus, overall, these results agree with our expectations. Next, we compared the proportion of attack and defend actions in different conditions and states with their Nash proportions. For hackers, the proportion of attack actions were significantly different from their Nash proportions across all conditions and states (effective patching and state *v*: *t*(49) = 12.43, *p* < .05, *r* = .87; effective patching and state *nv*: *t*(49) = 14.56, *p* < .05, *r* = .90); less-effective patching and state *v*: *t*(49) = 12.40, *p* < .05, *r* = .87; and, less-effective patching and state *nv*: *t*(49) = 12.12, *p* < .05, *r* = .86). Thus, these results for hackers agree with our expectations. For analysts, the proportion of defend actions was not significantly different from their Nash proportions in the *v* states across both effective and less-effective patching conditions (effective patching: *t*(49) = -1.95, *p* = .06, *r* = .26; less-effective patching: *t*(49) = -1.34, *p* = .18, *r* = .18). However, the proportion of defend actions was significantly different from their Nash proportions in the n*v* states across both effective and less-effective patching conditions (effective patching: *t*(49) = 3.76, *p* < .05, *r* = .28; less-effective patching: *t*(49) = 5.53, *p* < .05, *r* = .61). Overall, these results agree with our expectations in the *nv* states, but not in the *v* state.

[1] This refers to the effect size in the form of partial eta-squared.

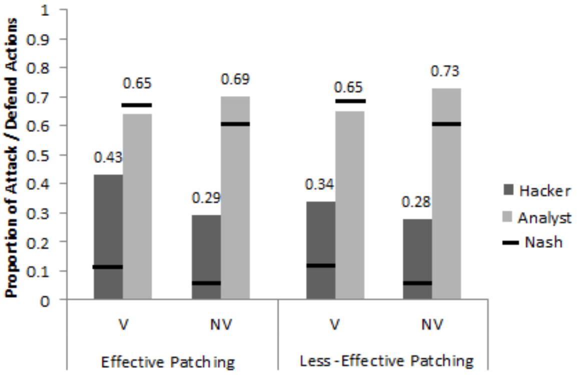

Figure 5. Proportion of attack/defend actions across the patching conditions and network states.

D. *Proportion of attack and defend actions across blocks in different conditions and states*

Next, across both conditions, we calculated the average proportion of attack and defend actions across 5-blocks (consisting of 10-rounds) in different *v* and *nv* states. Figure 6a and 6b shows the proportion of attack actions across 5-blocks in both effective and less-effective patching conditions. We compared the proportion of attack actions across blocks between *v* state and *nv* state. For hackers in effective patching condition, there was no significant difference between the proportion of attack actions across blocks in the *v* state and *nv* state (*F*(4, 192) = 0.19, *p* = .95, *ηp²* = .004). For hackers in less-effective patching condition, again, there was no significant difference between the proportion of attack actions across blocks in the *v* state and *nv* state (*F*(4, 192) = 1.01, *p* = .40, *ηp²* = .020). Overall, these results agree with our expectations.

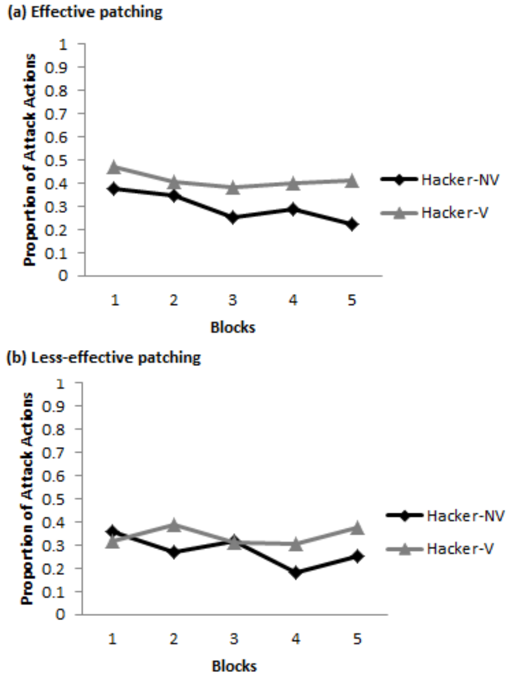

Figure 6. Average proportion of attack actions across rounds. (a) represents the effective patching condition and (b) represents less-effective patching condition

Figure 7a and 7b shows the proportion of defend actions across 5-blocks in both effective and less-effective patching conditions. We compared the proportion of defend actions across blocks between *v* state and *nv* state. For analysts in effective patching condition, there was no significant difference between the proportion of defend actions across blocks in the *v* state and *nv* state (*F*(4, 192) = 0.59, *p* = .67, *ηp²* = .012). For analysts in less-effective patching condition, again, there was no significant difference between the proportion of defend actions across blocks in the *v* state and *nv* state (*F*(4, 192) = 1.25, *p* = .29, *ηp²* = .024). Overall, again, these results agree with our expectations.

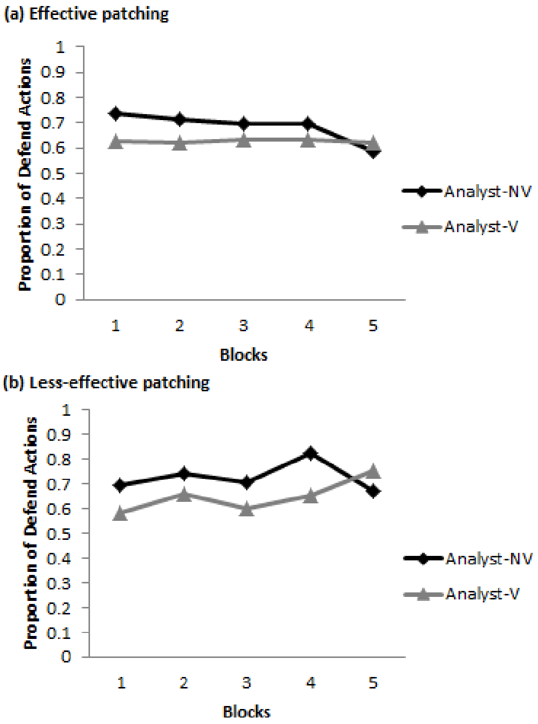

Figure 7. Average proportion of defend actions across rounds. (a) represents the effective patching condition and (b) represents less-effective patching condition

VI. DISCUSSION AND CONCLUSIONS

Due to the increase in cyber-attacks in the real-world, there is an urgent need to patch vulnerabilities present in computer networks [18]. However, this patching process may not be foolproof [19]. For example, in some cases, the patching may be effective and it may make the network less-vulnerable to cyber-attacks; however, in other cases, the patching may be less-effective and it may make the network vulnerable to cyber-attacks. In this paper, using a lab-based experiment, we investigated the influence of effectiveness of the patching processes on decision-making of hackers and analysts. Our results revealed that the proportion of attack and defend actions were similar between effective and less-effective patching conditions. Furthermore, although the proportion of defend actions were similar between vulnerable and non-vulnerable states, the proportion of attack actions were smaller in the non-vulnerable state compared to the vulnerable state. A majority of time, both players deviated significantly from their Nash equilibria in different conditions and states. We explain these results based upon expectations from Instance-based Learning Theory (IBLT) [5, 6, 7, 8] and Behavioral Game Theory (BGT) [12].

First, we found that the proportion of attack and defend actions were similar across the two patching conditions. A likely reason for this finding is the similarity of magnitude and valance of payoffs in the two patching conditions in our experiment. As mentioned above, according to IBLT, people maximize their perceived payoff across actions, which is determined by the blended values computed for different actions [5, 6, 7, 8]. As participants performing as hackers and analysts faced similar payoffs across different patching conditions, they likely possessed similar perceived payoffs in both conditions. Perhaps, it would be interesting to make the payoffs more deviant between the patching conditions as part of future research.

Second, we found that the attack and defend action proportions deviated significantly from their Nash proportions. Again, this expectation can be explained based upon IBLT. According to IBLT, human participants possess cognitive limitations on memory and recall processes and human beings tend to rely upon recency and frequency of outcomes to make their repeated decisions [6, 7]. It seemed that the reliance upon recency and frequency processes in our experiment did not allow participants to form optimal Nash expectations for their actions causing them to deviate significantly from the Nash proportions.

Third, our results revealed that although the proportion of defend actions were similar between vulnerable and non-vulnerable states, the proportion of attack actions were smaller in the non-vulnerable state compared to the vulnerable state. A likely reason for this result could be due to the similarity in defend action proportion across the two states. When analysts continue to patch computer systems in the non-vulnerable state, the hackers get penalized for their attack actions on the network. Also, the penalty for getting caught while attacking in the non-vulnerable state is much higher compared to that in the vulnerable state. As participants were asked to maximize their payoffs, as per BGT, they would tend to reduce those actions that minimize these payoffs. In summary, this reasoning likely caused hackers to reduce their attack proportions in the non-vulnerable state compared to the vulnerable state due to excessive losses in the non-vulnerable state.

Our results also revealed that analyst players did not deviate from their Nash proportions in the vulnerable state although we expected them to deviate from these Nash proportions. One likely reason for this result is that the Nash proportions were simply higher in the vulnerable state compared to those in the non-vulnerable state. As analysts continued to exhibit high patching proportions across both states, their action proportions seem to agree with the Nash proportions in the vulnerable state.

In this paper, we performed a lab-based experiment involving simple Markov security games. Although there are differences between lab-based environments and real-world environments, our results may have important implications for the real world. First, based upon our results, we expect that analysts would continue to excessively patch computer systems in the real-world irrespective of the optimality and the effectiveness of these patching decisions. Second, it seems that hackers, while attacking networks, do not seem to worry about whether computer systems are patched effectively or not. However, hackers do worry about the vulnerability of computer systems to their attacks. Thus, this perception of vulnerability is likely to influence hacker’s cyber-attack decisions. In the real-world, it may be important to showcase computer networks as less vulnerable to cyber-attacks. One could do so via a number of methods including social networks, newspapers, reports, and multimedia.

VII. FUTURE RESEARCH DIRECTIONS

There are a number of research directions that one could undertake as part of future research. Our results revealed that the perception of vulnerability influenced hacker's decisions and there could be several ways in which this perception could be shaped. For example, one could involve deception in computer networks via honeypots, where these honeypots pretend to be easily attackable systems by pretending to being more vulnerable to cyber-attacks. Second, one could involve intrusion-detection systems (IDSs) and provide the knowledge of their existence and accuracy to hackers. For example, if hackers are told that IDSs are not present or they are told that IDSs are present but these are less accurate, then this information is likely to influence the hacker's perception of network's vulnerability to her attacks. Again, in this case, the IDSs may be effective in making hackers attack certain systems (e.g., honeypots) over others and causing them to get caught while waging such attacks. Some of these ideas form the immediate next steps for us to undertake as part of our ongoing research program in game theory and cyber-security.

ACKNOWLEDGEMENT

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 number: SR/CSRI/28/2013(G)) to Varun Dutt and V. S. Chandrasekhar Pammi. Also, we are grateful to the Indian Institute of Technology, Mandi, for providing the necessary computational resources for this project.

REFERENCES

[1] S. Roy, C. Ellis, S. Shiva, D. Dasgupta, V. Shandilya and Q.Wu," A survey of game theory as applied to network security" in *System Sciences (HICSS), 43rd Hawaii International Conference on, 2010* (pp. 1-10). IEEE.

[2] T. Alpcan and T. Başar, *Network security: A decision and game-theoretic approach*. Cambridge University Press. 2010

[3] Cui .X, Tan. X, Yong. Z, & Xi. Z. . A Markov Game Theory-Based Risk Assessment Model for Network Information System. In proceeding of: International Conference on Computer Science and Software Engineering, CSSE 2008, Volume 3: Grid Computing / Distributed and Parallel Computing / Information Security, December 12-14, 2008, Wuhan, China.Source: DBLP

[4] NIST, Glossary of Key Information Security Terms. Available online: http://nvlpubs.nist.gov/nistpubs/ir/2013/NIST.IR.7298r2.pdf

[5] Arora, A., & Dutt, V. (2013). Cyber Security: Evaluating the Effects of Attack Strategy and Base Rate through Instance Based Learning. In *12th International Conference on Cognitive Modeling*. Ottawa, Canada.

[6] 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

[7] C. Gonzalez and V. Dutt," Instance-based learning: Integrating sampling and repeated decisions from experience" *Psychological review*, 2011,118(4), 523.

[8] C. Gonzalez, J. F. Lerch and C. Lebiere," Instance-based learning in dynamic decision making". *Cognitive Science*, 2003, 27(4), 591-635.

[9] C. Gonzalez and V. Dutt, "Refuting data aggregation arguments and how the IBL model stands criticism": A reply to Hills and Hertwig (2012)..

[10] A. Kaur and V. Dutt," Cyber Situation Awareness: Modeling the Effects of Similarity and Scenarios on Cyber Attack Detection" in Paper presented at the *12th International Conference on Cognitive Modeling*. Ottawa, Canada. 2013.

[11] T. Alpcan and T. Basar. An intrusion detection game with limited observations. In12th Int. Symp. on Dynamic Games and Applications, Sophia Antipolis, France 2006 Jul 1 (Vol. 26).

[12] C. Camerer, *Behavioral game theory*: Experiments in strategic interaction. Princeton University Press. 2003

[13] TechTarget, Information security threats. http://searchsecurity.techtarget.com/definition/hacker.

[14] Florida Tech. Cyber security analyst career guide. https://www.floridatechonline.com/blog/information-technology/cybersecurity-analyst-career-guide/

[15] J. Dunagan, R. Roussev, B. Daniels, A. Johnson, C. Verbowski, YM. Wang. Towards a self-managing software patching process using black-box persistent-state manifests. InAutonomic Computing, 2004. Proceedings. International Conference on 2004 May 17 (pp. 106-113). IEEE.

[16] Z. Maqbool, V.C. Pammi, & V Dutt. Cybersecurity: Effect of information availability in security games. In Cyber Situational Awareness, Data Analytics And Assessment (CyberSA), 2016 International Conference On 2016, June (pp. 1-5). IEEE.

[17] Z. Maqbool, N. Makhijani, V.C. Pammi, & V Dutt. Effects of Motivation: Rewarding Hackers for Undetected Attacks Cause Analysts to Perform Poorly. Human Factors, January 2017 59(3), 420-431.

[18] A. Humayed, J. Lin, F. Li, B. Luo. Cyber-Physical Systems Security--A Survey. IEEE Internet of Things Journal. 2017 May 10.

[19] CSO. Why patching is still a problem and how to fix it. https://www.csoonline.com/article/3025807/data-protection/why-patching-is-still-a-problem-and-how-to-fix-it.html

.