Stealthy Attacks Meets Insider Threats: A Three-Player Game Model

Xiaotao Feng*, Zizhan Zheng†, Pengfei Hu†, Derya Cansever‡ and Prasant Mohapatra†
*Department of Electrical and Computer Engineering, University of California, Davis, USA,
†Department of Computer Science, University of California, Davis, USA,
‡US Army CERDEC, USA
Email: {xtfeng, cszheng, pfhu, pmohapatra}@ucdavis.edu, derya.h.cansever.civ@mail.mil

***Abstract*—Advanced persistent threat (APT) is becoming a major threat to cyber security. As APT attacks are often launched by well funded entities that are persistent and stealthy in achieving their goals, they are highly challenging to combat in a cost-effective way. The situation becomes even worse when a sophisticated attacker is further assisted by an insider with privileged access to the inside information. Although stealthy attacks and insider threats have been considered separately in previous works, the coupling of the two is not well understood. As both types of threats are incentive driven, game theory provides a proper tool to understand the fundamental tradeoffs involved. In this paper, we propose the first three-player attacker-defender-insider game to model the strategic interactions among the three parties. Our game extends the two-player FlipIt game model for stealthy takeover by introducing an insider that can trade information to the attacker for a profit. We characterize the subgame perfect equilibria of the game with the defender as the leader and the attacker and the insider as the followers, under two different information trading processes. We make various observations and discuss approaches for achieving more efficient defense in the face of both APT and insider threats.**

I. Introduction

Protecting the nation's infrastructure systems and information technology systems from advanced and ever more sophisticated cyber attacks is a major concern. These attacks, often classified under the name of Advanced Persistent Threat (APT), are launched by well-funded entities and are persistent in pursuing their objectives. Moreover, they often act in a *stealthy* way to avoid being detected to maximize the long-term payoffs. In fact, some notorious cyber attacks remained undetected for months or even longer [1], [2]. Hence, traditional cyber defense techniques focusing on one-shot attacks of known types are insufficient in the face of continuous and stealthy attacks.

The defense task becomes even more challenging when sophisticated attackers are further assisted by insiders within an organization. An insider may have privileged access to system resources and sensitive data including the organization's security practices, and can potentially sell these information to an outside attacker for a profit. The inside information can be directly valuable to the attacker or can help reduce the attacker's cost in achieving its goal. According to the 2014 US State of Cybercrime Survey [3], 28% percent of electronic crime events are known or suspected to have been caused by an insider. Moreover, those crimes with insiders involved are often more costly or damaging than attacks from outsiders. *Therefore, it is critical to understand the interplay between advanced attacks and insider threats and design robust defense strategies accordingly.*

In this paper, we propose a three-player game to model the stealthy behavior of advanced attacks and its interplay with insiders. We observe that although APTs are well-funded, so does the defender is most cases. At the same time, they are both driven by incentives and subject to various constraints. The real challenge is to derive cost-effective defense strategies that strike a balance between the cost of defense and the loss from security breaches. Game theory provides a proper framework to reason about the strategic behavior of each side and help understand the fundamental tradeoffs involved.

Our game model is built upon the two-player FlipIt game model [4] proposed by RSA labs in 2012 for modeling stealthy takeover, by incorporating an insider. We consider a system resource under protection and a continuous time horizon, where at any time instance, either the defender or the attacker can make a move to take over the resource at some cost, and at any time t, the resource is under the control of the player that makes the last move before t. To capture the stealthy behavior of the players, we assume that neither the defender nor the attacker has any real-time feedback about the other side.

To extend this model to the three players setting, we assume that the attacker can purchase inside information to reduce the attack cost. That is, while the defender incurs a fixed cost of each move, the attacker can reduce the cost of move by utilizing the inside information. The insider makes a profit from selling information but also incurs a cost due to security breaches. Since neither the defender nor the attacker get any real-time feedback, it is reasonable to assume that they adopt simple non-adaptive strategies. In this paper, we consider a periodic strategy suggested in [4], where there is a random starting phase and a fixed inter-arrival time between two consecutive moves thereafter. When there is no real-time feedback, periodic defense is optimal against periodic attack, and vice versa. We study the subgame perfect equilibria of the game where the defender first determines and declares its strategy, the attacker and the insider then respond accordingly.

To have a complete game, we need to further state the

information trading process between the attacker and the insider. Two models are considered in this paper. In the first model, the insider first makes an offer to the attacker. The attacker may either accept it or decline it and then determines its strategy accordingly. The overall model then becomes a three-stage Stackerberg game. In the second model, the attacker and the insider are involved in a bargaining process, and their strategies are determined from the Nash bargaining solution [5]. We characterize the subgame perfect equilibria for both models and derive various insights accordingly.

Although game theoretical models have been extensively applied to cyber security [6], [7], [8], [9], [4], previous works mainly focus on one-shot attacks or attacks with known types. The FlipIt game is the first model that captures both the persistence and the stealthy behavior of advanced attacks. The basic model has been extended to the asymmetric information and the multi-node settings in several follow-up works [10], [11], [12], but they all consider the two-player attacker-defender scenario. To the best of our knowledge, the only work that considers both advanced attacks and insider threats is [13], where two separate differential games are used to model the defender-attacker interaction, and the competition among multiple insiders for selling information, respectively. However, the strategic interaction between attacker and insider is not modeled and the interplay between the three types of players is not considered.

Our main contribution can be summarized as follows.

- We propose a three-player FlipIt game model that captures the fundamental coupling of advanced attacks and insider threats.
- We model two types of information trading processes between attacker and insider with different bargaining power at each side, and derive the subgame perfect equilibrium for each case.
- Based on the equilibrium solutions derived, we make suggestions on achieving more cost-effective defense in the face of both advanced attacks and insider threats.

The remainder of this paper is organized as follows. We present the three-player FlipIt game model and discuss our choice of payoff functions in Section II. The analysis of the game is provided in Section III, where we first consider the simplified attacker-defender game without insider, and then study the two information trading models in detail. Various insights are derived from the analysis. We provide numerical results in Section IV, and conclude the paper in Section V.

II. Three-Player FlipIt Game

In this section, we present our three-player game model, which is inspired by the two-player FlipIt game by further introducing an insider that can trade inside information to the attacker.

A. Basic Model

The FlipIt game is a two-player attacker-defender game that was originally designed for modeling stealthy and persistent attacks against computing resources [4]. In the basic FlipIt game, both the attacker and the defender can take action ("flips") to get control over a resource, which generates a payoff to its current owner. There are several rules in this game: 1) The resource is either protected or compromised; 2) Each player is allowed to flip at anytime; 3) If players flip exactly at the same time, the ownership keeps unchanged; 4) Players earn profit for the length of time they control the resource; 5) Each player must pay a cost for every flip; 6) If a player already controls the resource and flips, the ownership is unchanged; 7) If a player does not control the resource and flips, it obtains the ownership; 8) Players do not know the current ownership of the resource before they flip. Figure 1 is a graphical representation of the two-player FlipIt game.

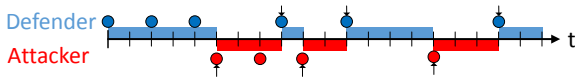

Figure 1: Two-player FlipIt game. *Blue* and *red circles* represent defender and attacker's actions, respectively. *Shaded rectangles* denote the state of resource, *blue* for protected, *red* for compromised. Takeovers are represented by arrows, reproduced from [4].

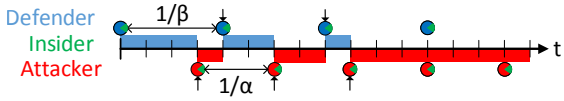

Figure 2: Three-player FlipIt game with insider. The attacker and defender employ periodic strategies at a rate of α and β, respectively. *Green sectors* in both sides indicate that the insider can earn benefit from both the protected system and the attacker.

In this work, we extend the two-player model to the three players case by introducing an insider. The insider can trade information to the attacker to make a profit. The inside information can reduce the attacker's cost as modeled below. On the other hand, similar to the defender, the insider also earns profit from the protected system. Therefore, the insider needs to balance the two types of payoffs. The three-player FlipIt game is graphically demonstrated in Figure 2.

B. Periodic Strategies

Various types of feedback can be studied under the FlipIt game framework. In this paper, we consider the setting where neither the defender nor the attacker can get any real-time feedback from the other side. It is then reasonable to assume that they adopt simple non-adaptive strategies. To this end, we focus on the *periodic strategy* defined in [4], where there is a random starting phase and a fixed inter-arrival time between two consecutive moves thereafter. As shown in [4], when there is no real-time feedback, periodic defense is optimal with respect to periodic attack, and vice versa. Moreover, this is the only strategy that has been analyzed in detail for the two-player FlipIt game with Nash Equilibria determined explicitly. We will consider more general types of feedback including the asymmetric information setting in our future work.

Let $x \in [0, 1]$ denote the long term average time when the resource is compromised, and $1 - x$ be the long-term average

time when the resource is protected. Under the periodic strategies, both the attacker and the defender need to decide their periods of moves. Let α denote the attack rate (i.e., $1/\alpha$ is the attack period), and β be the defense rate. Then x under the periodic strategies can be determined as follows [4]:

$$x = \begin{cases} 1 - \frac{\beta}{2\alpha} & \text{if } \alpha \geq \beta, \\ \frac{\alpha}{2\beta} & \text{if } \alpha \leq \beta. \end{cases}$$

C. Payoff Functions

We model the long-term time average utility gains for the attacker, defender and insider as follows.

Attacker: The attacker can earn benefit from the compromised system resource and incurs cost for taking actions and purchasing information from the insider. The attacker's payoff is defined as follows:

$$P_{\mathcal{A}}(\alpha, \beta, \gamma) = x - C_{\mathcal{A}}(1-\gamma)^b \alpha - \gamma \tag{1}$$

where $C_{\mathcal{A}}$ is the cost for each attack action. The first term in (1) denotes the benefit from the compromised system, the second term is the cost of launching attacks at rate α, and the last term is the cost of purchasing information from the insider at amount of γ. With the inside information, the cost of launching attack is reduced by a factor of $(1-\gamma)^b$. We introduce an exponent $b \geq 0$ to represent how information affects the cost. For instance, $b > 1$ means more information is more useful to the attacker. Due to the space constraint, we only provide results for $b = 2$ in this paper. Analysis of other values of b is given in our online technical report [14].

Defender: The defender can earn benefit from the protected system resource and incurs cost for taking actions. Its payoff is defined as follows:

$$P_{\mathcal{D}}(\alpha, \beta, \gamma) = (1 - x) - C_{\mathcal{D}}\beta \tag{2}$$

where $C_{\mathcal{D}}$ is the cost for each defense action. The first term in (2) denotes the profit from the protected system, and the second term is the cost for recapturing the compromised resources at rate β.

Insider: The insider gains benefit from the resource held by the defender, and also earns profit from selling information to the attacker. Its payoff is defined as follows:

$$P_{\mathcal{I}}(\alpha, \beta, \gamma) = \rho(1 - x) + \gamma \tag{3}$$

where ρ is a constant denoting the insider's proportion in the system. The first term in (3) denotes the profit from the protected system, and the second term is the profit of selling information to the attacker at amount of γ. To restrict the capability of the insider, we assume that $\gamma \leq \rho$.

Table I summarizes the notations used in the paper. We assume that C_A, C_D, and ρ are common knowledge among the players. Compared with previous works [10], [12], [13], we have intentionally minimized the number of parameters so that our model can be applied to more general settings.

Table I: List of Notations

Symbol	Meaning
x	average time when the system is compromised
α	attack rate
β	defense rate
γ	amount of inside information traded
ρ	insider's proportion in the system
C_A	attacker's cost per move
C_D	defender's cost per move

III. Subgame perfect Equilibria

In this section, we study the subgame perfect equilibrium [5] of the three-player game under two different information trading models. In both cases, the defender first determines and declares its defense period, optimized over the possible responses from the attacker and the insider. Note that with the random starting phase of the defense strategy, the actual times where defense actions are taken are unknown to the attacker and the insider. On the other hand, as we show below, the defender can actually gain from revealing (part of) its strategy. In the first information trading model, the insider then makes an offer to the attacker, which is optimized over the attacker's possible response. The attacker may either accept or reject the offer, and then determines its attack rate. In this case, we have found the three-level subgame perfect equilibria of the game. In the second model, given the defender's strategy, the insider and the attacker are involved in a bargaining process, where the amount of information traded is determined from the Nash bargaining solution (together with the attack rate). We have found the subgame perfect equilibria of the game together with the bargaining solutions.

Below we first characterize the subgame perfect equilibria of the attack-defender game when the insider does not exist in Section III-A, which is interesting by itself (only Nash equilibrium is studied in [4] for the two-player game) and also serves as a stepping stone to the more general cases. We then study the take-it-or-leave-it model in Section III-B, and the bargaining model in Section III-C. Based on the equilibrium solutions, we make some observations in Section III-D.

A. Attacker-Defender FlipIt Game

In this section, we study the subgame perfect equilibrium of the game when the insider does not exist. In this case, the payoff functions of attacker and defender are simplified to the follows:

$$P_{\mathcal{A}}(\alpha, \beta) = x - C_{\mathcal{A}}\alpha, \tag{4}$$

$$P_{\mathcal{D}}(\alpha, \beta) = (1 - x) - C_{\mathcal{D}}\beta. \tag{5}$$

To play the game, the defender first determines and declares β. The attacker observes β and then determines α accordingly. The following theorem specifies the subgame perfect equilibria of the two-player game.

Theorem 1. *The set of subgame perfect equilibria of the defender-attacker FlipIt game are:*

$$(\beta^*, \alpha^*) = \begin{cases} (\frac{C_{\mathcal{A}}}{8C_{\mathcal{D}}^2}, \frac{1}{4C_{\mathcal{D}}}), & \frac{C_{\mathcal{A}}}{C_{\mathcal{D}}} < 1, \\ (\frac{1}{2C_{\mathcal{A}}}, 0), & \frac{C_{\mathcal{A}}}{C_{\mathcal{D}}} \geq 1. \end{cases}$$

Proof: First assume $\alpha \geq \beta$. Taking the partial derivative of $P_{\mathcal{A}}(\alpha, \beta)$ in (4) w.r.t. α, we have : $\frac{\partial P_{\mathcal{A}}(\alpha,\beta)}{\partial \alpha} = \frac{\beta}{2\alpha^2} - C_{\mathcal{A}}$.

Hence, $P_{\mathcal{A}}(\alpha,\beta)$ is increasing when $\alpha \in [0, \sqrt{\frac{\beta}{2C_{\mathcal{A}}}}]$, and is decreasing when $\alpha \in [\sqrt{\frac{\beta}{2C_{\mathcal{A}}}}, +\infty)$, with the maximum value achieved at $\alpha_{max} = \max\left\{\beta, \sqrt{\frac{\beta}{2C_{\mathcal{A}}}}\right\}$. If $\alpha_{max} = \beta$, thus, $\beta \geq \frac{1}{2C_{\mathcal{A}}}$, defender's payoff becomes: $P_{\mathcal{D}}(\beta) = \frac{1}{2} - C_{\mathcal{D}}\beta$ with the maximum value $\frac{1}{2} - \frac{C_{\mathcal{D}}}{2C_{\mathcal{A}}}$ at $\beta_{max} = \frac{1}{2C_{\mathcal{A}}}$. Attacker will not take any action for a given $\beta = \frac{1}{2C_{\mathcal{A}}}$ since its payoff $P_{\mathcal{A}} = 0$, so there's no equilibria in this case.

If $\alpha_{max} = \sqrt{\frac{\beta}{2C_{\mathcal{A}}}}$, $P_{\mathcal{D}}(\beta) = \sqrt{\frac{C_{\mathcal{A}}\beta}{2}} - C_{\mathcal{D}}\beta$. Taking derivative of $P_{\mathcal{D}}(\beta)$ w.r.t. β, we have $\frac{\partial P_{\mathcal{D}}(\beta)}{\beta} = \sqrt{\frac{C_{\mathcal{A}}}{8\beta}} - C_{\mathcal{D}}$, $P_{\mathcal{D}}(\beta)$ is increasing when $\beta \in (0, \frac{C_{\mathcal{A}}}{8C_{\mathcal{D}}^2}]$, and decreasing when $\beta \in [\frac{C_{\mathcal{A}}}{8C_{\mathcal{D}}^2}, +\infty)$, it achieves the maximum at $\beta^* = \frac{C_{\mathcal{A}}}{8C_{\mathcal{D}}^2}$. Attacker will take action at rate $\alpha = \frac{1}{4C_{\mathcal{D}}}$ for a given β^* if $\alpha \geq \beta$ and $\frac{C_{\mathcal{A}}}{C_{\mathcal{D}}} < 1$

Next consider the case $\beta \geq \alpha$. We similarly take the partial derivative of $P_{\mathcal{A}}(\alpha,\beta)$ in (4) w.r.t. α and get $\frac{\partial P_{\mathcal{A}}(\alpha,\beta)}{\partial \alpha} = \frac{1}{2\beta} - C_{\mathcal{A}}$, Hence, $P_{\mathcal{A}}(\alpha,\beta)$ is decreasing when $\beta \in (\frac{1}{2C_{\mathcal{A}}}, +\infty)$, and non-decreasing when $\beta \in (-\infty, \frac{1}{2C_{\mathcal{A}}}]$. In both cases, defender achieves its maximum value $\frac{1}{2} - \frac{C_{\mathcal{D}}}{2C_{\mathcal{A}}}$ in (5) at $\beta^* = \frac{1}{2C_{\mathcal{A}}}$, for a given β^* in this case, the attacker plays at rate 0. However, if $\frac{C_{\mathcal{A}}}{C_{\mathcal{D}}} < 1$, defender's payoff $P_{\mathcal{D}}(\beta) < 0$, defender will not take actions, so there is no equilibria in the $\beta \geq \alpha$ and $\frac{C_{\mathcal{A}}}{C_{\mathcal{D}}} < 1$ case.

The theorem then follows from the above analysis by combining the two cases. Moreover, we can obtain the payoffs of the players at the equilibria as follows.
Case 1, $\frac{C_{\mathcal{A}}}{C_{\mathcal{D}}} < 1$: $P_{\mathcal{D}} = \frac{C_{\mathcal{A}}}{8C_{\mathcal{D}}}, P_{\mathcal{A}} = 1 - \frac{C_{\mathcal{A}}}{2C_{\mathcal{D}}}$,
Case 2, $\frac{C_{\mathcal{A}}}{C_{\mathcal{D}}} \geq 1$: $P_{\mathcal{D}} = 1 - \frac{C_{\mathcal{D}}}{2C_{\mathcal{A}}}, P_{\mathcal{A}} = 0$.

■

B. Three Level Sequential Game

We then study the general three-player game. In this section, we consider the case when the insider makes a take-it-or-leave-it offer to the attacker, and the game is played as follows:

1) Defender first determines the defense rate β;
2) Observing β, insider makes an offer γ to the attacker;
3) Observing both β and γ, attacker determines whether to accept the offer as well as the attack rate α.

To solve this game, we use backward induction:

1) Attacker finds $\alpha(\beta,\gamma) = \arg\max_{\alpha} P_{\mathcal{A}}(\alpha,\beta,\gamma)$ for any β and γ;
2) Insider finds $\gamma(\beta) = \arg\max_{\gamma} P_{\mathcal{I}}(\alpha(\beta,\gamma),\beta,\gamma)$ for any β;
3) $\beta^* = \arg\max_{\beta} P_{\mathcal{D}}(\alpha(\beta,\gamma(\beta)),\beta,\gamma(\beta))$ is defender's best strategy;
4) Insider's best response is $\gamma^* = \gamma(\beta^*)$;
5) Attacker's best response is $\alpha^* = \alpha(\beta^*,\gamma^*)$;
6) $(\alpha^*,\beta^*,\gamma^*)$ forms a three level subgame perfect equilibrium.

The following theorem characterizes the subgame perfect equilibria for the three-player game for $b = 2$. A complete analysis and proof for all values of b is provided in [14].

Theorem 2. *The set of subgame perfect equilibria of the three-player game with a take-it-or-leave-it trading between attacker and insider for $b = 2$ are:*

$$(\beta^*,\gamma^*,\alpha^*) = \begin{cases} (\frac{C_{\mathcal{A}}}{8C_{\mathcal{D}}^2}, 0, \frac{1}{4C_{\mathcal{D}}}), & \frac{C_{\mathcal{A}}}{C_{\mathcal{D}}} < 1, \\ (\frac{(1-\rho)^2 C_{\mathcal{A}}}{8C_{\mathcal{D}}^2}, \rho, \frac{1}{4C_{\mathcal{D}}}), & 1 \leq \frac{C_{\mathcal{A}}}{C_{\mathcal{D}}} < \frac{2}{1-\rho}, \\ (\frac{1}{2C_{\mathcal{A}}}, 0, 0), & \frac{C_{\mathcal{A}}}{C_{\mathcal{D}}} \geq \frac{2}{1-\rho}. \end{cases}$$

We further obtain the payoff of each player as follows:
Case 1, $\frac{C_{\mathcal{A}}}{C_{\mathcal{D}}} < 1$: $P_{\mathcal{D}} = \frac{C_{\mathcal{A}}}{8C_{\mathcal{D}}}, P_{\mathcal{I}} = \frac{\rho C_{\mathcal{A}}}{4C_{\mathcal{D}}}, P_{\mathcal{A}} = 1 - \frac{C_{\mathcal{A}}}{2C_{\mathcal{D}}}$.
Case 2, $1 \leq \frac{C_{\mathcal{A}}}{C_{\mathcal{D}}} < \frac{2}{1-\rho}$:

$$P_{\mathcal{D}} = \frac{(1-\rho)^2 C_{\mathcal{A}}}{8C_{\mathcal{D}}},$$
$$P_{\mathcal{I}} = \frac{\rho(1-\rho)^2 C_{\mathcal{A}}}{4C_{\mathcal{D}}} + \rho,$$
$$P_{\mathcal{A}} = (1-\rho)[1 - \frac{(1-\rho)C_{\mathcal{A}}}{2C_{\mathcal{D}}}].$$

Case 3, $\frac{C_{\mathcal{A}}}{C_{\mathcal{D}}} \geq \frac{2}{1-\rho}$: $P_{\mathcal{D}} = 1 - \frac{C_{\mathcal{D}}}{2C_{\mathcal{A}}}, P_{\mathcal{I}} = \rho, P_{\mathcal{A}} = 0$

C. Nash Bargaining between Insider and Attacker

In the previous section, we have considered the setting where the insider can make a take-it-or-leave-it offer to the attacker, which gives the insider some advantage. In this section, we consider a different scenario where the attacker and the insider are involved in a bargaining game with alternating offers. This is arguably more practical and is more fair for the attacker. We adopt the Nash bargaining solution [5] as the solution concept that describes the agreement between the attacker and the insider in equilibrium. For given β and γ, let $\alpha(\beta,\gamma) = \arg\max_{\alpha} P_{\mathcal{A}}(\beta,\gamma,\alpha)$. Let $p_A(\beta,\gamma) \triangleq P_{\mathcal{A}}(\beta,\gamma,\alpha(\beta,\gamma))$, and $p_I(\beta,\gamma) \triangleq P_{\mathcal{I}}(\beta,\gamma,\alpha(\beta,\gamma))$. The Nash bargaining solution is determined by

$$\max_{\gamma\in[0,\rho]} \Big(p_A(\beta,\gamma) - p_A(\beta,0)\Big) \cdot \Big(p_I(\beta,\gamma) - p_I(\beta,0)\Big) \quad (6)$$

To solve this game, we again use backward induction:

1) Attacker finds $\alpha(\beta,\gamma) = \arg\max_{\alpha} P_{\mathcal{A}}(\alpha,\beta,\gamma)$ for any β and γ;
2) For any β, insider and attacker work together to determine $\gamma(\beta)$ that maximizes (6);
3) $\beta^* = \arg\max_{\beta} P_{\mathcal{D}}(\alpha(\beta,\gamma(\beta)),\beta,\gamma(\beta))$ is defender's best strategy;
4) Insider's best response is $\gamma^* = \gamma(\beta^*)$;
5) Attacker's best response is $\alpha^* = \alpha(\beta^*,\gamma^*)$;
6) $(\alpha^*,\beta^*,\gamma^*)$ forms the subgame perfect equilibria of the game.

The following theorem specifies the subgame perfect equilibria of the game. The proof can be found in our technical report [14].

Theorem 3. *The set of subgame perfect equilibria of the three-player game with Nash bargaining between attacker and insider for $b = 2$ are:*

For $\rho \geq 1 - \frac{C_{\mathcal{D}}}{C_{\mathcal{A}}}$*:*

$$(\beta^*, \gamma^*, \alpha^*) = \begin{cases} \left(\frac{C_{\mathcal{A}}}{8C_{\mathcal{D}}^2}, 0, \frac{1}{4C_{\mathcal{D}}}\right), & \frac{C_{\mathcal{A}}}{C_{\mathcal{D}}} < 1, \\ \left(\frac{1}{8C_{\mathcal{A}}}, 1 - \frac{C_{\mathcal{D}}}{C_{\mathcal{A}}}, \frac{1}{4C_{\mathcal{D}}}\right), & 1 \leq \frac{C_{\mathcal{A}}}{C_{\mathcal{D}}} < \frac{1}{1-\rho}. \end{cases}$$

For $\rho < 1 - \frac{C_{\mathcal{D}}}{C_{\mathcal{A}}}$*:*

$$(\beta^*, \gamma^*, \alpha^*) = \begin{cases} \left(\frac{(1-\rho)^2 C_{\mathcal{A}}}{8C_{\mathcal{D}}^2}, \rho, \frac{1}{4C_{\mathcal{D}}}\right), & \frac{1}{1-\rho} \leq \frac{C_{\mathcal{A}}}{C_{\mathcal{D}}} < \frac{2}{1-\rho}, \\ \left(\frac{1}{2C_{\mathcal{A}}}, 0, 0\right), & \frac{C_{\mathcal{A}}}{C_{\mathcal{D}}} \geq \frac{2}{1-\rho}. \end{cases}$$

The corresponding payoffs are:

Case 1: $\rho \geq 1 - \frac{C_{\mathcal{D}}}{C_{\mathcal{A}}}$ and $\frac{C_{\mathcal{A}}}{C_{\mathcal{D}}} < 1$

$$P_{\mathcal{D}} = \frac{C_{\mathcal{A}}}{8C_{\mathcal{D}}}, P_{\mathcal{I}} = \frac{\rho C_{\mathcal{A}}}{4C_{\mathcal{D}}}, P_{\mathcal{A}} = 1 - \frac{C_{\mathcal{A}}}{2C_{\mathcal{D}}}.$$

Case 2: $\rho \geq 1 - \frac{C_{\mathcal{D}}}{C_{\mathcal{A}}}$ and $1 \leq \frac{C_{\mathcal{A}}}{C_{\mathcal{D}}} < \frac{1}{2-\rho}$

$$P_{\mathcal{D}} = \frac{C_{\mathcal{D}}}{8C_{\mathcal{A}}}, P_{\mathcal{I}} = \frac{\rho C_{\mathcal{D}}}{4C_{\mathcal{A}}} + 1 - \frac{C_{\mathcal{D}}}{C_{\mathcal{A}}}, P_{\mathcal{A}} = \frac{C_{\mathcal{D}}}{2C_{\mathcal{A}}}.$$

Case 3: $\rho < 1 - \frac{C_{\mathcal{D}}}{C_{\mathcal{A}}}$ and $\frac{1}{2-\rho} \leq \frac{C_{\mathcal{A}}}{C_{\mathcal{D}}} < \frac{2}{1-\rho}$

$$P_{\mathcal{D}} = \frac{(1-\rho)^2 C_{\mathcal{A}}}{8C_{\mathcal{D}}},$$

$$P_{\mathcal{I}} = \frac{\rho(1-\rho)^2 C_{\mathcal{A}}}{4C_{\mathcal{D}}} + \rho,$$

$$P_{\mathcal{A}} = (1-\rho)[1 - \frac{(1-\rho)C_{\mathcal{A}}}{2C_{\mathcal{D}}}].$$

Case 4: $\rho < 1 - \frac{C_{\mathcal{D}}}{C_{\mathcal{A}}}$ and $\frac{C_{\mathcal{A}}}{C_{\mathcal{D}}} \geq \frac{2}{1-\rho}$

$$P_{\mathcal{D}} = 1 - \frac{C_{\mathcal{D}}}{2C_{\mathcal{A}}}, P_{\mathcal{I}} = \frac{\rho}{2}, P_{\mathcal{A}} = 0.$$

D. Observations and Suggestions

Based on the theorems above, we can prove the following properties of our three-player game.

- The defender always obtains a higher payoff under subgame perfect equilibria compared with Nash equilibria derived in [4] for the two-player attacker-defender game.
- Comparing to the two-player attacker-defender game, the existence of an insider always improves attacker's payoff (if it accepts insider's offer), and at the same time, decreases defender's payoff in both the three-level subgame perfect model and the Nash bargaining model.
- The attacker always gets more benefit with the increasing proportion of insider in the system (ρ) in the Nash bargaining model. On the contrary, a larger ρ always reduces defender's payoff in both the three-level subgame perfect model and the Nash bargaining model.
- The attacker always gets more benefit from bargaining comparing to the take-it-or-leave-it model for any system parameters $C_{\mathcal{D}}$, $C_{\mathcal{A}}$ and ρ. Defender also gets more benefit although it does not participate in the bargaining game directly. This is because bargaining reduces γ, the amount of inside information traded in equilibrium, which improves defender's profit. On the contrary, bargaining always decreases insider's payoff.

From the above properties, we have the following suggestions to the defender for holding a more secure system:

- The defender can publicly announce (part of) its strategy (e.g., by revealing the protecting period but not the random starting phase) and let attacker move behind to obtain more benefit from the subgame perfect game.
- The defender should take effective practices (e.g., monitoring suspicious behavior, separation of duties, and secure backup) to prevent and detect insider thereby reduce its proportion in the system.
- Intuitively, the defender should get more benefit with a larger $C_{\mathcal{A}}/C_{\mathcal{D}}$ as in traditional attacker-defender games. However, this is not always true when there is an insider. In fact, the defender may sometimes prefer maintaining a relatively small $C_{\mathcal{A}}/C_{\mathcal{D}}$ to prevent the trading between attacker and insider. This phenomenon is further studied in simulations.

IV. Numerical Result

In this section, we examine our proposed three-player game with numerical study under different system configurations and scenarios. In addition to verifying the properties described in Section III-D, we make some further observations from the simulations results.

Subgame Perfect vs. Nash: We first compare the subgame perfect equilibrium found in SectionIII-A with the Nash equilibria derived in [4]. We plot the payoffs of both defender and attacker with $C_{\mathcal{A}}/C_{\mathcal{D}}$ varies between [0.5,2] in Figure 3. We observe that both defender and attacker obtain a higher payoff in the subgame perfect equilibrium case.

Impact of Insider: We then study the impact of an insider. We compare the payoffs in the three-level subgame perfect game with the non-insider case in Figure 4. The proportion of the insider in the system is fixed at $\rho = 0.3$. Figure 4a shows that the defender's payoff increases with $C_{\mathcal{A}}/C_{\mathcal{D}}$ smoothly in most cases, but there is a drop at $\frac{C_{\mathcal{A}}}{C_{\mathcal{D}}} = 1$ where the insider first joins the game (i.e., $\gamma = 0$ before that point), and a jump at $\frac{C_{\mathcal{A}}}{C_{\mathcal{D}}} = 2.5$ where both the insider and the attacker quit the game. From Figure 4b, we find that the existence of the insider provides a positive payoff to the attacker for $\frac{C_{\mathcal{A}}}{C_{\mathcal{D}}} \in [1, 2.5]$, where the attacker originally gets nothing in non-insider case.

We further demonstrate the impact of an increasing ρ in Figure 5. Figure 5a shows the payoff of each player in subgame perfect equilibrium under take-it-or-leave-it model with $\frac{C_{\mathcal{A}}}{C_{\mathcal{D}}} = 3$. Interestingly, a more powerful insider does not always help attacker to seize more profit (attacker's payoff decreases after $\rho = \frac{2}{3}$). This is because the insider is the leader and tends to sell as much information as possible to the attacker to maximize its benefit. The attacker may not need that much information, but it has to accept the offer, since otherwise it will get nothing. On the other hand, in the bargaining case shown in Figures 5b and 5c, attacker's payoff is non-decreasing as a function of ρ, due to the increased power of negotiation in this case.

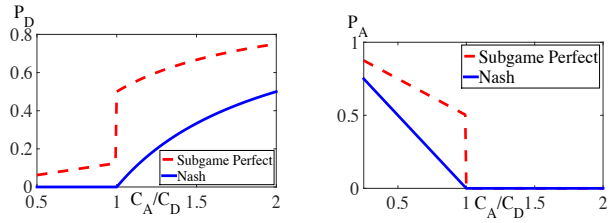

(a) Defender's payoff (b) Attacker's payoff

Figure 3: Payoff vs. $C_{\mathcal{A}}/C_{\mathcal{D}}$ under subgame perfect equilibrium and Nash equilibrium.

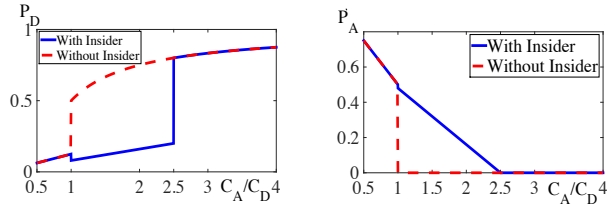

(a) Defender's payoff (b) Attacker's payoff

Figure 4: Payoff vs. $C_{\mathcal{A}}/C_{\mathcal{D}}$ without insider (subgame perfect equilibrium) and with insider (take-it-or-leave-it model), $\rho = 0.3$.

Take-it-or-Leave-it vs. Bargaining: Finally, we compare the payoffs of each of the three players in the take-it-or-leave-it model and the Nash bargaining model in Figure 6. We observe that the attacker always gets more benefit from bargaining while the insider prefers the take-it-or-leave-it case as expected. On the other hand, the defender can also benefit from bargaining even if it does not participate in the bargaining game directly, due to the fact that less information is traded in the bargaining case compared with the take-it-or-leave-it case.

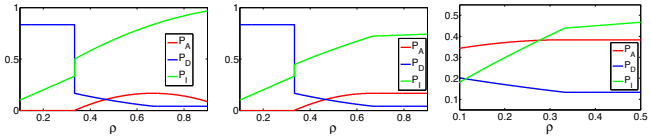

(a) TILI, $\frac{C_{\mathcal{A}}}{C_{\mathcal{D}}} = 3$ (b) bargaining, $\frac{C_{\mathcal{A}}}{C_{\mathcal{D}}} = 3$ (c) bargaining, $\frac{C_{\mathcal{A}}}{C_{\mathcal{D}}} = 1.5$

Figure 5: Payoff vs. ρ in the take-it-or-leave-it (TILI) model and the bargaining model.

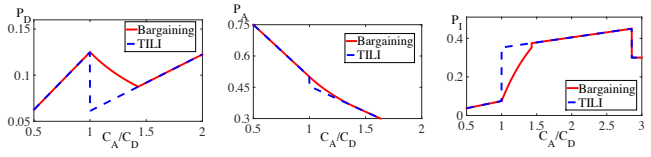

(a) Defender (b) Attacker (c) Insider

Figure 6: Payoff vs. $C_{\mathcal{A}}/C_{\mathcal{D}}$ in take-it-or-leave-it (TILI) model and Nash bargaining model. $\rho = 0.3$.

V. Conclusions

Advanced attacks with stealthy behavior and insider threats are two major concerns to cyber security. The coupling of the two can inflict even big damage to our nation's infrastructure and information technology systems. In this paper, we present the first three-player attacker-defender-insider game model to understand the interplay between stealthy attacks and insider threats. Our model is built upon the two-player FlipIt game where the attacker can purchase information from an insider. We characterize the subgame perfect equilibria of the game with defender as the leader and attacker and insider as the follower, under two different information trading processes. Various insights on achieving more cost-effective defense are derived.

VI. Acknowledgement

Research was sponsored by the Army Research Laboratory and was accomplished under Cooperative Agreement Number W911NF-13-2-0045 (ARL Cyber Security CRA). The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation here on.

References

[1] "ESET and Sucuri Uncover Linux/Cdorked.A: The Most Sophisticated Apache Backdoor," http://www.eset.com/int/about/press/articles/article/eset-and-sucuri-uncover-linuxcdorkeda-apache-webserver-backdoor-the-most-sophisticated-ever-affecting-thousands-of-web-sites/, 2013.

[2] B. Bencsáth, G. Pék, L. Buttyán, and M. Félegyházi, "The Cousins of Stuxnet: Duqu, Flame, and Gauss," *Future Internet*, vol. 4, pp. 971–1003, 2012.

[3] "2014 US State of Cybercrime Survey," http://www.pwc.com/us/en/increasing-it-effectiveness/publications/us-state-of-cybercrime.jhtml.

[4] M. van Dijk, A. Juels, A. Oprea, and R. L. Rivest, "FlipIt: The Game of "Stealthy Takeover"," *Journal of Cryptology*, vol. 26, no. 4, pp. 655–713, 2013.

[5] M. J. Osborne and A. Rubinstein, *A Course in Game Theory*. The MIT Press, 1994.

[6] T. Alpcan and T. Başar, *Network Security: A Decision and Game-Theoretic Approach*. Cambridge University Press, 2010.

[7] H. Kunreuther and G. Heal, "Interdependent security," *Journal of Risk and Uncertainty*, vol. 26, no. 2-3, 2003.

[8] A. Gueye, V. Marbukh, and J. C. Walrand, "Towards a Metric for Communication Network Vulnerability to Attacks: A Game Theoretic Approach," in *Proc. of Gamenets*, 2012.

[9] M. Tambe, *Security and Game Theory: Algorithms, Deployed Systems, Lessons Learned*. Cambridge University Press, 2011.

[10] A. Laszka, B. Johnson, and J. Grossklags, "Mitigating Covert Compromises: A Game-Theoretic Model of Targeted and Non-Targeted Covert Attacks," in *Proc. of WINE*, 2013.

[11] A. Laszka, G. Horvath, M. Felegyhazi, and L. Buttyán, "Flipthem: Modeling targeted attacks with flipit for multiple resources," in *Proc. of GameSec*, 2014.

[12] M. Zhang, Z. Zheng, and N. B. Shroff, "Stealthy attacks and observable defenses: A game theoretic model under strict resource constraints," in *Proc. of GlobalSIP*, 2014.

[13] P. Hu, H. Li, H. Fu, D. Cansever, and P. Mohapatra, "Dynamic defense strategy against advanced persistent threat with insiders," in *Proc. of INFOCOM*, 2015.

[14] X. Feng, Z. Zheng, P. Hu, D. Cansever, and P. Mohapatra, "Stealthy Attacks Meets Insider Threats: A Three-Player Game Model," Technical Report, available online at http://spirit.cs.ucdavis.edu/pubs/tr/mil15.pdf.