When to Follow the Tip: Security Games with Strategic Informants

Weiran Shen[1] , **Weizhe Chen**[2] , **Taoan Huang**[3] , **Rohit Singh**[4] and **Fei Fang**[1]
[1]Carnegie Mellon University
[2]Shanghai Jiao Tong University
[3]University of Southern California
[4]World Wide Fund for Nature
emersonswr@gmail.com, chenweizhe@sjtu.edu.cn, taoanhua@usc.edu, rsingh@wwf.sg, feif@cs.cmu.edu

Abstract

Although security games have attracted intensive research attention over the past years, few existing works consider how information from local communities would affect the game. In this paper, we introduce a new player – a strategic informant, who can observe and report upcoming attacks – to the defender-attacker security game setting. Characterized by a private type, the informant has his utility structure that leads to his strategic behaviors. We model the game as a 3-player extensive-form game and propose a novel solution concept of Strong Stackelberg-perfect Bayesian equilibrium. To compute the optimal defender strategy, we first show that although the informant can have infinitely many types in general, the optimal defense plan can only include a finite (exponential) number of different patrol strategies. We then prove that there exists a defense plan with only a linear number of patrol strategies that achieves the optimal defender's utility, which significantly reduces the computational burden and allows us to solve the game in polynomial time using linear programming. Finally, we conduct extensive experiments to show the effect of the strategic informant and demonstrate the effectiveness of our algorithm.

1 Introduction

Protecting wildlife and other natural resources from illegal activities, such as poaching, has been one of the world's common pressing challenges. However, due to insufficient funding and other supportive resources, the current low number in defensive units makes the protection even harder. Such a sharp contrast between protection need and available resources has led to intensive research efforts in applying game-theoretic approaches (especially Stackelberg Security Game (SSG)) to fighting these illegal activities.

Community engagement is one of the important factors that has largely been ignored in the existing literature. With proper design of patrol strategies, the local communities could serve as a surveillance and intelligence network with a much wider range. Information from local communities has been listed as one of the six pillars towards zero-poaching [WWF, 2015] and also plays an important role in other domains such as fighting urban crimes.

However, the local communities may have their own utility structures and can be unwilling to cooperate. For example, if some poaching activities are observed around a farm, the farmer may not give out such information, or even provide false information, if the farm is suffering loss from the wild-human conflict. Such strategic behavior poses additional challenges in designing better strategies for the defender.

To capture these strategic behaviors, we introduce a new player – a strategic informant – to the standard security game and aim to understand how this new player would affect the original game between the defender and the attacker. We first model the game as an extensive-form game and view the informant as an extra layer between the defender and the attacker: after the attacker sets a target but before he actually attacks, the informant strategically chooses whether to report and what message to report to the defender. When making a defense plan, the defender needs to take the strategic reasoning of both the attacker and the informant into consideration.

To better understand the structure of the 3-player game, we propose a new solution concept called *strong Stackelberg-perfect Bayesian equilibrium (SS-PBE)*. Essentially, such a solution concept is motivated by viewing the game from two different levels. At a lower level, we choose the perfect Bayesian equilibrium as the solution to the subgame between the attacker and the informant. At a higher level, we view the game as a Stackelberg game between the defender and the other two players. We show that the defender can actually enforce the best perfect Bayesian equilibrium of the subgame by introducing a slight perturbation to the defense plan.

The informant may report different messages under different situations. At first glance, this problem may seem intractable as the number of possible messages the defender needs to design depends on the number of different informant types, leading to a heavy computational burden. We borrow tools from the mechanism design domain and show that a variant of the revelation principle holds in our setting. Based on that, we further reduce the number of messages to $n + 1$ (n is the number of targets), which does not depend on the number of informant types.

Following the convention of solving Stackelberg games, we formulate the problem of finding the optimal defense plan as solving multiple linear programs. To show the effect of

strategic informants, we conduct experiments to evaluate the defender's utility and the number of defensive resources need, by changing the informant type distribution and show the impact if the defender fails to take into account the strategic behaviors. Our results provide a useful guideline to law enforcement agencies when facing strategic informants.

1.1 Related Works

The most relevant topic is the Stackelberg games [Basilico *et al.*, 2016; Conitzer and Sandholm, 2006]. Stackelberg Security Game [Tambe, 2011] has been applied to a variety of security problems [Rosenfeld and Kraus, 2017; Fang *et al.*, 2017; Schlenker *et al.*, 2018; Pita *et al.*, 2008]. Previous work on green security games with community engagement [Huang *et al.*, 2020] has considered the presence of non-strategic informants, whereas we consider informants with his own utility structures. Security games with the presence of alarm systems, drones and cameras that can provide real time information have also been studied [Basilico *et al.*, 2017; Ma *et al.*, 2018; Guo *et al.*, 2017]. [Gan *et al.*, 2019] studies security game with deceptive attackers. [Xu *et al.*, 2018; Bondi *et al.*, 2020] study a variant of security game that considers sensors with detection and signaling capability. Our work also makes use of the revelation principle [Myerson, 1981] in mechanism design and is closely related to finding equilibria in extensive form games [Černỳ *et al.*, 2018].

In criminology, [Smith and Humphreys, 2015; Moreto, 2015; Duffy *et al.*, 2015] investigate the role of community engagement in wildlife conservation. Based on the network of reliable informants, [Linkie *et al.*, 2015; Gill *et al.*, 2014] show the positive effects of community-oriented strategies.

In evolutionary game theory, [Short *et al.*, 2010] shows the effect of the presence of informants and [Short *et al.*, 2013] solves for the optimal informant recruitment strategy.

2 Preliminaries

We consider a game with 3 players: the defender, the attacker and the informant. We study the setting where all the 3 players are strategic. Let $T = \{1, 2, \dots, n\}$ be the set of targets and assume that the defender has r defensive resources. When a target t is attacked, and if it is covered by the defender, the attacker gets penalty P_t^a while the defender gets reward R_t^d; Otherwise, the attacker gets reward R_t^a and the defender gets penalty P_t^d. We assume $R_t^d > P_t^d$ and $R_t^a > P_t^a$ for all t.

Suppose that when the attacker plans to attack, the attack plan (i.e., the target t) may be observed by an informant with probability p_w. The informant's type θ is randomly drawn from a set Θ of all possible informant types according to a publicly known probability distribution $p(\theta)$. The informant gets utility $U_t^c(\theta)$ ($U_t^u(\theta)$) if the attacked target t is covered (uncovered), and we assume $U_t^c(\theta) \neq U_t^u(\theta)$.

The defender's defense plan is a tuple $d = (M, x, x^0)$ containing a routine patrol strategy x^0 (when no messages are reported), a set of possible messages M and a patrol strategy $x : M \mapsto \mathbb{R}^n$ that maps the reported message to a coverage probability. Here we assume that the message set M is also designed by the defender. We also assume that both the attacker and the informant have access to the defense plan d.

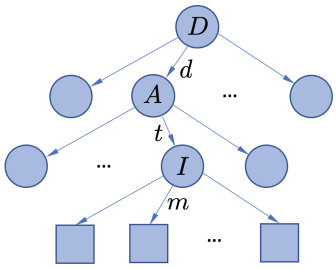

Figure 1: Game tree of the security game with strategic informants, where we omitted the step of Nature choosing the informant type.

After observing the attack plan, the informant can choose to send a certain message m to the defender. The defender then uses the tip-guided patrol strategy $x(m)$ against the attacker. Denote by $x_t(m)$ the coverage probability of target t. We consider the case where the informant is strategic, i.e., given both the defender's defense plan and the attacker's attack plan, the informant strategically chooses a message that maximizes his expected utility.

We assume that the attacker is rational (utility-maximizing) and aware of the existence of the strategic informant when deciding the attack strategy $s \in \Delta(T)$, and that the type of the informant is private information, i.e., only known to the informant himself. The goal of the defender is to design a defense plan (M, x, x^0) to maximize her expected utility.

Formally, we consider the security game with strategic community engagement defined below:

Definition 1. *The security game with strategic community engagement (Figure 1) proceeds as follows:*

1. *The defender announces a defense plan $d = (M, x, x^0)$;*
2. *Observing the defense plan, the attacker decides an attack strategy s, and chooses a target t according to s;*
3. *If the informant observes t, he can remain silent or send a message m to the defender (e.g., "Target t will be attacked" or "The attacker will go south");*
4. *According to the message m, the defender adopts the patrol strategy defined by the announced defense plan.*

3 Solution Concept

To establish our solution concept for the above game, we first focus on the subgame between the attacker and the informant. We call the subgame the "attacker-informant" game. In this game, the attacker moves first, and then the informant chooses a message according to his type. We consider *perfect Bayesian equilibrium* for this game.

Definition 2 (Perfect Bayesian Equilibrium (PBE))**.** *A perfect Bayesian equilibrium is a solution to an extensive form game if the following two conditions are satisfied:*

1. *Sequential rationality: each player's strategy should be optimal given the player's belief;*
2. *Belief consistency: each player's belief should be updated according to the Bayes' rule.*

For any target $t \in T$, informant type $\theta \in \Theta$, and any reported message $m \in M$, the expected utilities of the attacker, the defender and the informant can be written as:

$$\begin{aligned} U_a(t,m) =&(1-p_w)[x_t^0 P_t^a + (1-x_t^0)R_t^a]+ \\ & p_w[x_t(m)P_t^a + (1-x_t(m))R_t^a], \quad (1) \\ U_d(t,m) =&(1-p_w)[x_t^0 R_t^d + (1-x_t^0)P_t^d]+ \\ & p_w[x_t(m)R_t^d + (1-x_t(m))P_t^d], \quad (2) \\ U_i(t,m;\theta) =&(1-p_w)[x_t^0 U_t^c(\theta) + (1-x_t^0)U_t^u(\theta)]+ \\ & p_w[x_t(m)U_t^c(\theta) + (1-x_t(m))U_t^u(\theta)]. \quad (3) \end{aligned}$$

Lemma 1. *Let $m = m(t;\theta)$ be any strategy of the informant. There exists an attacker strategy s, such that (s,m) is a perfect Bayesian equilibrium of the attacker-informant game, if and only if m satisfies*[1]*:*

$$U_i(t,m;\theta) = \max_{m'}\{U_i(t,m';\theta)\}, \forall t,\theta. \quad (4)$$

Moreover, for any perfect Bayesian equilibrium (s,m), the utilities for all three players in the original game only depends on the attacker strategy s.

Proof. **Necessity.** The informant knows his own type and only the attacker has a belief about the informant's type. And since the attacker moves first and only moves once in the game, his belief will remain as the prior type distribution $p(\theta)$ during the two-step game. The informant has access to the attacker's actual action. As (s,m) is a perfect Bayesian equilibrium, by definition, we have $\forall t,\theta$:

$$\begin{aligned} m(t;\theta) \in& \arg\max_m\{U_i(t,m;\theta)\} \\ =& \begin{cases} \arg\max_m x_t(m) & \text{if } U_t^c(\theta) > U_t^u(\theta) \\ \arg\min_m x_t(m) & \text{if } U_t^c(\theta) < U_t^u(\theta) \end{cases}. \quad (5) \end{aligned}$$

Sufficiency. Since the above equation does not contain the attacker's strategy s, any informant strategy satisfying Equation (4) is actually a weakly dominant strategy. The above equation also implies that, for any t and θ, any such strategy $m(t;\theta)$ results in the same coverage probability $x_t(m)$ for target t. On the other hand, when the attacker moves, his strategy should optimize his expected utility:

$$\begin{aligned} & \mathop{\mathbf{E}}_{\theta,t}[U_a(t,m)] \\ =& \sum_\theta p(\theta) \sum_t s(t) \left\{(1-p_w)[x_t^0 P_t^a + (1-x_t^0)R_t^a] + \right. \\ & \left. p_w[x_t(m)P_t^a + (1-x_t(m))R_t^a]\right\} \\ =& \sum_t s(t) \Bigg\{ (1-p_w)[x_t^0 P_t^a + (1-x_t^0)R_t^a]+ \\ & p_w \left[P_t^a \sum_\theta p(\theta)x_t(m) + R_t^a \sum_\theta p(\theta)(1-x_t(m)) \right] \Bigg\}. \end{aligned}$$

Thus the attacker's expected utility only depends on the coverage probability $x_t(m)$ for each t and θ. To best respond, the attacker only needs to choose any distribution s over the set $\arg\max_t \mathbf{E}_{\theta,t}[U_a(t,m)]$.

The above analysis shows that switching to any other strategy satisfying Equation (4) does not change the utilities for both the attacker and the informant. Thus in any PBE (s,m) the expected utilities for both of them only depend on s. To show that the expected utility for the defender also only depends on s, simply notice that the defender's expected utility also only depends on the actual coverage probability of each target (Equation (2)). □

Corollary 1. *It is without loss of generality to only consider a pure strategy for the informant.*

Proof. Immediate from Lemma 1. □

According to Lemma 1, in the original game, the defender's expected utility may depend on how the attacker break ties. In the same spirit of the strong Stackelberg equilibrium, we consider the following solution concept:

Definition 3 (Strong Stackelberg-perfect Bayesian equilibrium (SS-PBE))**.** *A strategy profile (d,s,m) is a Strong Stackelberg-perfect Bayesian equilibrium if:*

1. *(s,m) is a perfect Bayesian equilibrium;*
2. *the attacker breaks ties in favor of the defender;*
3. *based on the above two conditions, d maximizes the defender's expected utility.*

Lemma 2. *It is without loss of generality to assume that the informant always reports a message in M after observing any attack plan, i.e., $\min_{m\in M} x_t(m) \le x_t^0 \le \max_{m\in M} x_t(m), \forall t$.*

Proof. Let (d,s,m) be any SS-PBE that does not satisfy the inequalities. Without loss of generality, assume $x_t^0 < \min_{m\in M}$ for some t. Then we can modify d by adding a new message m' to M with $\min_{m\in M} x_{t'}(m) \le x_{t'}(m') \le \max_{m\in M} x_{t'}(m), \forall t' \neq t$ and $x_t(m') = x_t^0$. Clearly, (d,s,m) is still an SS-PBE. And according to Equation (5) and Lemma 1, it would still be an SS-PBE if we slightly modify m so that the informant always break ties to favor reporting a message in M. For example, we can set $m_t(\theta) = m'$ for all θ with $U_t^c(\theta) < U_t^u(\theta)$. We can repeat the above process until the inequalities in the lemma are satisfied. In the final SS-PBE, all the three players' utilities are the same, and the informant always reports a message in M after observing any attack plan. □

4 Problem Analysis

Once the attacker has chosen a target t to attack, the expected utilities for both the attacker and the defender only depend on the defender's actual coverage probability x_t of target t, which, in turn, depends on the informant's reported message m. In general, the concrete meaning of the message is irrelevant as long as both the informant and the defender interpret it as the same patrol strategy $x(m)$. However, to help later analysis, we start with the case where $M = T \times \Theta$ and consider the following direct defense plan, analogous to the direct or revelation mechanism in the mechanism design literature.

[1] Throughout the paper, we assume $\max_{m'}\{U_i(t,m';\theta)\}$ always exists even if $|M| = \infty$. Otherwise, there can be no equilibrium.

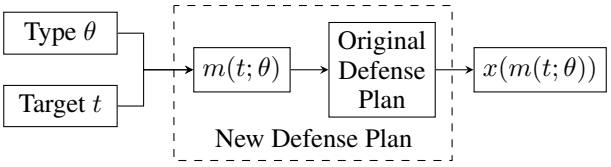

Figure 2: Framework of the revelation principle

Definition 4 (Direct Defense Plan)**.** *A direct defense plan is a tuple $(\bar{M}, \bar{x}, \bar{x}^0)$ where $\bar{M} = T \times \Theta$. A direct defense plan is truthful, if the informant's best strategy is to report the actual target of the attacker and his true type, i.e., $(t, \theta) = \bar{m}(t;\theta), \forall t \in T, \forall \theta \in \Theta$.*

Now we consider a variant of the well-known revelation principle [Myerson, 1981] that fits in our setting and also provide a brief proof for completeness.

Theorem 1 (Revelation Principle [Myerson, 1981])**.** *For any defense plan (M, x, x^0), there exists a truthful direct defense plan $(\bar{M}, \bar{x}, \bar{x}^0)$, such that for any target t chosen by the attacker, and any informant type θ, all the 3 players obtain the same expected utilities as in the original defense plan.*

The intuition behind the revelation principle is to let the mechanism "lie" for the informant (See Figure 2).

Proof. Let $m = m(t;\theta)$ be the informant's strategy in the original defense plan. Then the defender uses the patrol strategy $x(m(t;\theta))$. Let $\bar{x}^0 = x^0$ and define $\bar{x}(t,\theta) = x(m(t;\theta)), \forall t, \forall \theta$. It is easy to see that the direct plan $(\bar{M}, \bar{x}, \bar{x}^0)$ is truthful. Otherwise, assume that reporting a different (t', θ') leads to a strictly higher informant's utility, i.e.,

$$\begin{aligned} &c(t) + p_w[\bar{x}_t(t',\theta')U_t^c(\theta) + (1 - \bar{x}_t(t',\theta'))U_t^u(\theta)] \\ &\quad > c(t) + p_w[\bar{x}_t(t,\theta)U_t^c(\theta) + (1 - \bar{x}_t(t,\theta))U_t^u(\theta)], \end{aligned}$$

where $c(t) = (1-p_w)[x_t^0 U_t^c + (1-x_t^0)U_t^u]$. This means that in the original defense plan, we have:

$$\begin{aligned} &c(t) + p_w\{x_t(m(t';\theta'))U_t^c(\theta) + [1 - x_t(m(t';\theta'))]U_t^u(\theta)\} \\ &> c(t) + p_w\{x_t(m(t;\theta))U_t^c(\theta) + [1 - x_t(m(t;\theta))]U_t^u(\theta)\}, \end{aligned}$$

which implies that $m(t', \theta')$ is a better strategy for the informant, contradicting Equation (4).

Now we show that all the three players have the same expected utility under the new direct defense plan. Once the defense plan is given, according to Equation (1), (2) and (3), for any target t and informant type θ, all three players' utilities in the game only depend on the actual coverage probability $x_t(m)$ for the target. For any t and θ, the informant will report $m(t;\theta)$ and (t,θ) in both the two plan, resulting in coverage probabilities $x_t(m(t;\theta))$ and $\bar{x}_t(t,\theta)$. And we have $\bar{x}(t,\theta) = x(m(t;\theta))$ by definition. □

According to Theorem 1, it is without loss of generality to focus on truthful direct plans. We remark that this is only for ease of analysis, while in actual deployment, it may be more appropriate to still use the original format of defense plans.

Although focusing on truthful direct defense plans simplifies our analysis, it is still challenging to compute the optimal plan for the defender. For each message m, we need to specify a patrol strategy, which contains $n = |T|$ variables. And we have $n|\Theta|$ possible messages, which means we need to determine $n^2|\Theta|$ different variables. This could lead to a heavy computational burden if $|\Theta|$ is very large.

However, we claim that it is possible to achieve the optimal defender's utility with only $n+1$ messages (no longer a direct defense plan of course), even though the number of different informant types cannot be controlled by the defender.

We view the game from the defender's perspective and define the *partial outcome* of the game to be the parameterized mapping $y : T \times \Theta \mapsto [0,1]$, that maps a target to its coverage probability, parameterized by the informant's type θ, or equivalently $y(t,\theta) = x_t(m(t;\theta))$.

The following lemma is useful for proving the above claim.

Lemma 3. *Given any message set M and any defender strategy $x(m)$, there are at most 2^n different outcomes, or equivalently, we only need to consider at most 2^n different informant types, since the outcome is parameterized by it.*

Proof sketch. For each t and each θ, the partial outcome $y(t,\theta) = x_t(m)$ depends on the message $m(t;\theta)$, which in turn only depends on whether $U_t^c(\theta)$ is greater than $U_t^u(\theta)$ or not. So for each target t, there are at most 2 different x_t's. And there can be at most 2^n different partial outcomes. □

Now we are ready to show that $|M| = n + 1$ is sufficient to achieve the optimal defender's utility.

Theorem 2. *There exists a defender strategy $d = (M, x, x^0)$, with $|M| = n+1$, that achieves the optimal defender's utility.*

Proof. Let $\hat{d} = (\hat{M}, \hat{x}, \hat{x}^0)$ be an optimal truthful direct defense plan. We will construct a new defense plan based on $\hat{d}$. To ensure truthfulness, $\hat{x}$ must satisfy:

$$\begin{aligned} &\hat{x}_t(t,\theta) \geq \hat{x}_t(t,\theta'), \forall \theta', \forall \theta \text{ with } U_t^c(\theta) > U_t^u(\theta), \\ &\hat{x}_t(t,\theta) \leq \hat{x}_t(t,\theta'), \forall \theta', \forall \theta \text{ with } U_t^c(\theta) < U_t^u(\theta). \end{aligned}$$

Therefore, if two different types θ and θ' both satisfy $U_t^c(\theta) > U_t^u(\theta)$ and $U_t^c(\theta') > U_t^u(\theta')$, then we must have $\hat{x}_t(t,\theta) = \hat{x}_t(t,\theta')$. The coverage probabilities of other targets are irrelevant as long as they guarantee truthfulness.

We construct the new defense plan (M, x, x^0) as follows: first set $x^0 = \hat{x}^0$. Then for each t, let $\theta^+(t)$ be any informant type with $U_t^c(\theta^+(t)) > U_t^u(\theta^+(t))$. We add a message $m^+(t) = (t, \theta^+(t))$ for each t to M, and set $x(m^+(t)) = \hat{x}(t, \theta^+(t))$. In the end, we add another m^- to M, and set $x_t(m^-) = \min_{m \in \hat{M}} \hat{x}_t(m), \forall t$.

With the above construction, we have $|M| = n + 1$. It is easy to check that $\sum_t x_t(m) \leq r, \forall m \in M$. Now we show that this new defense plan has the same expected utility for the defender as in $\hat{d}$. For any target t, and any informant type θ, the informant will either choose to report $m^+(t)$ or m^-. For example, if θ satisfies $U_t^c(\theta) > U_t^u(\theta)$, then $x_t(m^+(t)) = \hat{x}_t(t, \theta^+(t)) = \max_{\theta'}\{\hat{x}_t(t,\theta')\} \geq \max_{m'}\{x_t(m')\}$. If $U_t^c(\theta) > U_t^u(\theta)$, it is clear that $x_t(m^-) = \min_{m \in \hat{M}} \hat{x}_t(m) = \min_{m \in M} x_t(m)$. This means that in defense plan d, no matter which target t the attacker

chooses to attack, the informant will always choose a message that gives exactly the same expected defender (and also attacker and informant) utility as in $\hat{d}$. Taking expectation over t completes the proof. □

Definition 5 (Defender-Aligned and Attacker-Aligned Informant Types)**.** *An informant type θ is said to be defender-aligned if $U_t^c(\theta) > U_t^u(\theta), \forall t$, and attacker-aligned if $U_t^c(\theta) < U_t^(\theta), \forall t$.*

Lemma 4. *If all informant types are attacker-aligned, There exists an optimal defense plan where the defender always uses the routine patrol strategy x^0, i.e., sets $M = \emptyset$.*

Proof. For any defense plan $d = (M, x, x^0)$, the expected defender's utility is:

$$U_d = \sum_\theta p(\theta) \sum_t s(t) \Big\{ (1-p_w) \left[x_t^0 R_t^d + (1-x_t^0) P_t^d \right] + p_w \left[x_t(m(t;\theta)) R_t^d + (1 - x_t(m(t;\theta))) P_t^d \right] \Big\}.$$

Since all informant types are attacker-aligned, we have $U_t^c(\theta) < U_t^u(\theta), \forall t, \theta$, which means:

$$x_t(m(t;\theta)) = \min_{m'} x_t(m') \le x_t^0, \forall t, \theta,$$

where the inequality is by Lemma 2. Thus,

$$\begin{aligned} U_d =& \sum_\theta p(\theta) \sum_t s(t) \Big\{ (1-p_w) \left[x_t^0 R_t^d + (1-x_t^0) P_t^d \right] + \\ & p_w \left[x_t(m(t;\theta))(R_t^d - P_t^d) + P_t^d \right] \Big\} \\ \le & \sum_\theta p(\theta) \sum_t s(t) \Big\{ (1-p_w) \left[x_t^0 (R_t^d - P_t^d) + P_t^d \right] + \\ & p_w \left[x_t^0 (R_t^d - P_t^d) + P_t^d \right] \Big\} \\ = & \sum_t s(t) [x_t^0 U_t^c + (1-x_t^0) U_t^u], \end{aligned}$$

where the inequality is because $R_t^d > P_t^d$, and the last term is the defender's expected utility of always using x_t^0. Since the above analysis is true for any defense plan, it also holds for any optimal defense plan $\hat{d} = (\hat{M}, \hat{x}, \hat{x}^0)$. Thus always using $\hat{x}^0$ gives a weakly better utility, which implies that $\hat{x}^0$ alone is also optimal. □

Example 1 (Effect of different informants)**.** *Suppose there are two targets and the defender has $r = 1$ resource. Consider the following symmetric, zero-sum instance:*

	1	2
1	$d, -d$	$-d, d$
2	$-d, d$	$d, -d$

The defender and the attacker are the row player and the column player, respectively. Clearly, when there is no informant, the defender will just use a strategy $(0.5, 0.5)$, which gives both the defender and the attacker a utility of 0. If there is a defender-aligned informant with $p_w = 1$, then in the optimal defense plan, the defender will always listen to the informant and allocate the 1 unit of defensive resource accordingly, which gives a utility of d. And if there is an attacker-aligned informant, according to Lemma 4, the optimal defense strategy is still to use $(0.5, 0.5)$, leading to a 0 utility. However, if the defender does not know that the informant is attacker-aligned but still listens to him, then the defender will end up with $-d$ utility.

We now consider how to compute the optimal defense plan. When deciding the attack strategy to optimize Equation (1), the attacker cannot observe the informant's type. Thus similar to the standard Stackelberg setting, the optimal attacker strategy can be achieved with a pure strategy, i.e., attacking a certain target with probability 1. We break ties in favor of the defender when attacking multiple targets gives the attacker the same expected utility.

With Theorem 2, we can index the $n+1$ messages such that $m_t = m^+(t)$, and $m_{n+1} = m^-$. To ensure that the informant always chooses m_t and m_{n+1} when t is the target, we need to guarantee that $x_t(m_t) \le x_t(m') \le x_t(m_{n+1}), \forall m' \in M$.

To compute the optimal defense plan, we follow [Conitzer and Sandholm, 2006] and solve a linear program for each target t, and then choose the best defense plan.

maximize:
$$\sum_\theta p(\theta) \Big\{ (1-p_w) \left[x_t^0 R_t^d + (1-x_t^0) P_t^d \right] + p_w \left[x_t(m(t;\theta)) R_t^d + (1 - x_t(m(t;\theta))) P_t^d \right] \Big\}$$
subject to:
$$\begin{aligned} & \sum_\theta p(\theta) \Big\{ (1-p_w) \left[x_t^0 P_t^a + (1-x_t^0) R_t^a \right] + p_w \left[x_t(m(t;\theta)) P_t^a + (1 - x_t(m(t;\theta))) R_t^a \right] \Big\} \\ \ge & \sum_\theta p(\theta) \Big\{ (1-p_w) \left[x_{t'}^0 P_{t'}^a + (1-x_{t'}^0) R_{t'}^a \right] + p_w \left[x_{t'}(m(t';\theta)) P_{t'}^a + (1 - x_{t'}(m(t';\theta))) R_{t'}^a \right] \Big\} && \forall t' \in T \\ & x_{t'}(m_{n+1}) \le x_{t'}(m') && \forall m' \in M, t' \in T \\ & x_{t'}(m_{n+1}) \le x_{t'}^0 && \forall t' \in T \\ & x_{t'}(m_{t'}) \ge x_{t'}(m') && \forall m' \in M, t' \in T \\ & x_{t'}(m_{t'}) \ge x_{t'}^0 && \forall t' \in T \\ & \sum_{t'} x_{t'}^0 \le r \\ & 0 \le x_{t'}^0 \le 1 && \forall t' \in T \\ & \sum_{t'} x_{t'}(m) \le r && \forall m \in M \\ & 0 \le x_{t'}(m) \le 1 && \forall m \in M, t' \in T \end{aligned}$$

In the above linear program, the first constraint ensures that choosing target t is the best strategy for the attacker, and $m(t;\theta)$ is either m_t or m_{n+1}, which only depends on the informant type θ given any target, and can be pre-computed based on $U_t^c(\theta)$ and $U_t^u(\theta)$.

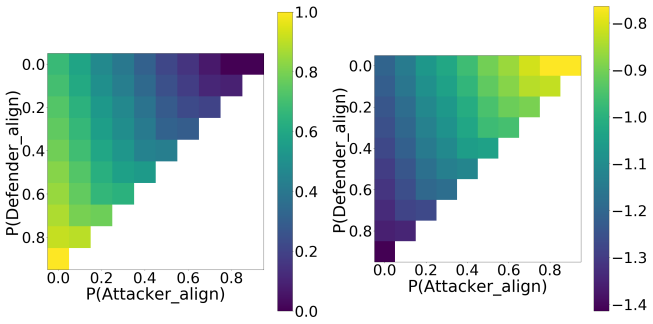

(a) Defender's utility (b) Attacker's utility

Figure 3: Utility heatmaps

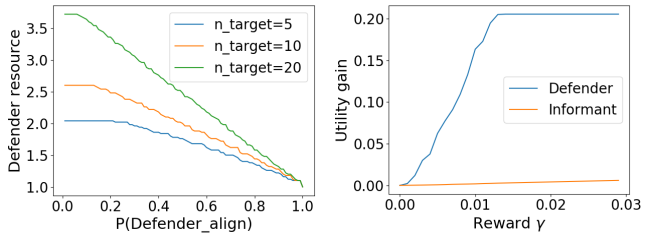

Figure 4: Left: defensive resources needed with a strategic informant. Right: defender's utility gain from rewarding the informant.

5 Experiments

All our results shown in this section are averaged over 50 randomly generated game instances. In each instance, there is a defender with $r = 5$, an attacker, and an informant with $p_w = 0.3$. The rewards and penalties for both the defender and the attacker are drawn from $U[0, 1]$ and $U[-1, 0]$. There are 3 different informant types: a defender-aligned type (θ_1), an attacker-aligned type (θ_2) and a random type (θ_3). All the informant utilities are randomly drawn from $U[0, 0.2]$. In our experiments, we change probabilities for θ_1 and θ_2 to simulate the process of changing from mostly attacker-aligned informants to mostly defender-aligned informants. We implemented our linear program with Python using Gurobi 9.0 [Gurobi Optimization, 2020] as the solver[2]. The running time of our algorithm is listed in Table 1.

5.1 Utility vs. Informant Type

We enumerate all possible type distributions satisfying $p(\theta) \in \{0.1, 0.2, \ldots, 1.0\}, \forall\theta$. We compute the corresponding utilities for both the defender and the attacker to analyze the effect of having different types of informants. As shown in Figure 3, the defender gets higher utilities and the attacker gets lower utilities as the informant goes from fully attacker-aligned to fully defender-aligned, which shows that the existence of the informant could significantly affect the game.

[2]The code can be found at https://github.com/AIandSocialGoodLab/securitygamewithinformants

# targets	10	30	100	200	300	400	500
time/s	0.2	3.4	133.2	1.1k	3.7k	9.6k	18.1k

Table 1: The running time of our algorithm.

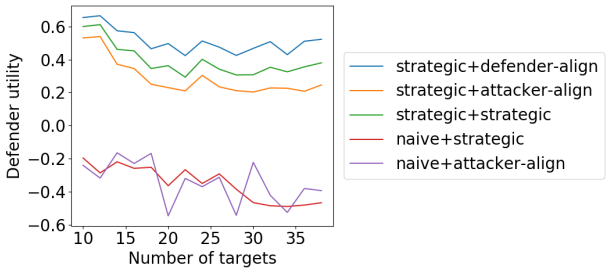

Figure 5: The defender's utilities of different defenders meeting different informants, where "strategic+def_align" means strategic defender and defender-aligned informant.

5.2 Number of Resources vs. Informant Type

We explore how the informant could influence the game in another dimension. We set $p(\theta_3) = 0$, and only change $p(\theta_1)$ and $p(\theta_2)$. The points on the same curve in the left figure of Figure 4 correspond to the same defender's utility. For example, when there are 5 targets, in order to achieve the same utility of having $r = 1$ with a fully defender-aligned informant, a defender needs to have about 2 resources when faced with a fully attacker-aligned informant. And this number goes up quickly when the number of targets increases. This implies that when there is a large number of targets, a defender-aligned informant is worth many defensive resources.

5.3 Effect of Rewarding the Informant

In this experiment, we analyze the effect of incentivizing the informant by giving away rewards. Assume that the defender give a reward γ to the informant, if the defender successfully defends the attack following the informant's message. In this case, the informant will report $m^+(t)$ if $U_t^c(\theta) + \gamma \geq U_t^u(\theta)$. We consider a setting with $p(\theta_3) = 1$. We show in the right figure of Figure 4 that the expected reward given to the informant grows almost linearly with respect to γ. But the defender's utility gain grows quickly in the beginning, and then stops as the informant is already fully defender-aligned.

5.4 Effect of Misclassfying the Informant

Figure 5 shows the results when different defenders meet different informants in . When the defender knows the informant types ("strategic"), the defender's utility goes down when the number of targets increases, and the more defender-aligned the informant is, the more utility the defender gets. However, when the informant is not fully defender-aligned (the "rand" and "att_align"), the defender can suffer a huge loss by blindly following the informant's messages ("naive"). Again, this experiment shows that the strategic behaviors of the informant can have a huge impact on the defender's utility.

Acknowledgments

This work is supported in part by NSF grant IIS-1850477 and a research grant from Lockheed Martin.

References

[Basilico *et al.*, 2016] Nicola Basilico, Stefano Coniglio, and Nicola Gatti. Methods for finding leader-follower equilibria with multiple followers. In *AAMAS'16*, pages 1363–1364, 2016.

[Basilico *et al.*, 2017] Nicola Basilico, Andrea Celli, Giuseppe De Nittis, and Nicola Gatti. Coordinating multiple defensive resources in patrolling games with alarm systems. In *AAMAS'17*, pages 678–686, 2017.

[Bondi *et al.*, 2020] Elizabeth Bondi, Hoon Oh, Haifeng Xu, Fei Fang, Bistra Dilkina, and Milind Tambe. To signal or not to signal: Exploiting uncertain real-time information in signaling games for security and sustainability. 2020.

[Černỳ *et al.*, 2018] Jakub Černỳ, Branislav Boỳanskỳ, and Christopher Kiekintveld. Incremental strategy generation for stackelberg equilibria in extensive-form games. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, pages 151–168, 2018.

[Conitzer and Sandholm, 2006] Vincent Conitzer and Tuomas Sandholm. Computing the optimal strategy to commit to. In *Proceedings of the 7th ACM conference on Electronic commerce*, pages 82–90. ACM, 2006.

[Duffy *et al.*, 2015] Rosaleen Duffy, Freya AV St John, Bram Büscher, and DAN Brockington. The militarization of anti-poaching: undermining long term goals? *Environmental Conservation*, 42(4):345–348, 2015.

[Fang *et al.*, 2017] Fei Fang, Thanh Hong Nguyen, Rob Pickles, Wai Y. Lam, Gopalasamy R. Clements, Bo An, Amandeep Singh, Brian C. Schwedock, Milind Tambe, and Andrew Lemieux. PAWS - A deployed game-theoretic application to combat poaching. *AI Magazine*, 2017.

[Gan *et al.*, 2019] Jiarui Gan, Haifeng Xu, Qingyu Guo, Long Tran-Thanh, Zinovi Rabinovich, and Michael Wooldridge. Imitative follower deception in stackelberg games. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pages 639–657, 2019.

[Gill *et al.*, 2014] Charlotte Gill, David Weisburd, Cody W Telep, and Trevor Bennett. Community-oriented policing to reduce crime, disorder and fear and increase satisfaction and legitimacy among citizens: A systematic review. *Journal of Experimental Criminology*, 2014.

[Guo *et al.*, 2017] Qingyu Guo, Boyuan An, Branislav Bosansky, and Christopher Kiekintveld. Comparing strategic secrecy and stackelberg commitment in security games. In *IJCAI-17*, 2017.

[Gurobi Optimization, 2020] LLC Gurobi Optimization. Gurobi optimizer reference manual, 2020.

[Huang *et al.*, 2020] Taoan Huang, Weiran Shen, David Zeng, Tianyu Gu, Rohit Singh, and Fei Fang. Green security game with community engagement. In *Proceedings of the 2020 International Conference on Autonomous Agents and Multiagent Systems*, 2020.

[Linkie *et al.*, 2015] Matthew Linkie, Deborah J. Martyr, Abishek Harihar, Dian Risdianto, Rudijanta T. Nugraha, Maryati, Nigel Leader-Williams, and Wai-Ming Wong. Editor's choice: Safeguarding sumatran tigers: evaluating effectiveness of law enforcement patrols and local informant networks. *Journal of Applied Ecology*, 2015.

[Ma *et al.*, 2018] Xiaobo Ma, Yihui He, Xiapu Luo, Jianfeng Li, Mengchen Zhao, Bo An, and Xiaohong Guan. Camera placement based on vehicle traffic for better city security surveillance. *IEEE Intelligent Systems*, 33(4):49–61, Jul 2018.

[Moreto, 2015] William D Moreto. Introducing intelligence-led conservation: bridging crime and conservation science. *Crime Science*, 4(1):15, 2015.

[Myerson, 1981] Roger B Myerson. Optimal auction design. *Mathematics of operations research*, 6(1):58–73, 1981.

[Pita *et al.*, 2008] James Pita, Manish Jain, Janusz Marecki, Fernando Ordóñez, Christopher Portway, Milind Tambe, Craig Western, Praveen Paruchuri, and Sarit Kraus. Deployed armor protection: the application of a game theoretic model for security at the los angeles international airport. In *AAMAS: industrial track*, 2008.

[Rosenfeld and Kraus, 2017] Ariel Rosenfeld and Sarit Kraus. When security games hit traffic: Optimal traffic enforcement under one sided uncertainty. In *IJCAI*, pages 3814–3822, 2017.

[Schlenker *et al.*, 2018] Aaron Schlenker, Omkar Thakoor, Haifeng Xu, Fei Fang, Milind Tambe, Long Tran-Thanh, Phebe Vayanos, and Yevgeniy Vorobeychik. Deceiving cyber adversaries: A game theoretic approach. In *AAMAS*, 2018.

[Short *et al.*, 2010] MB Short, PJ Brantingham, and MR D?orsogna. Cooperation and punishment in an adversarial game: How defectors pave the way to a peaceful society. *Physical Review E*, 82(6):066114, 2010.

[Short *et al.*, 2013] Martin B Short, Ashley B Pitcher, and Maria R D'Orsogna. External conversions of player strategy in an evolutionary game: A cost-benefit analysis through optimal control. *European Journal of Applied Mathematics*, 24(1):131–159, 2013.

[Smith and Humphreys, 2015] MLR Smith and Jasper Humphreys. *The Poaching Paradox: Why South Africa's 'Rhino Wars' Shine a Harsh Spotlight on Security and Conservation*. Ashgate Publishing Company, 2015.

[Tambe, 2011] Milind Tambe. *Security and game theory: algorithms, deployed systems, lessons learned*. Cambridge University Press, 2011.

[WWF, 2015] WWF. Developing an approach to community-based crime prevention. http://zeropoaching.org/pdfs/Community-based-crime%20prevention-strategies.pdf, 2015.

[Xu *et al.*, 2018] Haifeng Xu, Kai Wang, Phebe Vayanos, and Milind Tambe. Strategic coordination of human patrollers and mobile sensors with signaling for security games. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.