

Green Security Game with Community Engagement*

Taoan Huang
University of Southern California
taoanhua@usc.edu

Weiran Shen
Carnegie Mellon University
emersonswr@gmail.com

David Zeng
Carnegie Mellon University
dzeng@andrew.cmu.edu

Tianyu Gu
Carnegie Mellon University
tianyug@andrew.cmu.edu

Rohit Singh
World Wide Fund for Nature
rsingh@wwfnet.org

Fei Fang
Carnegie Mellon University
feif@cs.cmu.edu

ABSTRACT

While game-theoretic models and algorithms have been developed to combat illegal activities, such as poaching and over-fishing, in green security domains, none of the existing work considers the crucial aspect of community engagement: community members are recruited by law enforcement agencies as informants and can provide valuable tips, e.g., the location of ongoing illegal activities, to assist patrols. We fill this gap and (i) introduce a novel two-stage security game model for community engagement, with a bipartite graph representing the informant-attacker social network and a level- κ response model for attackers inspired by cognitive hierarchy; (ii) provide complexity results and exact, approximate, and heuristic algorithms for selecting informants and allocating patrollers against level- κ ($\kappa < \infty$) attackers; (iii) provide a novel algorithm to find the optimal defender strategy against level- ∞ attackers, which converts the problem of optimizing a parameterized fixed-point to a bi-level optimization problem, where the inner level is just a linear program, and the outer level has only a linear number of variables and a single linear constraint. We also evaluate the algorithms through extensive experiments.

KEYWORDS

Security Game; Computational Sustainability; Community Engagement

ACM Reference Format:

Taoan Huang, Weiran Shen, David Zeng, Tianyu Gu, Rohit Singh, and Fei Fang. 2020. Green Security Game with Community Engagement. In *Proc. of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020)*, Auckland, New Zealand, May 9–13, 2020, IFAAMAS, 9 pages.

1 INTRODUCTION

Despite the significance of protecting natural resources to environmental sustainability, a common lack of funding leads to an extremely low density of law enforcement units (referred to as defenders) to detect and deter illegal activities such as wildlife poaching and overfishing (referred to as attacks). Insufficient sanctions to the poachers and fishermen doing overfishing (referred to as attackers) make it even more challenging to effectively combat these activities [8, 9]. To improve patrol efficiency, law enforcement

agencies often recruit informants from local communities who can provide tips to the defender [10]. Since attackers are often from the same communities and their activities can sometimes be observed by informants through social interactions, such tips contain detailed information about ongoing or upcoming attacks and can be directly used to guide the patrols. In fact, community engagement is listed by World Wildlife Fund for Nature as one of the *six pillars towards zero poaching* [25]. The importance of community engagement goes beyond these green security domain and extends to domains such as fighting urban crimes [4, 22].

Previous research in computational game theory have led to models and algorithms that can help the defenders allocate limited resources in the presence of attackers, with applications to enforce traffic [17], combat oil-siphoning [23], and deceive cyber adversaries [18] in addition to protecting critical infrastructure [16] and combating wildlife crime [3]. However, none of the work has considered this essential element of community engagement.

Community engagement leads to fundamentally new challenges that do not exist in previous literature. First, the defender not only needs to determine how to patrol but also needs to decide whom to recruit as informants. Second, there can be multiple attackers, and the existence of informants makes the success or failure of their attacks interdependent since any tip about other attackers' actions can change the defender's patrol. Third, because of the combinatorial nature of the tips, representing the defender's strategy requires exponential space, making the problem of finding optimal defender strategy extremely challenging. Fourth, attackers may notice the patrol pattern over time and adapt their strategies accordingly.

In this paper, we provide the first study to fill the gap and provide a novel two-stage security game model for community engagement which represents the social network between potential informants and attackers with a bipartite graph. In the first stage of the game, the defender recruits a set of informants under a budget constraint, and in the second stage, the defender chooses a set of targets to protect based on tips from recruited informants. Inspired by the quantal cognitive hierarchy model [24], we use a level- κ response model for attackers, taking into account the fact that the attacker can take iterative actions and the attacker's strategy will impact the actual marginal strategy of the defender.

Our second contribution includes complexity results and algorithms for computing optimal defender strategy against level- κ ($\kappa < \infty$) attackers. We show that the problem of selecting the optimal set of informants is NP-Hard. Further, based on sampling techniques, we develop an approximation algorithm to compute the optimal patrol strategy and a heuristic algorithm to find the optimal set of informants to recruit. For an expository purpose, we

*The full version of the paper can be found at <https://arxiv.org/abs/2002.09126>.

mainly describe the algorithms for level-0 attackers and provide the extension to level- κ ($0 < \kappa < \infty$) attackers in the last section.

The third contribution is a novel algorithm to find the optimal strategy against level- ∞ attackers, which is extremely challenging: an attacker’s strategy affects the defender’s marginal strategy, which in turn affects the attackers’ strategies and level- ∞ attackers are defined through a fixed-point argument. Thus the defender’s utility relies crucially on solving a parameterized fixed-point problem. A naïve mathematical programming-based formulation is prohibitively large to solve. We instead reduce the program to a bi-level optimization problem, where both levels become more tractable.

Finally, we conduct extensive experiments. We compare the running time and solution quality of different algorithms. We show that our bi-level optimization algorithm achieves better performance than the algorithm adapted from previous works. We also compare level-0 attackers and the case with insider threat (i.e., the attacker is fully aware of the informants), where we formulate the problem as a mathematical program and solve it by adapting an algorithm from previous works. We show that the defender suffers from utility loss if the insider threat is not taken into consideration and the defender still assumes a naïve attacker model (level-0).

2 RELATED WORK AND BACKGROUND

Community engagement is studied in criminology. Duffy et al. [2], Moreto [13], Smith and Humphreys [20] investigate the role of community engagement in wildlife conservation. Gill et al. [4], Linkie et al. [10] show the positive effects of community-oriented strategies. However, they lack a mathematical model for strategic defender-attacker interactions.

Recruitment of informants has also been proposed to study societal attitudes in relation to crimes using evolutionary game theory models. Short et al. [19] formulate the problem of solving recruitment strategies as an optimal control problem to account for limited resources and budget. In contrast to their work, we emphasize the synergy of community engagement and allocation of defensive resources and aim to find the best strategy of recruiting informants and allocating defensive resources.

In security domains, Stackelberg Security Game (SSG) has been applied to a variety of security problems [21], with variants accounting for alarm systems, surveillance cameras, and drones that can provide information in real time [1, 5, 11]. Unlike the sensors that provide location-based information as studied in previous works, the kind of tips the informants can provide depends on their social connections, an essential feature about community engagement.

Other than the full rationality model, boundedly rational behavioral models such as quantal response (QR) [12, 26] and subjective utility quantal response [15] have been explored in the study of SSG. Our model and solution approach are compatible with most existing behavioral models in the SSG literature, but for an expository purpose, we only focus on the QR model.

3 MODEL

In this section, we introduce our novel two-stage green security game with community engagement. The key addition is the consideration of informants. They can be recruited and trained by the defender to provide tips about ongoing or upcoming attacks.

Following existing works on SSG [6, 7], we consider a game with a set of targets $T = [n] = \{1, \dots, n\}$. The defender has r units of defensive resources and each can protect or cover one target with no scheduling constraint. An attacker can choose a target to attack. If target i is attacked, the defender (attacker) receives $R_i^d > 0$ ($P_i^a < 0$) if it is covered, otherwise receives $P_i^d < 0$ ($R_i^a > 0$).

Informants recruited by the defender can provide tips regarding the exact targets in ongoing or upcoming attacks but tip frequency and usefulness may vary due to heterogeneity in the informants’ social connections. We model the interactions and connections between potential informants X (i.e., members of the community that are known to be non-attacker and can be recruited by the defender) and potential attackers Y using a bipartite graph $G_S = (X, Y, E)$ with $X \cap Y = \emptyset$. Here we assume the defender has access to a list of potential attackers which could be provided by the conservation site manager, since the deployment of our work relies on the manager’s domain knowledge, experience, and understanding of the social connections among community members.

When an attacker decides to launch an attack, an informant who interacted with the attacker previously may know his target location. Formally, for each $v \in Y$, we assume that v will attack a target with probability p_v but the target is unknown without informants and each attacker takes actions independently. An edge $(u, v) \in E$ is associated with an information sharing intensity w_{uv} , representing the probability of attack activities of attacker v being reported by u , given v attacks and u is recruited as an informant.

In the first stage, the defender recruits k informants, and in the second stage, the defender receives tips from the informants and allocates r units of defensive resources. The defender’s goal is to maximize the expected utility defined as the summation of the utilities for each attack.

Let U denote the set of recruited informants in the first stage where $|U| \leq k$, and $V = \{v \mid \exists u \in U, (u, v) \in E\}$ denote the set of attackers that are connected with at least one informant in U . We represent tips as a vector of disjoint subsets of attackers $\mathbf{V} = (V_1, \dots, V_n)$, where V_i is the set of attackers who are reported to attack target $i \in T$ such that $V_i \subseteq V$, $V_i \cap V_j = \emptyset$ for any $i, j \in T$. An attacker v is *reported* if there exists $i \in T$ such that $v \in V_i$, otherwise he is *unreported*. We also denote by $V_0 = \bigcup_{i \in T} V_i$ the set of reported attackers. It is possible that $V_0 = \emptyset$ and we say the defender is *informed* if $V_0 \neq \emptyset$. Note that \mathbf{V} is a compact representation of the tips received by the defender as it neglects the identity of the informants, which is not crucial in the defender’s decision making given that all the tips are assumed to be correct.

In practice, tips are infrequent and the defender is often very protective of the informants. Thus, the attackers are often not aware of the existence of informants unless there is a significant insider threat. In addition, patrols can be divided into two categories – routine patrols and ambush patrols, where the latter are in response to tips from informants. Ambush patrols are costly, often requiring rangers to lie in wait for many hours for the possibility of catching a poacher. If not informed, the defender follows her routine patrol strategy $\mathbf{x}_0 = (x_1, \dots, x_n)$ with x_i denoting the probability that target i is covered. Naturally, under this assumption the defender should use a strategy \mathbf{x}_0 that is optimal against the QR model, which can be computed by following [26]. If informed she uses different

strategies $\mathbf{x}(V)$ based on the tip V . Assume that each attacker, if deciding to attack a target, will respond to the defender's strategy following a known behavioral model – the QR model. We define $\text{QR}(\mathbf{x}') := (q'_1, \dots, q'_n)$, where q'_i is the probability of attacking target i defined by

$$q'_i = \frac{e^{\lambda[x'_i P_i^a + (1-x'_i) R_i^a]}}{\sum_{j \in T} e^{\lambda[x'_j P_j^a + (1-x'_j) R_j^a]}}, \quad (1)$$

and \mathbf{x}' is the attacker's *subjective belief* of the coverage probabilities. In the above equation, $\lambda \geq 0$ is the precision parameter [12] fixed throughout the paper. We discuss the relaxation of the some of the assumptions mentioned above in Section 8.

3.1 Level- κ Response Model

Motivated by the costly ambush patrols and inspired by the cognitive hierarchy theory, we propose the level- κ response model as the attackers' behavior model.

When the informants' report intensities are negligible, the attackers are almost always faced with the routine patrol \mathbf{x}_0 . But when the informants' report intensities are not negligible, the attackers' behavior will change the marginal probability that a target is covered. Thus we assume that level-0 attackers just play the quantal response against the routine patrol \mathbf{x}_0 : $\mathbf{q}^0 = \text{QR}(\mathbf{x}_0)$. Then the defender will likely get informed with different tips V , and respond with $\mathbf{x}(V)$ accordingly. Over time, the attackers will learn about the change in the frequency that a target is covered. We denote the induced defender's marginal strategy at level 0 by $\hat{\mathbf{x}}^0 = \text{MS}(\mathbf{x}_0, \mathbf{x}, \mathbf{q}^0)$. After observing $\hat{\mathbf{x}}^0$ at level 0, level-1 attackers will update their strategies from \mathbf{q}^0 to $\mathbf{q}^1 = \text{QR}(\hat{\mathbf{x}}^0)$. Similarly, attackers at level κ ($0 < \kappa < \infty$) will use quantal response against the defender's marginal strategy at level $\kappa - 1$, i.e., $\mathbf{q}^\kappa = \text{QR}(\hat{\mathbf{x}}^{\kappa-1})$, where $\hat{\mathbf{x}}^{\kappa-1} = \text{MS}(\mathbf{x}_0, \mathbf{x}, \mathbf{q}^{\kappa-1})$. In Section 5, we also define level- ∞ attackers.

Denote by $\text{DefEU}(U)$ the defender's optimal utility when they recruit a set of informants U and use the optimal defending strategy. The key questions raised given this model are (i) how to recruit a set U of at most k informants and (ii) how to respond to the provided tips to maximize the expected $\text{DefEU}(U)$?

4 DEFENDING AGAINST LEVEL-0 ATTACKERS

In this section, we first tackle the case where all attackers are level-0 by providing complexity results and algorithms to find the optimal set of informants. Designing efficient algorithms to solve this computationally hard problem is particularly challenging due to the combinatorial nature of the tips and exponentially many possibilities of informant selections. Furthermore, in the general case, attackers are heterogeneous and we do not know which attackers will be reported, making it hard to compute $\text{DefEU}(U)$.

4.1 Complexity Results

Let $\mathbf{q}^0 = (q_1, \dots, q_n)$. Before presenting our complexity results, we first define some useful notations. Given the set of informants U and the tips $V = (V_1, \dots, V_n)$, we denote by $\tilde{p}_v(V_0)$ the probability of $v \in Y$ attacking a target given V_0 such that $V_0 = \bigcup_{i \in T} V_i$. We

can compute $\tilde{p}_v(V_0)$ with

$$\tilde{p}_v(V_0) = \begin{cases} 1 & v \in V_0 \\ \frac{(1-\tilde{w}_v)p_v}{(1-\tilde{w}_v)p_v+1-p_v} & v \in V \setminus V_0, \\ p_v & v \in Y \setminus V \end{cases}$$

where $\tilde{w}_v = 1 - \prod_{(u,v) \in E, u \in U} (1 - w_{uv})$ is the probability of v being reported given he attacks. Given V_0 and $t_i = |V_i|$ reported attacks on each target i , we compute the expected utility on i if i is covered with $\text{EU}_i^c(t_i, V_0) := \left(t_i + q_i \sum_{v \in Y \setminus V_0} \tilde{p}_v(V_0)\right) R_i^d$. We compute the expected utility if i is uncovered, $\text{EU}_i^u(t_i, V_0)$, similarly. Then, the expected gain of the target if covered can be written as $\text{EG}_i(t_i, V_0) := \text{EU}_i^c(t_i, V_0) - \text{EU}_i^u(t_i, V_0)$.

THEOREM 4.1. *When the defender is informed by informants U , the optimal allocation of defensive resources can be determined in $O(|Y| + n)$ time given the tips $V = (V_1, \dots, V_n)$.*

The proof of Theorem 4.1 can be found in the full version of the paper. Given tips from recruited informants, the defender can find the optimal resource allocation by greedily protecting the targets with the highest expected gains. However, the problem of computing the optimal set of informants is still hard.

THEOREM 4.2. *Computing the optimal set of informants to recruit is NP-Hard.*

The proof of Theorem 4.2 can be found in the full version of the paper, which focuses on a relatively simple case and constructs a reduction from the maximum coverage problem (MCP).

4.2 Finding the Optimal Set of Informants

In this subsection, we develop exact and heuristic informant selection algorithms to compute the optimal set of informants. To find the U that maximizes $\text{DefEU}(U)$, we first focus on computing $\text{DefEU}(U)$ by providing a dynamic programming-based algorithm and approximate algorithms.

4.2.1 Calculating $\text{DefEU}(U)$. Let DefEU_0 be the expected utility when using the optimal regular defending strategy against a single attack, which can be obtained by the algorithms introduced in [26]. Then $\text{DefEU}(U)$ can be explicitly written as

$$\begin{aligned} \text{DefEU}(U) &= \Pr[V_0 = \emptyset] \text{DefEU}_0 \\ &+ \Pr[V_0 \neq \emptyset] \mathbb{E} \left[\sum_{i \in [n]} \mathbf{x}_i(V) \text{EG}_i(t_i, V_0) + \text{EU}_i(t_i, V_0) \mid V_0 \neq \emptyset \right]. \end{aligned}$$

To directly compute $\text{DefEU}(U)$ from the above equation is formidable due to the exponential number of tips combinations. However, it is possible to reduce a significant amount of enumeration by handling the calculation carefully. We first develop an Enumeration and Dynamic Programming-based Algorithm (EDPA) to compute the exact $\text{DefEU}(U)$ as shown in Algorithm 1.

First, we compute the utility when the defender is not informed (lines 4-6). Then, we focus on calculating the total utility $\text{DefEU}'(U)$ in the case when the defender is informed. By the linearity of expectation, $\text{DefEU}'(U)$ can be computed as the summation of the expected utility obtained from all targets. Therefore, we focus on the calculation of the expected utility of a single target i . For each

target i , Algorithm 1 enumerates all possible types of tips (lines 2-7). We denote each type of tip by a tuple (t_i, V_0) , which encodes the set of reported attackers $V_0 \neq \emptyset$ and the number of reported attackers t_i targeting location i . The probability of receiving (t_i, V_0) can be written as

$$\Pr(t_i, V_0|U) = P_{V_0} \binom{|V_0|}{t_i} q_i^{t_i} (1 - q_i)^{|V_0| - t_i},$$

where

$$P_{V_0} = \prod_{v \in V_0} (\tilde{w}_v p_v) \prod_{v \in V \setminus V_0} (1 - \tilde{w}_v p_v) \quad (2)$$

is the probability of having V_0 being the set of reported attackers given U (line 3). Let $P_{i,r}$ be the probability of i being among the r targets with the highest expected gain given (t_i, V_0) and U (lines 12-13). For a given tip type (t_i, V_0) , the expected contribution to $\text{DefEU}'(U)$ of target i is

$$\begin{aligned} & \Pr(t_i, V_0|U) \cdot \text{EU}_i(t_i, V_0) + P_{V_0} \binom{|V_0|}{t_i} q_i^{t_i} \cdot P_{i,r} \text{EG}_i(t_i, V_0) \\ &= P_{V_0} \binom{|V_0|}{t_i} q_i^{t_i} \left((1 - q_i)^{|V_0| - t_i} \text{EU}_i(t_i, V_0) + P_{i,r} \text{EG}_i(t_i, V_0) \right). \end{aligned}$$

We can then compute $\text{DefEU}'(U)$ by summing over all possible $t_i, V_0 \neq \emptyset$.

The calculation of $P_{i,r}$ is all that remains. This can be done very efficiently via Algorithm 2, a dynamic programming-based calculation. Let $\{i_1, \dots, i_{n-1}\}$ denote the set of targets apart from i , i.e., $T \setminus \{i\}$ (line 1) and $y!$ $f(s, x, y)$ be the probability of having y reported attacks among the first s targets with x of the targets having expected gain higher than EG_i given the tips of type (t_i, V_0) . Therefore, $f(s, x, y)$ can be neatly written as

$$f(s, x, y) = \sum_{\substack{a_1 + \dots + a_s = y, \\ \sum_{j=1}^s \mathbb{1}_{[\text{EG}_{i_j}(a_j, V_0) > \text{EG}_i(t_i, V_0)]} = x}} \frac{q_{i_1}^{a_1} q_{i_2}^{a_2} \dots q_{i_s}^{a_s}}{a_1! a_2! \dots a_s!},$$

which can be calculated using dynamic programming (line 5-11). Computing $f(s, x, y)$ is done in a similar way by counting the number of s -partitions on integer y , where we also consider the constraint brought in by the limitation on the number of resources. To calculate $f(s, x, y)$, we enumerate a_s as \tilde{y} (line 6) and compare $\text{EG}_{i_s}(a_s, V_0)$ with $\text{EG}_i(t_i, V_0)$ (line 8). If $\text{EG}_{i_s}(a_s, V_0) > \text{EG}_i(t_i, V_0)$, we check the value of $f(s - 1, x - 1, y - \tilde{y})$ (line 9), otherwise check $f(s - 1, x, y - \tilde{y})$ (line 11). Thus, we have $P_{i,r} = (|V_0| - t_i)! \left(\sum_{x=0}^{r-1} f(s, x, |V_0| - t_i) \right)$. The time complexity for Algorithm 2 is $O(nr|Y|^2)$ and $O(2^{|Y|} n^2 r |Y|^3)$ for Algorithm 1.

Since EDPA runs in exponential time, we introduce approximation methods to estimate $\text{DefEU}(U)$. Let $\text{DefEU}(U, C)$ be the estimated defender's utility returned by Algorithm 1 if only subsets of reported attackers V_0 with $|V_0| < C$ are enumerated in line 2. We denote by C-TRUNCATED this approach of estimating $\text{DefEU}(U)$. Next, we show that $\text{DefEU}(U, C)$ is close to the exact $\text{DefEU}(U)$ when it is unlikely to have many attacks happening at the same time. Formally, assume that the expected number of attacks is bounded by a constant C' , that is $\sum_{v \in Y} p_v \leq C'$, $\text{DefEU}(U, C)$ for $C > C'$ is an estimation of $\text{DefEU}(U)$ with bounded error.

Algorithm 1 Calculate $\text{DefEU}(U)$

```

1: EU ← 0
2: for all possible sets of reported attackers  $V_0 \subseteq V$  do
3:    $P_{V_0} \leftarrow \prod_{v \in V_0} (\tilde{w}_v p_v) \prod_{v \in V \setminus V_0} (1 - \tilde{w}_v p_v)$ 
4:   if  $V_0 = \emptyset$  then
5:     EU = EU +  $P_{V_0} \sum_{v \in Y} \tilde{p}_v(V_0) \text{DefEU}_0$ 
6:     Continue to line 2
7:   for target  $i \in T$  and  $0 \leq t_i \leq |V_0|$  do
8:     Calculate  $f(\cdot)$  given  $|V_0|, i, t_i$ 
9:      $\text{EG}_i \leftarrow (t_i + q_i \sum_{v \in Y \setminus V_0} \tilde{p}_v(V_0)) (R_i^d - P_i^d)$ 
10:     $\text{EU}_i^u \leftarrow (t_i + q_i \sum_{v \in Y \setminus V_0} \tilde{p}_v(V_0)) P_i^d$ 
11:     $P_{i,r} \leftarrow (|V_0| - t_i)! \left( \sum_{x=0}^{r-1} f(s, x, |V_0| - t_i) \right)$ 
12:    EU = EU +  $P_{V_0} \binom{|V_0|}{t_i} q_i^{t_i} \cdot P_{i,r} \cdot \text{EG}_i$ 
13:    EU = EU +  $P_{V_0} \binom{|V_0|}{t_i} q_i^{t_i} (1 - q_i^{t_i}) \text{EU}_i^u$ 
14:  $\text{DefEU}(U) \leftarrow \text{EU}$ 

```

Algorithm 2 Calculate $f(\cdot)$ given $|V_0|, i, t_i$

```

1:  $\{i_1, \dots, i_{n-1}\} \leftarrow T \setminus \{i\}$ 
2:  $\text{EG}_i \leftarrow (t_i + q_i \sum_{v \in Y \setminus V_0} \tilde{p}_v(V_0)) (R_i^d - P_i^d)$ 
3: Initialize  $f(s, x, y) \leftarrow 0$  for all  $s, x, y$ 
4:  $f(0, 0, 0) \leftarrow 1$ 
5: for  $s$  in  $[1, n - 1]$ ,  $x$  in  $[0, \min(s, r)]$ ,  $y$  in  $[0, |V_0| - t_i]$  do
6:   for  $\tilde{y}$  in  $[0, y]$  do
7:      $\text{EG}_{i_s} \leftarrow (\tilde{y} + q_{i_s} \sum_{v \in Y \setminus V_0} \tilde{p}_v(V_0)) (R_{i_s}^d - P_{i_s}^d)$ 
8:     if  $\text{EG}_{i_s} > \text{EG}_i$  then
9:        $f(s, x, y) += \frac{q_{i_s}^{\tilde{y}}}{\tilde{y}!} f(s - 1, x - 1, y - \tilde{y})$ 
10:    else
11:       $f(s, x, y) += \frac{q_{i_s}^{\tilde{y}}}{\tilde{y}!} f(s - 1, x, y - \tilde{y})$ 

```

LEMMA 4.3. Assume that $\sum_{v \in Y} p_v \leq C'$ and $|P_i^d|, |R_i^d| \leq Q$, the error of estimation $|\text{DefEU}(U, C) - \text{DefEU}(U)|$ for $C > C'$ is at most:

$$Q \cdot e^{-2(C-C')^2/|Y|} \left(C + \frac{1}{1 - e^{-4(C-C')/|Y|}} \right).$$

The proof of Lemma 4.3 can be found in the full version. The time complexity of C-TRUNCATED is given by $O(n^2 r |Y|^{C+3})$.

However, for the case where $\sum_{v \in Y} p_v$ is large, we have to set C to be larger than $\sum_{v \in Y} p_v$ for C-TRUNCATED in order to obtain a high-quality solution; otherwise the error will become unbounded. To mitigate this limitation, we also propose an alternative sampling approach, T-SAMPLING, to estimate $\text{DefEU}(U)$ for general cases without restrictions on $\sum_{v \in Y} p_v$. Instead of enumerating all possible V_0 as EDPA does, in T-SAMPLING, we draw T i.i.d. samples of the set of reported attackers where each sample V_0 is drawn with probability P_{V_0} . T-SAMPLING takes the average of the expected defender's utility when having V_0 as the reported attackers over all samples as the estimation of $\text{DefEU}(U)$. We can sample V_0 as follows: (i) Let $V_0 = \emptyset$ initially; (ii) For each $v \in V$, add v to V_0 with probability $\tilde{w}_v p_v$; (iii) Return V_0 as a sample of the set of reported attackers. From Equation (2), the above sampling process is consistent with the distribution of V_0 . T-SAMPLING returns an estimation of $\text{DefEU}(U)$ in $O(Tn^2 r |Y|^3)$ time.

PROPOSITION 4.4. Let $\text{DefEU}^{(T)}(U)$ be the estimation of $\text{DefEU}(U)$ given by T-SAMPLING using T samples. We have: $\lim_{T \rightarrow \infty} \text{DefEU}^{(T)}(U) = \text{DefEU}(U)$

4.2.2 *Selecting Informants U .* Given the algorithms for computing $\text{DefEU}(U)$, a straightforward way of selecting informants is through enumeration (denoted as SELECT).

When using C-TRUNCATED as a subroutine to compute $\text{DefEU}(U)$, the solution quality of the selected set of informants is guaranteed by the following theorem.

THEOREM 4.5. Assume that $\sum_{v \in Y} p_v \leq C'$ and $|P_i^d|, |R_i^d| \leq Q$. Let U_{OPT} and U' be the optimal set of informants and the one chosen by C-TRUNCATED. Then for $C > C'$, the error $|\text{DefEU}(U_{\text{OPT}}) - \text{DefEU}(U')|$ can be bounded by:

$$2Q \cdot e^{-2(C-C')/|Y|} \left(C + \frac{1}{1 - e^{-4(C-C')/|Y|}} \right).$$

PROPOSITION 4.6. Using T-SAMPLING to estimate DefEU , the optimal set of informants can be found when $T \rightarrow \infty$.

Algorithm 3 Search(U')

- 1: **if** $|U'| = k$ **then**
 - 2: Update OPT with $(U', \text{DefEU}(U'))$
 - 3: **return**
 - 4: $u_1 \leftarrow \arg \max_{u \in X} \text{DefEU}(U' \cup \{u\})$
 - 5: $u_2 \leftarrow \arg \max_{u \in X \setminus \{u_1\}} \text{DefEU}(U' \cup \{u\})$
 - 6: Search($U' \cup \{u_1\}$), Search($U' \cup \{u_2\}$)
-

Based on existing results in submodular optimization [14], one may expect a greedy algorithm that step by step adds an informant that leads to the largest utility to work well. However, the set function $\text{DefEU}(U)$ in our problem violates submodularity (see the full version of the paper for an example) and such greedy algorithm will not guarantee an approximation ratio of $1 - 1/e$. Therefore, we propose GSA (Greedy-based Search Algorithm) for the selection of informants as shown in Algorithm 3. GSA starts by calling Search(\emptyset). While $|U'| < k$, Search(U') expands the current set of informants U' by adding u_1, u_2 to U' and recursing, where u_1 and u_2 are the two informants that give the largest marginal gain in DefEU (line 4-5); Otherwise, it updates the optimal solution with U' (line 1-3).

We identify a tractable case to conclude the section.

LEMMA 4.7. Given the set of recruited informants U , the defender's expected utility $\text{DefEU}(U)$ can be computed in polynomial time if $w_{uv} = 1 \forall (u, v) \in E$. When k is a constant, the optimal set of informants can be computed in polynomial time.

This represents the case where the informants have strong connections with a particular group of attackers and can get full access to their attack plans. We refer to the property of $w_{uv} = 1$ for all u, v as SISI (Strong Information Sharing Intensity). Denote by ASISI (Algorithm for SISI) the polynomial-time algorithm in Lemma 4.7. We provide more details about the SISI case in the full version of the paper.

We provide a summary of the time complexity of all algorithms for computing the optimal U in the full version.

5 DEFENDING AGAINST LEVEL- ∞ ATTACKERS

As discussed in Section 3.1, a level- κ attacker may keep adapting to the new marginal strategy formed by his current level of behavior. In this section, we first show in Theorem 5.1 that there exists a fixed-point strategy for the attacker in our level- κ response model, and then use that to define the level- ∞ attackers.

We formulate the problem of finding the optimal defender's strategy for this case as a mathematical program. However, such a program can be too large to solve. We propose a novel technique that reduces the program to a bi-level optimization problem, with both levels much more tractable.

THEOREM 5.1. Let $\Delta_n = \{\mathbf{q} \mid \mathbf{q} \in [0, 1]^n, \mathbf{1}^\top \mathbf{q} \leq 1\}$. Given defender's strategies \mathbf{x}_0 and $\mathbf{x}(\mathbf{V})$, there exists $\mathbf{q}^* \in \Delta_n$ such that $\mathbf{q}^* = \text{QR}(\text{MS}(\mathbf{x}_0, \mathbf{x}, \mathbf{q}^*))$.

PROOF. Since Δ_n is a compact convex set and $\text{QR}(\text{MS}(\mathbf{x}_0, \mathbf{x}, \mathbf{q}^*))$ is a continuous function of \mathbf{q} , by Brouwer fixed-point theorem, there exists $\mathbf{q}^* \in \Delta_n$ such that $\mathbf{q}^* = \text{QR}(\text{MS}(\mathbf{x}_0, \mathbf{x}, \mathbf{q}^*))$. \square

According to the definition of level- κ attackers, we have $\mathbf{q}^{\kappa+1} = \text{QR}(\text{MS}(\mathbf{x}_0, \mathbf{x}, \mathbf{q}^\kappa))$. Slightly generalizing the definition, we define a level- ∞ attacker as:

Definition 5.2 (level- ∞ attacker). Given the defender's strategies \mathbf{x}_0 and $\mathbf{x}(\mathbf{V})$, the strategy \mathbf{q} of a level- ∞ attacker satisfies $\mathbf{q} = \text{QR}(\text{MS}(\mathbf{x}_0, \mathbf{x}, \mathbf{q}))$.

Remark 5.3. Note that Definition 5.2 is not obtained by taking the limit of the level- κ definition, since such a limit may not even exist (see an example in the full version).

Remark 5.4. Although the level- ∞ attacker is defined through a fixed point argument, we still stick to the Stackelberg assumption: the defender leads and the attacker follows. Notice that in the equation $\mathbf{q} = \text{QR}(\text{MS}(\mathbf{x}_0, \mathbf{x}, \mathbf{q}))$, \mathbf{q} will only be defined after the defender commits to strategies \mathbf{x}_0 and \mathbf{x} . However, it is different from the standard Strong Stackelberg Equilibrium [7] in that the attacker is following a level- ∞ response model, as defined by the fixed point equation.

Also, as we will discuss in Section 7.1.3 on our experiments, when $r = n$, the defender's optimal strategy is not to use up all the available resources. This is clearly different from a Nash equilibrium, as the defender still has incentives to use more resources.

5.1 Convergence Condition for the Level- κ Response Model

We focus on the single-attacker case, where there are only n different types of tips. We use \mathbf{V}_i to denote the tips where the attacker is reported to attack target i . When the attacker is using strategy \mathbf{q} , the probability of receiving \mathbf{V}_i is $\Pr\{\mathbf{V}_i\} = wq_i$.

THEOREM 5.5. Let $\bar{x}_i = \max_j \{x_i(\mathbf{V}_j)\}$. In the single attacker case, if there exists constant $L \in [0, 1)$, such that $\bar{x}_i \leq \frac{L}{n\lambda(R_i^a - P_i^a)}$, $\forall i$, then level- κ agents converge to level- ∞ agents as κ approaches infinity.

The proof of Theorem 5.5 is omitted since it is immediate from the following lemma:

LEMMA 5.6. *In the single attacker case, if there exists constant $L \in [0, 1)$, such that $\bar{x}_i \leq \frac{L}{n\lambda(R_i^a - P_i^a)}$ for all i , then $g(\mathbf{q})$ is L -Lipschitz with respect to the L^1 -norm, i.e., $g(\mathbf{q})$ is a contraction.*

The proof of Lemma 5.6 can be found in the full version.

COROLLARY 5.7. *In the single attacker case, if there exists a constant $L \in [0, 1)$, such that $\frac{L}{n\lambda(R_i^a - P_i^a)} > 1, \forall i$, then level- κ agents converge to level- ∞ agents as κ goes to infinity.*

5.2 A Bi-Level Optimization for Solving the Optimal Defender's Strategy

In this section, we still consider the single attacker case and assume the defender has $r \geq 1$ resources. Clearly, the optimal set of informants should contain the ones with the highest information sharing intensities. It remains to compute the optimal strategies \mathbf{x}_0 and $\mathbf{x}(\mathbf{V})$. Given the optimal set of informants U^* , the probability of receiving a tip is $w = 1 - \prod_{u \in U^*} (1 - w_{u1})$. Let $\Pr\{\mathbf{V}\}$ be the probability of receiving tips \mathbf{V} , which depends \mathbf{q} . Let $\mathbf{x}(\mathbf{V}) = (x_1(\mathbf{V}), \dots, x_n(\mathbf{V}))$ be the defender strategy when receiving tips \mathbf{V} .

Let $\mathbf{q} = (q_1, \dots, q_n)$ be the strategy of the level- ∞ attacker. Given \mathbf{V} and the corresponding t_i 's, the expected number of attackers that are going to attack target i is $d_i = t_i + (1 - \sum_j t_j) \tilde{p}_v(\mathbf{0}) q_i$. Therefore, given $\hat{\mathbf{x}}$ we have the defender's expected utility $\text{DefEU}(\mathbf{x}_0, \mathbf{x})$ as

$$\text{DefEU}(\mathbf{x}_0, \mathbf{x}) = \sum_{\mathbf{V}} \Pr\{\mathbf{V}\} d_i \left[P_i^d + x_i(\mathbf{V}) (R_i^d - P_i^d) \right].$$

Then the problem of finding the optimal defender strategy can be formulated as the following mathematical program:

$$\max \text{DefEU}(\mathbf{x}_0, \mathbf{x}) \quad \text{s.t.} \quad \mathbf{q} = \text{QR}(\text{MS}(\mathbf{x}_0, \mathbf{x}, \mathbf{q})).$$

In the single-attacker case, we need n and n^2 variables to represent \mathbf{x}_0 and \mathbf{x} . We can use the QRI-MILP algorithm¹ to find the solution. However, this approach needs to solve a mixed integer program and does not scale well.

To tackle the problem, we focus on the defender's marginal strategy instead of the full strategy representation, and decompose the above program into a bi-level optimization problem.

Let $\hat{\mathbf{x}} = \text{MS}(\mathbf{x}_0, \mathbf{x}, \mathbf{q}) = \sum_{\mathbf{V}} \Pr\{\mathbf{V}\} \mathbf{x}(\mathbf{V})$, where we slightly abuse notation and use $\mathbf{V} = \emptyset$ to denote the case of receiving no tip, $\mathbf{x}(\emptyset)$ to denote \mathbf{x}_0 . The bi-level optimization method works as follows. At the inner level, we fix an arbitrary feasible $\hat{\mathbf{x}}$, and solve the following mathematical program:

$$\begin{aligned} \max \quad & \text{DefEU}(\hat{\mathbf{x}}) \\ \text{s.t.} \quad & \sum_{\mathbf{V}} \Pr\{\mathbf{V}\} \mathbf{x}(\mathbf{V}) = \hat{\mathbf{x}}, \quad \mathbf{q} = \text{QR}(\hat{\mathbf{x}}) \\ & \mathbf{1}^T \mathbf{x}(\mathbf{V}) \leq r, \quad \mathbf{x}(\mathbf{V}) \in [0, 1]^n, \quad \forall \mathbf{V} \end{aligned}$$

Since $\hat{\mathbf{x}}$ is fixed, \mathbf{q} and $\Pr\{\mathbf{V}\}$ are also fixed. Thus, the program above becomes a linear program, with $\mathbf{x}(\mathbf{V})$ as variables. We can always find a feasible solution to it by simply setting $\mathbf{x}(\mathbf{V}) = \hat{\mathbf{x}}, \forall \mathbf{V}$. Solving this linear program gives us the optimal defender's utility $\text{DefEU}(\hat{\mathbf{x}})$ for any possible $\hat{\mathbf{x}}$. To find the optimal defender strategy, we solve the outer-level optimization problem below:

$$\max \text{DefEU}(\hat{\mathbf{x}}) \quad \text{s.t.} \quad \hat{\mathbf{x}} \text{ is feasible.}$$

¹An algorithm that computes an approximate defender's optimal strategy against a variant of level-0 attackers who take into account the impact of informants when determining the target they attack. See the full version of the paper for more details.

Since the feasible region of $\hat{\mathbf{x}}$ is continuous, we can use any known algorithm (e.g., gradient descent) to solve the outer-level program. The inner-level linear program still suffers from the scalability problem. However, when there are multiple attackers, the optimal objective value can be well-approximated by simply sampling a subset of possible \mathbf{V} 's, or focusing only on the \mathbf{V} 's with the highest probabilities. For those \mathbf{V} 's that are not considered, we can always use \mathbf{x}_0 as the default strategy for $\mathbf{x}(\mathbf{V})$.

6 DEFENDING AGAINST INFORMANT-AWARE ATTACKERS

We now consider a variant of our model where attackers take into account the impact of informants when determining the target they attack. Specifically, we assume the attackers follow the QR behavior model but incorporate the probability of being discovered when determining their expected utility for attacking a target.² In this setting, the attackers' subjective belief \mathbf{x}' of the target coverage probability does not necessarily satisfy $\sum_i x'_i \leq r$. Consider the example of a single attacker and a single informant with report intensity 1. Assume that the defender has $r = 1$ and always protects the target being reported with probability 1. Then no matter which target the attacker chooses to attack, it will always be covered.

We focus on the single attacker case with $r \geq 1$. We first consider the problem of computing the optimal defender strategy when given the set of informants U and associated probability of receiving a tip $w = 1 - \prod_{u \in U} (1 - w_{u1})$. In the general case with multiple attackers, we will need to specify the defender strategy for each combination of tips received. However, when there is only one attacker, we can succinctly describe the defender strategy by their default strategy without tips, \mathbf{x} , and their probability of defending a location after receiving a tip for that location, \mathbf{z} . Then, under the QR adversary model, the probability q_i of the attacker targeting location i will be:

$$q_i = \frac{e^{\lambda\{(1-w)x_i + wz_i\}P_i^a + [1-(1-w)x_i - wz_i]R_i^a}}{\sum_{j \in \mathcal{T}} e^{\lambda\{(1-w)x_j + wz_j\}P_j^a + [1-(1-w)x_j - wz_j]R_j^a}}.$$

This leads to the following optimization problem, QRI, to compute the optimal defender strategy:

$$\begin{aligned} \max_{\mathbf{x}, \mathbf{y}, \mathbf{a}} \quad & \frac{\sum_{i \in \mathcal{T}} e^{\lambda R_i^a} e^{-\lambda(R_i^a - P_i^a)y_i} \left[(R_i^d - P_i^d)y_i + P_i^d \right]}{\sum_{i \in \mathcal{T}} e^{\lambda R_i^a} e^{-\lambda(R_i^a - P_i^a)y_i}} \\ \text{subject to} \quad & y_i = (1-w)x_i + wz_i, \quad \forall i \in \mathcal{T} \quad (3) \\ & \sum_{i \in \mathcal{T}} x_i \leq r \quad (4) \\ & 0 \leq x_i, z_i \leq 1, \quad \forall i \in \mathcal{T} \quad (5) \end{aligned}$$

We can compute the optimal defender strategy by adapting the approach used in the PASAQ algorithm [26]. The description of the algorithm can be found in the full version of the paper.

7 EXPERIMENT

In this section, we demonstrate the effectiveness of our proposed algorithms through extensive experiments. In our experiments, all

²Consider attackers that have had experience playing against the defender. Over time, the attacker might start to consider their expected utility in practice, which is affected by informants.

reported results are averaged over 30 randomly generated game instances. We provide details about how we generate game instances and parameters in the full version. Unless specified otherwise, all game instances are generated in this way.

7.1 Experimental Results

We compare the scalability and the solution quality of SELECT using EDPA, C-TRUNCATED, T-SAMPLING to obtain DefEU and GSA for different settings of the problems against level-0 attackers.

First, we test the case where $\sum_{v \in Y} p_v < 3$. We set $|X| = 6$, $k = 4$, $n = 8$, $r = 3$ and enumerate $|Y|$ from 2 to 16. The results are shown in Figure 1a. We also include GREEDY as a baseline that always chooses the informants that maximizes the probability of receiving tips. We can see that T-SAMPLING performs the best in term of runtime, but fails to provide high-quality solutions. While C-TRUNCATED is slower than T-SAMPLING, it performs the best with no error on all test cases. However, when there is no restriction on $\sum_{v \in Y} p_v$, as shown in Figure 1b, C-TRUNCATED performs badly, even worse than GREEDY for large $|Y|$, while T-SAMPLING performs a lot better and GSA performs the best. We also fix $|X| = 7$, $|Y| = 10$, $k = 3$, $r = 5$ and change the number of targets n from 5 to 25 for $\sum_{v \in Y} p_v < 3$. The results are shown in Figure 1c. GSA is the fastest but provides slightly worse solutions than C-TRUNCATED does. The runtime of GREEDY is less than 0.3s for all instances tested.

We then perform a case study to show the trade-off between the optimal number of resources to allocate and the optimal number of informants to recruit with budget constraints when defending against level-0 attackers. We set $|X| = |Y| = n = 6$ and generate an instance of the game. We set the cost of allocating one defensive resource $C_r = 3$ and the cost of hiring one informant $C_i = 1$. Given a budget B , we can recruit k informants and allocate r resources when $k \cdot C_i + r \cdot C_r \leq B$. The trade-off between the optimal k and r is shown in Figure 1d. In the same instance, we study how the defender’s utility would change by increasing the number of recruited informants with fixed r . Given a fixed number of resources, the defender should recruit as many informants as possible. We can also see that assuming the defender can acquire sufficient resources, the importance of recruiting additional informants is diminished. This result provides useful guidance to defenders such as conservation agencies in allocating their budget and recruiting informants.

We run additional experiments for the SISI case and do a case study to show the errors of the estimations for all $U \subseteq X$ on 2 instances. The results can be found in the full version.

7.1.1 Level-0 vs. Level- ∞ attackers. We set $|X| = n = 6$, $|Y| = p_1 = 1$, and G_S to be fully connected. We set $r = 2, 4, 6$ and vary k from 1 . . . 6. We first fix the defender’s strategy to the one against level-0 attackers and compare the utility achieved by the defender when defending against a level-0 attacker and a level- ∞ attacker. We show how the defender utility varies with the number of informants and defensive resources in Figure 1e. On average, we see that the defender utility against a level- ∞ attacker is lower than that against a level-0 attacker. We also show the utility of the defender using her optimal strategy against a level- ∞ attacker. We can see that when facing a level- ∞ attacker, the defender utility when using the optimal strategy is higher by a margin than using the one against level-0 attackers.

7.1.2 Level-0 vs. informant-aware defenders. We set $|X| = n = 6$, $|Y| = p_1 = 1$, and G_S to be fully connected. We vary r from 1 . . . 6 and k from 0 . . . 6. We assume that the defender recruits the k informants with the highest information sharing intensity w_{u1} . The optimal defender strategy against the informant-aware attacker case is found using QRI-MILP. The defender strategy against the level-0 attacker case is computed using PASAQ [26]. The defender utility against the level-0 attacker is found by first computing q_i , DefEU₀ and then using the results to compute DefEU(U).

In Figure 1g, we show how the defender utility in the two cases varies with the number of informants and defensive resources. On average, we see that the defender utility is marginally higher against the level-0 attacker than against the informant-aware attacker, particularly when the defender has either very few or very many defensive resources. We also compare the defender’s utility of the level-0 defender (defending against level-0 attackers) and the informant-aware defender (defending against informant-aware attackers). The results of these experiments can be found in the full version of the paper.

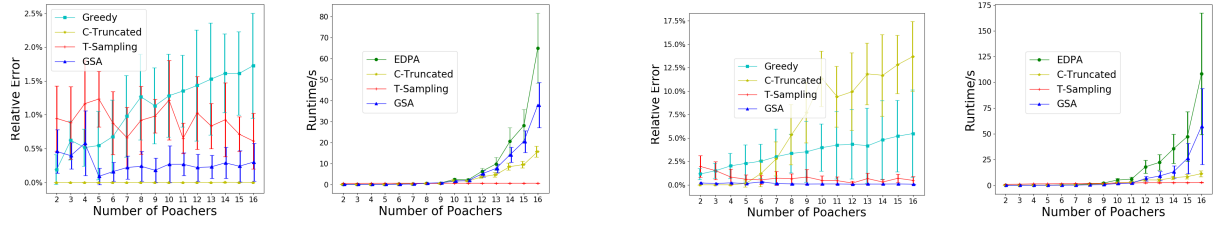
7.1.3 Comparison between the Bi-Level Algorithm and QRI-MILP. We empirically compare the bi-level optimization algorithm with QRI-MILP. We set $|X| = n = 6$, $|Y| = p_1 = 1$, and G_S to be fully connected. We vary r from 1 . . . 6 and k from 0 . . . 6.

In both cases, we assume that the defender recruits the k informants with the highest information sharing intensity w_{u1} . The results are shown in Figure 1f. In general, our bi-level algorithm gives higher expected defender utilities than the QRI-MILP algorithm, except when $r = 1$. Our results show that both increasing the number of resources and hiring more informants increase the defender’s utility. However, as the number of resources (r) increases, the utility gain from hiring more informants diminishes.

Intuitively, if the number of resources equals the number of targets, the defender should always cover all the targets. Interestingly, during our experiments, we observed that in this case, the optimal defender strategy may not always use all his resources to cover all the targets. The reason is that in a general sum game, by decreasing the probability of protecting a certain target on purpose, the defender can lure the attacker into attacking the target more frequently, and thus increase his expected utility. Such strategies can be found in real-world wildlife protections where the patrollers may sometimes deliberately ignore the tips. This is also reflected in our bi-level algorithm. If the defender always uses all his resources, then both the defender’s and the attacker’s strategies are fixed, and hiring more informants does not increase the defender’s expected utility. But if the defender strategy does not always use all his resources, then hiring more informants could help (see the bi-level algorithm for the $r = 6$ case in Figure 1f).

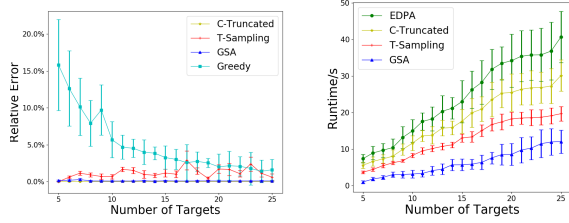
8 DISCUSSION AND CONCLUSION

In this paper, we introduced a novel two-stage security game model and a multi-level QR behavioral model that incorporated community engagement. We provided complexity results, developed algorithms to find (sub-)optimal groups of informants to recruit against level-0 attackers and evaluated the algorithms through extensive experiments. Our results also generalize to the case where informants have heterogeneous recruitment costs and to different kinds of

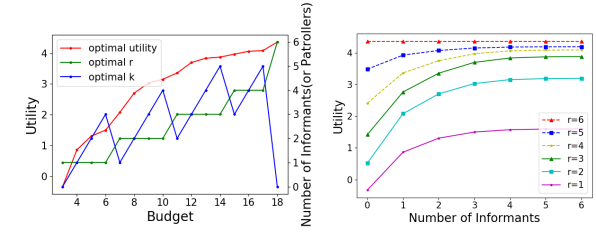


(a) Runtime and solution quality increasing $|Y|$ with $\sum p_v < 3$.

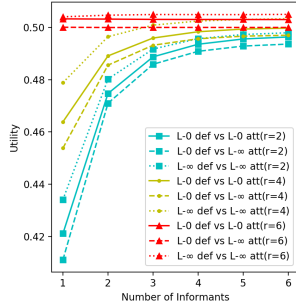
(b) Runtime and solution quality increasing $|Y|$ for General Cases.



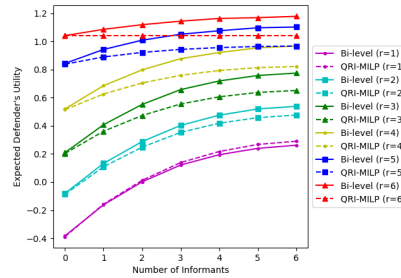
(c) Runtime and solution quality increasing n with $\sum p_v < 3$.



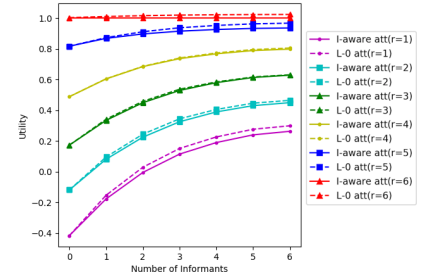
(d) Trade-off between r and k , and increase of utility with fixed r ($|X| = 6, |Y| = 6, n = 6$).



(e) Comparison between the defender utility against level-0 and level- ∞ attackers. “L- ∞ /0 def” means that the defender uses the optimal strategy against a Level- ∞ /0 attacker.



(f) Comparison between the bi-level optimization algorithm and QRI-MILP.



(g) Comparison between defender utility against level-0 and informant-aware attackers.

Figure 1: Experimental Results.

attacker response models, such as SUQR model [15], which can be done by calculating the attacker’s response correspondingly. Also, in the full version of the paper, we show how to extend our algorithms to defend against level- κ ($\kappa < \infty$) attackers. In Section 5, we defined a more powerful type of attacker that could respond to the marginal strategy and developed a bi-level optimization algorithm to find the optimal defender’s strategy in this case.

In the anti-poaching domain, some conservation site managers utilize the so-called “intelligence” operations that rely on informants in nearby villages to alert rangers when they know the poachers’ plans in advance. The deployment of the work relies on the site manager to provide their understanding of the social connections among community members. The edges and parameters of the bipartite graph in our model can be extracted from a local social media application or data collected by site managers. Recruiting and training reliable informants is costly and managers may only afford a limited number of them. Our model and solution can help

the managers efficiently recruit informants, make the best use of tips and evaluate the trade-off between allocating budget to hiring rangers and recruiting informants in a timely fashion.

For future work, one can use historical records as training data and learn the attackers’ behavior in different domains. It is worth considering the case where the informants could provide inaccurate tips or other types of tips, e.g., some subset of targets will be attacked instead of a single location. Another important future direction is to model the informants as strategic agents who have their own utility structures and to design reward mechanism to elicit true information and maximize the defender’s utility.

ACKNOWLEDGEMENT

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

REFERENCES

- [1] Nicola Basilico, Andrea Celli, Giuseppe De Nittis, and Nicola Gatti. 2017. Coordinating multiple defensive resources in patrolling games with alarm systems. In *AAMAS'17*. 678–686.
- [2] Rosaleen Duffy, Freya AV St John, Bram Büscher, and DAN Brockington. 2015. The militarization of anti-poaching: undermining long term goals? *Environmental Conservation* 42, 4 (2015), 345–348.
- [3] 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. 2017. PAWS - A Deployed Game-Theoretic Application to Combat Poaching. *AI Magazine* (2017). <http://www.aaai.org/ojs/index.php/aimagazine/article/view/2710>
- [4] Charlotte Gill, David Weisburd, Cody W Telep, and Trevor Bennett. 2014. Community-oriented policing to reduce crime, disorder and fear and increase satisfaction and legitimacy among citizens: A systematic review. *Journal of Experimental Criminology* (2014).
- [5] Qingyu Guo, Boyuan An, Branislav Bosansky, and Christopher Kiekintveld. 2017. Comparing strategic secrecy and Stackelberg commitment in security games. In *IJCAI-17*.
- [6] Manish Jain, Jason Tsai, James Pita, Christopher Kiekintveld, Shyamsunder Rathi, Milind Tambe, and Fernando Ordóñez. 2010. Software assistants for randomized patrol planning for the lax airport police and the federal air marshal service. *Interfaces* (2010).
- [7] Dmytro Korzhuk, Zhengyu Yin, Christopher Kiekintveld, Vincent Conitzer, and Milind Tambe. 2011. Stackelberg vs. Nash in security games: An extended investigation of interchangeability, equivalence, and uniqueness. *Journal of Artificial Intelligence Research* 41 (2011), 297–327.
- [8] Bertrand Le Gallic and Anthony Cox. 2006. An economic analysis of illegal, unreported and unregulated (IUU) fishing: Key drivers and possible solutions. *Marine Policy* 30, 6 (2006), 689–695.
- [9] N Leader-Williams and EJ Milner-Gulland. 1993. Policies for the enforcement of wildlife laws: the balance between detection and penalties in Luangwa Valley, Zambia. *Conservation Biology* 7, 3 (1993), 611–617.
- [10] Matthew Linkie, Deborah J. Martyr, Abishek Harihar, Dian Risdianto, Rudijanta T. Nugraha, Maryati, Nigel Leader-Williams, and Wai-Ming Wong. 2015. EDITOR'S CHOICE: Safeguarding Sumatran tigers: evaluating effectiveness of law enforcement patrols and local informant networks. *Journal of Applied Ecology* (2015).
- [11] Xiaobo Ma, Yihui He, Xiapu Luo, Jianfeng Li, Mengchen Zhao, Bo An, and Xiaohong Guan. 2018. Camera Placement Based on Vehicle Traffic for Better City Security Surveillance. *IEEE Intelligent Systems* 33, 4 (Jul 2018), 49–61. <https://doi.org/10.1109/mis.2018.223110904>
- [12] Richard D McKelvey and Thomas R Palfrey. 1995. Quantal response equilibria for normal form games. *Games and economic behavior* (1995).
- [13] William D Moreto. 2015. Introducing intelligence-led conservation: bridging crime and conservation science. *Crime Science* 4, 1 (2015), 15.
- [14] George L Nemhauser, Laurence A Wolsey, and Marshall L Fisher. 1978. An analysis of approximations for maximizing submodular set functions—I. *Mathematical programming* 14, 1 (1978), 265–294.
- [15] Thanh Hong Nguyen, Rong Yang, Amos Azaria, Sarit Kraus, and Milind Tambe. 2013. Analyzing the Effectiveness of Adversary Modeling in Security Games.. In *AAAI*.
- [16] James Pita, Manish Jain, Janusz Marecki, Fernando Ordóñez, Christopher Portway, Milind Tambe, Craig Western, Praveen Paruchuri, and Sarit Kraus. 2008. Deployed ARMOR protection: the application of a game theoretic model for security at the Los Angeles International Airport. In *AAMAS: industrial track*.
- [17] Ariel Rosenfeld and Sarit Kraus. 2017. When Security Games Hit Traffic: Optimal Traffic Enforcement Under One Sided Uncertainty.. In *IJCAI*. 3814–3822.
- [18] Aaron Schlenker, Omkar Thakoor, Haifeng Xu, Fei Fang, Milind Tambe, Long Tran-Thanh, Phebe Vayanos, and Yevgeniy Vorobeychik. 2018. Deceiving Cyber Adversaries: A Game Theoretic Approach. In *AAMAS*.
- [19] Martin B Short, Ashley B Pitcher, and Maria R D'Orsogna. 2013. External conversions of player strategy in an evolutionary game: A cost-benefit analysis through optimal control. *European Journal of Applied Mathematics* 24, 1 (2013), 131–159.
- [20] MLR Smith and Jasper Humphreys. 2015. *The Poaching Paradox: Why South Africa's 'Rhino Wars' Shine a Harsh Spotlight on Security and Conservation*. Ashgate Publishing Company.
- [21] Milind Tambe. 2011. *Security and game theory: algorithms, deployed systems, lessons learned*. Cambridge University Press.
- [22] Rebecca Tublitz and Sarah Lawrence. 2014. The Fitness Improvement Training Zone Program. (2014).
- [23] Xinrun Wang, Bo An, Martin Strobel, and Fookwai Kong. 2018. Catching Captain Jack: Efficient Time and Space Dependent Patrols to Combat Oil-Siphoning in International Waters. (2018). <https://www.aaai.org/ocs/index.php/AAAI/AAAI18/paper/view/16312>
- [24] James R Wright and Kevin Leyton-Brown. 2014. Level-0 meta-models for predicting human behavior in games. In *Proceedings of the fifteenth ACM conference on Economics and computation*. ACM, 857–874.
- [25] WWF. 2015. Developing an approach to community-based crime prevention. <http://zeropoaching.org/pdfs/Community-based-crime%20prevention-strategies.pdf>. (2015).
- [26] Rong Yang, Fernando Ordóñez, and Milind Tambe. 2012. Computing optimal strategy against quantal response in security games. In *AAMAS*.