When Can the Defender Effectively Deceive Attackers in Security Games?*

Thanh Nguyen[1] **and Haifeng Xu**[2]

[1]Department of Computer and Information Science , University of Oregon, USA
[2]Department of Computer Science, University of Virginia, USA
thanhhng@cs.uoregon.edu, hx4ad@virginia.edu

Abstract

This paper studies *defender patrol deception* in general Stackelberg security games (SSGs), where a defender attempts to alter the attacker's perception of the defender's patrolling intensity so as to influence the attacker's decision making. We are interested in understanding the *complexity* and *effectiveness* of optimal defender deception under different attacker behavior models. Specifically, we consider three different attacker strategies of response (to the defender's deception) with increasing sophistication, and design efficient polynomial-time algorithms to compute the equilibrium for each. Moreover, we prove formal separation for the effectiveness of patrol deception when facing an attacker of increasing sophistication, until it becomes even harmful to the defender when facing the most intelligent attacker we consider. Our results shed light on *when* and *how* deception should be used in SSGs. We conduct extensive experiments to illustrate our theoretical results in various game settings.

Introduction

SSGs are a well-known class of games which are used in various real-world security domains (Basilico, Gatti, and Amigoni 2009; Letchford and Vorobeychik 2011; Tambe 2011). This paper investigates the complexity and effectiveness of a basic patrol deception strategy of the defender in SSGs, where the defender can conduct patrols in a deceptive manner to mislead the attacker. In particular, we assume that the defender can strategically *disguise* or *exaggerate* (to a limited extent) the protection status at each target to fool the attacker. One of our major motivating domains, among others, is the wildlife conservation where SSG models have been deployed (Fang et al. 2016) and where deception approaches have been considered in previous work (Bondi et al. 2020; Xu et al. 2018). For example, rangers can hide themselves from being seen by poachers (Guo et al. 2017). They can also hire local villagers or use conservation drones to make the patrolling appear more intense (Bondi et al. 2020).

Motivated by these practices, we study a new and basic deception game model, coined Security Games with Deceptive Coverage (SeGDCs), which augments an *arbitrary* security game with the additional deception component. In SeGDCs, the defender can make the attacker observe a different (deceptive) patrolling intensity rather than the true ones. Concretely, the defender can deceit the attacker's perception of the protection probability of each target, and the deceptive probability lies within an interval, containing the true protection probability. The interval bounds at each target capture the defender's capability limit in influencing the attacker's perception at that target. The defender's deception capability at each target is usually determined by the target's own characteristics such as terrain features and thus can be considered to be independent of other targets'. Under this scenario, we aim to answer the following research question:

> *How do the* complexity *and* effectiveness *of the optimal deception evolve as the attacker becomes increasingly more sophisticated?*

We consider three natural types of attack strategies, with increasing sophistication: (i) *Ignorant attack strategy* — the attacker is unaware of defender's deception and simply responds to the deceptive strategy; (ii) *Maximin attack strategy* — the attacker is aware of the defender's deception and addresses it via a robust maximin approach; and (iii) *Equilibrium attack strategy* — the attacker is aware of and reasons about the defender's deception, leading to a sequential equilibrium of the game. Our models of attacker strategies are motivated by real-world domains such as wildlife protection in which different types of poachers with different levels of sophistication exist, ranging from local villagers to professional groups (Van Uhm 2016).

Within these three natural attack models, we provide a complete set of answers to the aforementioned question. We show that the optimal equilibrium deception strategy can be computed in polynomial time whenever the standard SSG (without deception) can be solved efficiently. This conveys a clean message that *deception against more sophisticated attackers is "no harder"*. Furthermore, we show that the defender deception becomes *strictly* less powerful, in terms of securing the defender's expected utility, as the attacker becomes more sophisticated, until to the point where deception even leads to a decrease in the defender's utility. Indeed, deception may cause the defender to lose her commitment power (as the attacker cannot observe defender's underlying true strategy any more due to her deception) and leads to a harmful outcome when facing a sophisticated attacker. To our knowledge, this is the first quantitative study on un-

*The two authors contribute equally.

derstanding the effectiveness of deception in SSGs against attackers of different levels of sophistication. This illustrates another conceptual message: *deception may be a double-edged sword for the defender, and needs to be carefully used.*

Finally, we conduct extensive experiments to evaluate the impact of the defender's deception. Our empirical results align with our theoretical findings, which show that the defender obtains a significant benefit while the attacker suffers a significant loss due to the defender's deception when the attacker plays the Ignorant or Maximin strategies. Conversely, the defender suffers a substantial loss when the attacker chooses the equilibrium-based strategy.

Additional Discussions on Related Work

For the study of deception in general security domains, we refer curious readers to a recent survey (Fraunholz et al. 2018) and references therein. Within SSGs, many recent works have investigated the problem of deception from various perspectives. Such deception typically originates from the asymmetric information among players (Guo et al. 2017; Rabinovich et al. 2015; Sinha et al. 2018; Xu et al. 2015). On one hand, the attacker can strategically change his attack behavior to mislead the defender given that the defender is uncertain about some of the attacker's characteristics (Gan et al. 2019; Nguyen and Xu 2019; Nguyen et al. 2019). On the other hand, the defender can also strategically reveal or disguise her information to the attacker (Guo et al. 2017; Rabinovich et al. 2015; Xu et al. 2015; Zhuang, Bier, and Alagoz 2010). Our work belongs to this second line of research, i.e., the defender's deception. However, different from previous works which all focus on designing the optimal deception strategy for a particular model, our work seeks to understand the power of deception under different attacker response models.

Among existing works, the work in (Guo et al. 2017) is the closest to ours in the sense that they also consider a defender who can disguise security resources. However, they take a Bayesian approach and assume that there is a prior probability distribution over the defender's types (each corresponds to a particular number of resources). The defender then commits to a signaling mechanism to influence the attacker's belief over the types. However, the defender in our setting does not have the power to commit to signaling schemes, and thus deception and the attacker's inference are simultaneously. Such simultaneous move makes the defender to lose commitment power to mixed strategies and is the intrinsic reason for why deception can sometimes be harmful to the defender.

Preliminaries

Stackelberg security games. (Tambe 2011) In SSGs, a *defender* (she) allocates limited number of security resources to protect a set of critical targets $\{1,,2,\dots,n\} = [n]$ from the attack of an *attacker* (he). When the attacker attacks a target i, if the defender is protecting i, the attacker receives a penalty P_i^a while the defender gets a reward R_i^a. Conversely, if i is not protected, the attacker obtains a reward $R_i^a (> P_i^a)$ while the defender receives a penalty $P_i^d (< R_i^d)$.

A pure strategy of the defender can be viewed as a *subset* of targets that can be covered by her security resources. We consider generic SSGs with *arbitrary resource allocation constraints*. We use a binary vector $\boldsymbol{e} \in \{0,1\}^{[n]}$ to represent a defender pure strategy where $e_i = 1$ if and only if target i is protected by $\boldsymbol{e}$. Let $\mathcal{E} \subset \{0,1\}^{[n]}$ denote a set of all these strategies. A defender mixed strategy is a distribution over $\mathcal{E}$. The set $\mathcal{E}$ is typically exponentially large. Nevertheless, previous work shows that the optimal defender strategy may still be computed efficiently through optimization techniques when $\mathcal{E}$ is nicely structured (Tambe 2011).

An important concept in SSGs is the *marginal protection probability*. Let $\{p_{\boldsymbol{e}}\}_{\boldsymbol{e}\in\mathcal{E}}$ denote any mixed defender strategy where $p_{\boldsymbol{e}}$ is the probability of taking pure strategy $\boldsymbol{e}$. Then $\mathbf{x} = \sum_{\boldsymbol{e}\in\mathcal{E}} p_{\boldsymbol{e}} \cdot \boldsymbol{e}$ is the vector of marginal probabilities where x_i is the marginal probability that target i is protected by mixed strategy $\{p_{\boldsymbol{e}}\}_{\boldsymbol{e}\in\mathcal{E}}$. Let

$$\mathbf{X} = \{\mathbf{x} = \sum_{\boldsymbol{e}\in\mathcal{E}} p_{\boldsymbol{e}} \cdot \boldsymbol{e} : \sum_{\boldsymbol{e}\in\mathcal{E}} p_{\boldsymbol{e}} = 1, p_{\boldsymbol{e}} \geq 0, \forall \boldsymbol{e}\} \quad (1)$$

denote the set of all possible marginal probability vectors. In SSGs, both players' utilities only depend on the marginal probabilities. Therefore, for convenience, we will use $\mathbf{x}$ as a defender mixed strategy from the feasible set $\mathbf{X}$. The rational attacker is assumed to be able to observe the defender's mixed strategy and then choose one target to attack. Given $\mathbf{x}$, the defender's and attacker's expected utility when the attacker attacks a target i is computed as follows:

$$U_i^d(x_i) = x_i(R_i^d - P_i^d) + P_i^d$$
$$U_i^a(x_i) = x_i(P_i^a - R_i^a) + R_i^a$$

Strong Stackelberg equilibrium (SSE). Given $\mathbf{x}$, we denote by $\Gamma(\mathbf{x}) = \{i : U_i^a(x_i) \geq U_j^a(x_j), \forall j\}$ the set of targets that maximize the attacker's expected utilities. Formally, a pair of strategies $(\mathbf{x}^0, i^0)$ forms an SSE if: (i) The attacker plays a best response to $\mathbf{x}^0$, i.e., $i^0 \in \Gamma(\mathbf{x}^0)$, and breaks ties in favor of the defender, i.e., $U_{i^0}^d(x_{i^0}^0) \geq U_i^d(x_i^0), \forall i \in \Gamma(\mathbf{x}^0)$; and (ii) $\mathbf{x}^0$ is optimal for the defender: $U_{i^0}^d(x_{i^0}^0) \geq U_i^d(x_i), \forall i \in \Gamma(\mathbf{x})$ and $\mathbf{x} \in \mathbf{X}$.

Nash equilibrium for simultaneous-move security games. We denote by $\mathbf{q} = (q_1, q_2, \dots, q_n)$ the attacker's mixed strategy where $q_i \in [0,1]$ is the probability target i is attacked. We also denote by $\mathbf{Q} = \{\mathbf{q} : q_i \in [0,1] \text{ and } \sum_i q_i = 1\}$ the set of these strategies. Formally, a pair $(\mathbf{x}^0, \mathbf{q}^0)$ forms a Nash equilibrium if: (i) The defender plays an optimal strategy, given $\mathbf{q}^0$: $\sum_i q_i^0 U_i^d(x_i^0) \geq \sum_i q_i^0 U_i^d(x_i), \forall \mathbf{x} \in \mathbf{X}$; and (ii) The attacker plays an optimal strategy, given $\mathbf{x}^0$: $\sum_i q_i^0 U_i^a(x_i^0) \geq \sum_i q_i U_i^a(x_i^0), \forall \mathbf{q} \in \mathbf{Q}$.

The Model of Patrol Deception Games

This work studies the problem of defender patrol deception in which the defender can alter the attacker's perception of the defender's strategy. Essentially, for each strategy, $\mathbf{x}$, the defender can make the attacker observe a *different* strategy $\mathbf{c}$. Formally, $x_i - \alpha_i \leq c_i \leq x_i + \beta_i$ and $c_i \in [0,1]$ for all targets i. The parameter $\alpha_i \in [0,1)$ represents the defender's

ability to *disguise* resources and $\beta_i \in [0, 1)$ captures her ability to *exaggerate* resources at i. This include a special case where the defender can only disguise or only exaggerate the coverage. It is usually difficult to significantly alter attacker's perception, so $\{\alpha_i, \beta_i\}_i$ can be very small. Nevertheless, we show later that a tiny α_i can at times be more powerful than having $(n-1)$ additional resources.

Given $\mathbf{x}$, we denote by $\Omega(\mathbf{x}) = \{\mathbf{c} : \max\{x_i - \alpha_i, 0\} \leq c_i \leq \min\{x_i + \beta_i, 1\}\})$ the set of possible *implementable* deceptive strategies. Here, we assumed that the defender's deception capabilities at each target are independent, i.e., c_i only depends on x_i, α_i, β_i. These α_i, β_i are typically small and are determined by the characteristics of the target. For example, in wildlife conservation areas with more curvy or hilly terrain, it is easier for the rangers to hide. For areas with more nearby local villagers and bushy terrains, it is easier to get villagers or to use drones to exaggerate the coverage intensity. Moreover, despite such flexibility of picking a deception strategy, we show later that the defender's deception may nevertheless become harmful to herself when playing against a sufficiently sophisticated attacker. Therefore, a more constrained deception strategy will only do worse.

Definition 1 (Defender strategy)**.** *The defender strategy in our* SeGDC *model is a pair* $(\mathbf{x}, \mathbf{c})$ *where* $\mathbf{x}$ *is the defender's actual strategy and* $\mathbf{c}$ *is the deceptive strategy that the defender intends the attacker to perceive. Let* $\Omega = \{(\mathbf{x}, \mathbf{c}) : \mathbf{x} \in \mathbf{X}, \mathbf{c} \in \Omega(\mathbf{X})\}$ *be the set of all feasible defender strategies.*

One interesting property of our deception game is that it has both "Stackelberg component" and "Nash component". In particular, $\mathbf{c}$ is always "observed" (more precisely, *perceived*) by the attacker, which consequently leads to the defender's commitment to $\mathbf{c}$. However, the mapping from true coverage $\mathbf{x}$ to $\mathbf{c}$, which is precisely the deception strategy, is *unobservable* by the attacker.[1] This results in a simultaneous move between the leader's *deception strategy* and the *attacker's response* to $\mathbf{c}$. Our study of the defender's patrol deception will focus on: (i) computing the optimal defender strategy; and (ii) investigating the benefit and loss of the players as a result of the defender's patrol deception.

Deception Against Ignorant Attacks

Our study starts with the basic situation where the attacker is unaware of the defender's deception, and thus would believe the deceptive protection coverage $\mathbf{c}$. We refer to such an attacker as *ignorant attacker*. Note that the attacker here does not respond to true $\mathbf{x}$ simply because he does not know that the coverage he observed is deceptive (thus the term "ignorance"). The ignorant attacker will respond optimally to his observation $\mathbf{c}$. Examples of ignorant attackers include native poachers in conservation domains who are local agricultural villagers and tend to conduct simple hunting activities (Van Uhm 2016). We first study the algorithmic question of computing the equilibrium defender strategy in SeGDCs against an ignorant attacker (Theorem 1).

[1] An alternative way is to assume the defender can commit to a mapping from $\mathbf{x}$ to $\mathbf{c}$. However, this assumption requires the attacker to observe the entire mapping, which is unrealistic.

Theorem 1 (Efficient Solvability)**.** *If the standard* SSE *(*without *deception) can be computed in polynomial time, then the optimal defender strategy against an ignorant attacker in* SeGDCs *can also be computed in polynomial time.*

All formal proofs in this paper are deferred to the appendix due to space limit. Proof sketches are provided for a representative subset of the results.

Next, we investigate the power of the defender's deception. We compare the following two situations, which represent two different types of defender capabilities:

1. **Situation 1**: The defender has *unlimited* security resources to implement any marginal probabilities in $[0, 1]^n$, but *cannot* deceive the attacker about protection.
2. **Situation 2**: The defender can hide at most an ϵ amount of her protection at a target.[2]

One might think that in this situation, the defender can do a "perfect job" by simply fully covering all the targets. Surprisingly, we show that *deception* — particularly, the capability of hiding security resources — can yield strictly better utility than such perfect protection, illustrating that the ability of deception can be more powerful than getting more resources. Intuitively, this comes from the following insight: perfect protection may not be ideal as that may "scare away" the attacker from those targets, of which the payoffs are better for the defender. However, deceptively hiding security resources can achieve this goal by making those targets appear more attractive.

Theorem 2 (The Power of Deception)**.** *For any* ϵ*, there are game instances such that:*

1. *The defender uses* 1 *resource in* ***Situation 2*** *but at least* $(n-1)$ *resources in* ***Situation 1****;*
2. *Nevertheless, the defender achieves* strictly *higher utility in* ***Situation 2*** *than in* ***Situation 1****.*

Proof Sketch. Given any ϵ, the proof proceeds by constructing such a game instance with n targets and only *one* security resource. Pick any $n > \max\{3, 1/\epsilon^2\}$. The proof constructs the players' payoffs as in the following table where target 2 to n have identical payoffs. Target 1 is the only special target with larger payoff scales for both players.

Targets	R_i^d	P_i^d	R_i^a	P_i^a
1	n^3	$2-2n$	$2n-2$	$-(n-1)(n-2)$
2	0	$-n$	n	0
·	·	·	·	·
n	0	$-n$	n	0

What remains is to compute the defender utilities in the two situations. In **Situation 1**, simple calculation shows that the optimal defender strategy is to protect target 2 to n with probability 1 and target 1 with probability $2/n$. This makes all targets to have an attacker utility of zero. The defender achieves utility $\frac{2}{n} \cdot n^3 - \frac{n-2}{n} \cdot 2(n-1) > 0$ (since $n > 3$) at target 1 whereas the defender utility at other targets is 0 due

[2] The defender may be able to exaggerate her protection as well, but that is not necessary for our next result.

to their zero-sum payoff structure. Therefore, the attacker will break ties in favor of the defender and attacks target 1.[3]

In **Situation 2**, the defender can hide protection probability by at most ϵ. Calculations show that in this case, the defender can cover target 1 with probability $p_1 = \frac{1}{n} + \epsilon\frac{(n-1)^2}{1+(n-1)^2}$ and other $(n-1)$ targets with probability $(1-p_1)/(n-1)$. These probabilities are chosen such that $p_1 - \epsilon$ protection of target 1 is equally attractive to the attacker as $(1-p_1)/(n-1)$ protection for each of target 2 to n, and moreover the attacker will break ties by attacking target 1 as well. Since $n > \max\{3, 1/\epsilon^2\}$, and thus $\epsilon > 1/\sqrt{n}$. Therefore, we obtain $p_1 = \frac{1}{n} + \epsilon\frac{(n-1)^2}{1+(n-1)^2} \geq \frac{1}{n} + \frac{1}{\sqrt{n}} \cdot \frac{9}{10} > \frac{2}{n}$, which is target 1's protection probability in **Situation 1**. Therefore, the defender's utility here is strictly higher than **Situation 1**, concluding our proof. □

Deception Against Sophisticated Attacks

We now turn to the case of a more sophisticated attacker who is aware of the defender's deception. We consider two natural types of attack strategies to combat the defender's deception: (i) **Maximin attack strategy** — the attacker is aware of the defender's deception and take a Maximin-based robust approach to deal with his uncertainty about the defender's true strategy; and (ii) **Equilibrium attack strategy** — the attacker follows an equilibrium strategy approach to counteract the defender's deception.

Maximin Attack Strategy

When the attacker is aware of the defender's deception, one natural approach for the attacker to counteract is to consider the possible set $[c_i - \beta_i, c_i + \alpha_i]$ that x_i is from, given that c_i is observed. The attacker can attempt to deal with the worst-case scenario within this uncertainty set.

Computing Optimal Defender Deception. To compute an optimal deception strategy, we first analyze the attacker's attack maximin strategy. Given the observed $\{c_i\}$, the attacker chooses the target i^* that maximizes his worst-case utility over all possible $\mathbf{x}$'s, defined as follows:

$$i^* \in \text{argmax}_i \left[\min_{\mathbf{z}} \; R_i^a(1-z_i) + P_i^a z_i\right]$$
$$\text{s.t. } 0 \leq z_i \leq 1, c_i - \beta_i \leq z_i \leq c_i + \alpha_i, \forall i.$$

Despite such more intricate behavior of the attacker, we show that the optimal defender strategy can still be computed efficiently in very general setups.

Theorem 3. *If the standard* SSE *can be computed in polynomial time, then the optimal defender strategy against a maximin attacker can also be computed in polynomial time.*

Proof Sketch. The proof starts with a characterization of the optimal solution to the following maximin problem:

$$i^* \in \text{argmax}_i \left[\min_{\mathbf{z}} \; R_i^a(1-z_i) + P_i^a z_i\right]$$
$$\text{s.t. } 0 \leq z_i \leq 1, c_i - \beta_i \leq z_i \leq c_i + \alpha_i, \forall i.$$

[3]It is worthwhile to mention that if the defender tries to protect target 1 with a probability higher than $2/n$, the attacker will then attack a target $i > 1$, reducing the defender utility to 0. This is the reason why more security resources are not necessarily helpful.

We observe that for any given target i, the optimal solution of the inner minimization in the above maximin problem is $z_i^{\text{worst}} = \min\{c_i + \alpha_i, 1\}$. Therefore, the attacker will choose the target i which has the highest expected utility w.r.t the worse case coverage $\min\{c_i + \alpha_i, 1\}$.

Utilizing this observation, we can formulate the problem of finding the optimal defender strategy against a maximin attacker as an optimization program (OP).

$$\max_{\mathbf{x}, p_{\boldsymbol{e}}, \mathbf{q}} \; U^d$$
$$\text{s.t. } \max\{x_i - \alpha_i, 0\} \leq c_i \leq \min\{x_i + \beta_i, 1\}, \forall i$$
$$U^d \leq x_i R_i^d + (1-x_i)P_i^d + (1-q_i)M, \forall i$$
$$\mathbf{x} = \sum_{\boldsymbol{e} \in \mathcal{E}} p_{\boldsymbol{e}} \boldsymbol{e}, \sum_{\boldsymbol{e} \in \mathcal{E}} p_{\boldsymbol{e}} = 1, p_{\boldsymbol{e}} \geq 0 \; \forall \boldsymbol{e} \in \mathcal{E}$$
$$\sum\nolimits_i q_i = 1, q_i \in \{0,1\}, \forall i$$
$$U^a \leq \min\{c_i + \alpha_i, 1\}(P_i^a - R_i^a) + R_i^a + (1-q_i)M, \forall i$$
$$U^a \geq \min\{c_i + \alpha_i, 1\}(P_i^a - R_i^a) + R_i^a, \forall i$$

There are two major difficulties in solving the above OP. First, it has a *non-convex* constraint, i.e. the last constraint, since $\min\{c_i + \alpha_i, 1\}$ is a concave function. To overcome this barrier, we analyze the above OP and show that it always admits an optimal solution where $\min\{c_i + \alpha_i, 1\} = c_i + \alpha_i$ for any i. This property allows us to substitute all the $\min\{c_i + \alpha_i, 1\}$ in the above OP by linear term $c_i + \alpha_i$, resulting in a MILP with $\mathbf{q}$ as the only integer variable. Since $\mathbf{q}$ only takes at most n different values, we can enumrate each possibility and solve the corresponding LP it induces (Conitzer and Sandholm 2006). The second challenge is that this LP however has exponentially many variables due to $\{p_{\boldsymbol{e}}\}_{\boldsymbol{e} \in \mathcal{E}}$. To resolve this challenge, we leverage linear program duality to devise a polynomial-time Turing reduction from solving this large LP to solving the standard security game version without deception, i.e., the special case with $\alpha_i, \beta_i = 0, \forall i$. □

We now study the power of deception against a maximin attacker. Let $\bar{\mathbf{x}}$ be the defender SSE strategy in **Situation 1**. Note that even though the defender can use as many resources in **Situation 1**, it is not always beneficial for the defender to protect all targets all the time. In other words, typically $\bar{\mathbf{x}} \neq \mathbf{1}$. Recall that $\Gamma(\mathbf{x}) = \{i : U_i^a(x_i) \geq U_j^a(x_j), \forall j\}$ is the *attack set*, containing all targets at which the attacker has the highest expected utility with respect to a strategy $\mathbf{x}$ of the defender. Similarly, we denote by $\Gamma^d(\mathbf{x}) = \{i \in \Gamma(\mathbf{x}) : U^d(x_i) \geq U^d(x_j), \forall j \in \Gamma(\mathbf{x})\}$ the set of equally good targets within $\Gamma(\mathbf{x})$ for the defender.[4] When the attacker breaks tie in favor of the defender, the attacker will attack one of the targets in $\Gamma^d(\mathbf{x})$. We define the following quantity:

$$L = \min_{i \in \Gamma^d(\bar{\mathbf{x}}): \bar{x}_i \geq \alpha_i} \left[\bar{x}_i + \sum\nolimits_{j \neq i} \max\{0, \bar{x}_j - \alpha_j - \beta_j\}\right] \quad (2)$$

If $\bar{x}_i < \alpha_i$ for all $i \in \Gamma^d(\bar{\mathbf{x}})$, then we simply set $L = +\infty$ (for the sake of presentation). Note that $\bar{\mathbf{x}}$

[4]Note that $\Gamma^d(\mathbf{x})$ can have multiple targets. For example, $\Gamma^d(\mathbf{x}) = \Gamma(\mathbf{x})$ in zero-sum games.

is the defender SSE strategy in **Situation 1**, therefore the term $\left[\bar{x}_i+\sum_{j\neq i}\max\{0,\bar{x}_j-\alpha_j-\beta_j\}\right]$ is roughly the total amount of resources needed in $\bar{\mathbf{x}}$, lessened by the deception capabilities. Theorem 4 compares the effectiveness of deception to that of using more resources as in **Situation 1**.

Theorem 4 (The Power of Deception)**.** *For any $(\alpha_i,\beta_i)\neq 0$, the utility the defender obtains for playing deceptively does not exceed the utility she would achieve in **Situation 1**. In particular, if the total number of defender resources in* SeGDC*: $k<L$, the defender's utility in **Situation 1** is strictly better than in* SeGDC*. If $k\geq L$, her utility is the same in both cases when there is no resource-allocation constraint.*

Proof Sketch. The proof of Theorem 4 is rather involved. Our argument for the situation of $k\geq L$ proceeds by (i) first showing that we can carefully craft a deceptive strategy $(\mathbf{x},\mathbf{c})$, based on the SSE strategy $\bar{\mathbf{x}}$ in **Situation 1**, which leads to a greater defender utility in SeGDC, leveraging that we have enough resources, i.e., $k\geq L$ for the crafted strategy, and (ii) then showing that we can do the same for crafting a defense strategy $\mathbf{x}^1$ which leads to a greater defender utility in **Situation 1**, based on the equilibrium strategy $(\mathbf{x}^*,\mathbf{c}^*)$ in SeGDC. First, let i' be the target which corresponds to the value of L in (2). Based on $\bar{\mathbf{x}}$, we can generate the defender's strategy $(\mathbf{x},\mathbf{c})$ to play in SeGDC as follows:

$$x_{i'}=\bar{x}_{i'}$$
$$x_j=\max\{0,\bar{x}_j-\alpha_j-\beta_j\},\forall j\neq i'$$
$$c_{i'}=\max\{x_{i'}-\alpha_{i'},0\}=\bar{x}_{i'}-\alpha_{i'}\text{ (since }\bar{x}_{i'}\geq\alpha_{i'})$$
$$c_j=x_j+\beta_j=\max\{\beta_j,\bar{x}_j-\alpha_j\},\forall j\neq i'.$$

When the defender plays this $(\mathbf{x},\mathbf{c})$ in SeGDC, the maximin attacker responds w.r.t the worst case coverage, $\mathbf{c}+\alpha$. According to the definition of the target i', we can easily obtain $i'\in\Gamma(\mathbf{c}+\alpha)$. As a result, the defender will receive an utility that is at least $U_{i'}^d(x_{i'})$, which is exactly the defender's equilibrium utility in **Situation 1**.

Conversely, give the optimal strategy $(\mathbf{x}^*,\mathbf{c}^*)$ in SeGDC, we can generate a strategy of the defender at **Situation 1** as follows: $x_i^1=c_i^*+\alpha_i$ for all i. When the defender plays $\{x_i^1\}$ in **Situation 1**, target i^* is also the best response of the attacker, and the defender's utility at this target, $U_{i^*}^d(c_{i^*}^*+\alpha_{i^*})$, is no less than the defender's utility in SeGDC, $U_{i^*}^d(x_i^*)$. Combining with the previous result, the defender's equilibrium utility thus must be the same in both cases.

Finally, our argument for the situation of $k<L$ proceeds by proving that — either (i) a strategy $\mathbf{x}^1$ for the defender in **Situation 1** can be carefully crafted based on the deceptive strategy $\mathbf{c}^*$, which leads to the same attacked target i^* but with a strictly higher coverage $x_{i^*}^1>x_{i^*}^*$; or (ii) the defender's coverage at i^* satisfies: $x_{i^*}^*\leq\bar{x}_{i^*}$ and $i^*\in\Gamma(\bar{\mathbf{x}})\setminus\Gamma^d(\bar{\mathbf{x}})$, i.e., i^* is not attacked in **Situation 1**. Both (i) and (ii) imply that the defender achieves a strictly higher utility in **Situation 1** compared to SeGDC. □

Theorem 4 shows that **Situation 1** is never worse than SeGDC (as opposed to Theorem 2 for ignorant attacker) when facing a maximin attacker. In fact, **Situation 1** is strictly better when: (1) the defender does not have enough resources, i.e., $k<L\neq\infty$; or (2) $\bar{x}_i<\alpha_i$ for all $i\in\Gamma^d(\bar{\mathbf{x}})$. The later situation is particularly interesting and is due to the following reason. The SSE strategy places $\bar{x}_i(<\alpha_i)$ coverage on the attacked target, however, the maximin attacker will always overestimate its coverage probability because he thinks target i is covered with at least α_i. As a result, the defender is not able to induce the attacker to attack her favorable targets from set $\Gamma^d(\bar{\mathbf{x}})$, thus resulting in utility decrease.

Equilibrium Attack Strategy

We now turn to the most sophisticated situation where each player reasons about its opponent's strategy to the infinity. Given the attacker can only observe the defender's committed strategy $\mathbf{c}$ but not the actual strategy $\mathbf{x}$, this leads to our study of the sequential equilibrium (Shoham, Leyton-Brown et al. 2009). We denote by $\Omega^{true}(\mathbf{c})=\{\mathbf{x}\in\mathbf{X}:c_i-\beta_i\leq x_i\leq c_i+\alpha_i,\forall i\}$ the set of all possible actual defense strategies given an observed strategy $\mathbf{c}$. We start with a few necessary definitions which we use to define a sequential equilibrium.

Definition 2 (Attacker Behavior Strategy)**.** *For each observed strategy $\mathbf{c}$ (aka. information set), a behavior strategy of the attacker is a randomization over targets to attack, denoted by $q(i\mid\mathbf{c})\in[0,1]$, where $\sum_i q(i\mid\mathbf{c})=1$.*

Definition 3 (Defender Behavior Strategy)**.** *A behavior strategy of the defender at his information set (which is the $\emptyset$ information set at the beginning of the game) is a randomization over her strategies $(\mathbf{x},\mathbf{c})$ in* SeGDC*, denoted by $p(\mathbf{x},\mathbf{c})\in[0,1]$, where $\int_{(\mathbf{x},\mathbf{c})\in\Omega}p(\mathbf{x},\mathbf{c})d(\mathbf{x},\mathbf{c})=1$.*

Definition 4 (Bayes Belief Update)**.** *Given a pair of strategies $(\mathbf{p},\mathbf{q})$, for each observed deceptive strategy $\mathbf{c}$, the attacker can update his belief, using the Bayes rule as follows:*

$$b(\mathbf{x}\mid\mathbf{c})\propto p(\mathbf{x},\mathbf{c}),\forall\mathbf{x}\in\Omega^{true}(\mathbf{c})$$

We are now ready to define the sequential equilibrium.

Definition 5 (Sequential equilibrium)**.** *A pair $(\mathbf{p}^*,\mathbf{q}^*)$ forms a sequential equilibrium if and only if there exist probability distributions $\mathbf{b}^*$ such that:*

1. *$(\mathbf{p}^*,\mathbf{q}^*,\mathbf{b}^*)=\lim_{n\to\infty}(\mathbf{p}^n,\mathbf{q}^n,\mathbf{b}^n)$ for some sequence $(\mathbf{p}^1,\mathbf{q}^1,\mathbf{b}^1),\ldots$ where $\mathbf{p}^n,\mathbf{q}^n$ is fully mixed, and the belief $\mathbf{b}^n$ is* consistent *with $(\mathbf{p}^n,\mathbf{q}^n)$ (i.e., this belief is precisely the one defined by Bayes' rule).*
2. *The attacker obtains the highest expected utility based on his belief update $\mathbf{b}^*$ at each of his information set (aka. each of possible deceptive strategy $\mathbf{c}$ he would observe):*

$$\int_{x\in\Omega^{true}(\mathbf{c})}b^*(\mathbf{x}\mid\mathbf{c})\sum\nolimits_i q^*(i\mid\mathbf{c})U_i^a(x_i)d\mathbf{x}$$
$$\geq\int_{x\in\Omega^{true}(\mathbf{c})}b^*(\mathbf{x}\mid\mathbf{c})U_i^a(x_i)d\mathbf{x},\forall i$$

3. *The defender obtains the highest expected utility at his information set (which is the $\emptyset$ information set at the beginning of the game) against the attacker's strategy $\mathbf{q}^*$:*

$$\int_{(\mathbf{x},\mathbf{c})\in\Omega}p^*(\mathbf{x},\mathbf{c})\sum\nolimits_i q^*(i\mid\mathbf{c})U_i^d(x_i)d(\mathbf{x},\mathbf{c})$$
$$\geq\sum\nolimits_i q^*(i\mid\mathbf{c})U_i^d(x_i),\forall(\mathbf{x},\mathbf{c})\in\Omega$$

Since the defender strategy space is infinite, it is not straightforward to show the existence of a sequential equilibrium. Nevertheless, our next result shows that there always exists a sequential equilibrium of the game which is equivalent to the Nash equilibrium of SSGs without deception.

Theorem 5. *Given any non-deceptive Nash equilibrium* $(\mathbf{x}^0, \mathbf{q}^0)$*, there always exists a sequential equilibrium,* $(\mathbf{p}^*, \mathbf{q}^*)$*, of the game, which is* equivalent *to* $(\mathbf{x}^0, \mathbf{q}^0)$ *in the following sense:*

1. *Attacker:* $q^*(i \mid \mathbf{c}) = q_i^0, \forall i$ *and* $\mathbf{c}$
2. *Defender:* $\int_{(\mathbf{x},\mathbf{c})\in\Omega} p^*(\mathbf{x}, \mathbf{c}) x_i d(\mathbf{x}, \mathbf{c}) = x_i^0, \forall i$

Proof Sketch. We first construct behavior strategies for the players $(\mathbf{p}, \mathbf{q})$ based on $(\mathbf{x}^0, \mathbf{q}^0)$ as follows:

Attacker behavior strategy. Given any observed $\mathbf{c}$, the attacker plays $q(i \mid \mathbf{c}) = q_i^0$.

Defender behavior strategy. We construct the distribution $p(\mathbf{x}, \mathbf{c}) = p(\mathbf{c})p(\mathbf{x} \mid \mathbf{c})$ based on the determination of $p(\mathbf{c})$ and $p(\mathbf{x} \mid \mathbf{c})$ as follows: (i) $p(\mathbf{c})$ can be any distribution supported on the domain $\Omega(\mathbf{x}^0)$; and (ii) the distribution $p(\mathbf{x} \mid \mathbf{c})$ is determined such that:

$$\forall \mathbf{c} \in \Omega(\mathbf{x}^0) : \int_{\mathbf{x}\in\Omega^{true}(\mathbf{c})} p(\mathbf{x} \mid \mathbf{c}) x_i d\mathbf{x} = x_i^0, \forall i,$$

$$\forall \mathbf{c} \notin \Omega(\mathbf{x}^0) : p(\mathbf{x} \mid \mathbf{c}) \text{ is arbitrary on the domain } \Omega^{true}(\mathbf{c})$$

We follow the trembling-hand approach to build a sequential equilibrium $(\mathbf{p}^*, \mathbf{q}^*, \mathbf{b}^*)$. We only need to examine the information sets of the attacker which has a a zero probability of occurrence. Note that only observation histories of the attacker (aka. information sets) which correspond to $\mathbf{c} \notin \Omega(\mathbf{x}^0)$ have a zero probability of occurrence. Therefore, for each $\epsilon > 0$, we construct a new strategy of the defender $\mathbf{p}_\epsilon$ with $p_\epsilon(\mathbf{x}, \mathbf{c}) = p_\epsilon(\mathbf{c})p_\epsilon(\mathbf{x} \mid \mathbf{c})$, as follows:

- $p_\epsilon(\mathbf{c}) = \epsilon \cdot \frac{\text{volume}(\Omega(\mathbf{X}))}{\text{volume}(\Omega(\mathbf{x}^0))}$ if $\mathbf{c} \notin \Omega(\mathbf{x}^0)$
- $p_\epsilon(\mathbf{c}) = p(\mathbf{c}) - \epsilon$ if $\mathbf{c} \in \Omega(\mathbf{x}^0)$
- $p_\epsilon(\mathbf{x} \mid \mathbf{c})$ is defined the same as $p(\mathbf{x} \mid \mathbf{c})$

where $\Omega(\mathbf{X})$ is the entire feasible domain of deceptive strategies $\mathbf{c}$ and $\Omega(\mathbf{x}^0)$ is the feasible domain of $\mathbf{c}$ with respect to the actual defense strategy $\mathbf{x}^0$. Based on this behavior strategy $\mathbf{p}_\epsilon$ of the defender, we construct the new belief of the attacker at each observation history $\mathbf{c} \notin \Omega(\mathbf{x}^0)$ as follows:

$$b^*(\mathbf{x} \mid \mathbf{c}) = \lim_{\epsilon\to 0} p_\epsilon(\mathbf{x} \mid \mathbf{c}) = \lim_{\epsilon\to 0} \frac{p_\epsilon(\mathbf{x}, \mathbf{c})}{\int_{\mathbf{x}'\in\Omega^{true}(\mathbf{c})} p_\epsilon(\mathbf{x}', \mathbf{c}) d\mathbf{x}'}$$
$$= \frac{p(\mathbf{x} \mid \mathbf{c})}{\int_{\mathbf{x}'\in\Omega^{true}(\mathbf{c})} p(\mathbf{x}' \mid \mathbf{c}) d\mathbf{x}'}$$

while keeping beliefs at other observation histories $\mathbf{c} \in \Omega(\mathbf{x}^0)$ unchanged, i.e., $b^*(\mathbf{x} \mid \mathbf{c}) = b(\mathbf{x} \mid \mathbf{c})$. We now construct a sequential equilibrium $(\mathbf{p}^*, \mathbf{q}^*)$ which is the same as $(\mathbf{p}, \mathbf{q})$ except for the attacker strategies at observation histories with a zero probability, $\mathbf{c} \notin \Omega(\mathbf{x}^0)$ — we replace these strategies with the best response of the attacker with respect to the new belief $b^*(\mathbf{x} \mid \mathbf{c})$. This is straightforward to compute since we just need to find the target that maximizes the attacker expected utility with respect to this belief at $\mathbf{c}$. We can easily verify that $p^*(\mathbf{x}, \mathbf{c}) = \lim_{\epsilon\to 0} p_\epsilon(\mathbf{x}, \mathbf{c})$ and $b^*(\mathbf{x} \mid \mathbf{c}) = \lim_{\epsilon\to 0} p_\epsilon(\mathbf{x} \mid \mathbf{c})$ (as defined).

Now, we only need to prove that (i) the attacker plays a best response at every observation history, $\mathbf{c} \in \Omega(\mathbf{x}^0)$, with a non-zero probability of occurrence; and (ii) the defender plays a best response at his information set (which is the $\emptyset$ information set at the beginning of the game). The detail of this part is included in the appendix. □

As a result, a sequential equilibrium can be computed from the Nash equilibrium for the standard SSG, which is known to admit efficient polynomial-time algorithms for general SSGs (Korzhyk et al. 2011; Xu 2016). Moreover, (Korzhyk et al. 2011) show that the attacker will achieve the same utility in any Nash equilibrium as in the SSE whereas the defender will achieve lower utility under mild assumptions. These results and Theorem 5 yield the following characterization about SeGDCs.

Corollary 1 (Characterization of Equilibrium Utility)**.** *Under mild non-degeneracy assumption,*[5] *compared to the standard* SSE *without deception, the attacker's utility will not change while the defender obtains a lower utility in the sequential equilibrium strategies as determined in Theorem 5 for any* SeGDC*.*

Remark 1. *The key conceptual message from Theorem 5 is that the defender loses her first-mover advantage during deception despite that she can still commit to the deceptive coverage* $\mathbf{c}$*. This is due to the attacker's uncertainty of the mapping* $\mathbf{x} \to \mathbf{c}$*. It forces the attacker to do Nash equilibrium reasoning, which turns out to be harmful to the defender. These results illustrate the double-edged role of deception in* SSGs*, which depends on the sophistication of the attacker.*

Experiments

Our experiments are conducted on a High Performance Computing (HPC) cluster, with processors are dual E5-2690v4 (28 cores) and 128 GB memory. Our experiments use the standard covariance game generator GAMMUT (http://gamut.stanford/edu), to generate payoff matrices. A covariance value r governs the correlation between the defender and attacker's payoffs. In particular, when $r = -1.0$, the generated games are zero-sum. When $r = 0.0$, the payoffs of players are not correlated. The players' rewards and penalties are within $[1, 10]$ and $[-10, -1]$, respectively. We consider four cases in our evaluation: (i) SSE — the defender is not deceptive; (ii) Dec.Ignorant; (iii) Dec.Maximin; and (iv) Dec.Equilibrium. In the last three cases, the defender plays deceptively while the attacker plays the Ignorant, Maximin, and Equilibrium strategies, respectively. We use Cplex to solve our optimization programs (https://www.ibm.com/analytics/cplex-optimizer). Each data point is averaged over 200 games. Our results are statistically significant (t-boostrap with $\alpha = 0.05$ (Wilcox 2003)).

[5]Precisely, the assumption is $P_i^a \neq r^*, \forall i$ where r^* is the minimax attacker utility. This is a standard non-degeneracy assumption for security games; See (Korzhyk et al. 2011).

Solution quality. Our results are shown in Figure 1(a-h). The x-axis is the ratio of the number of resources to targets ($\frac{k}{n}$) or the covariance value. The y-axis is the defender or attacker's utility on average. We consider two cases of the defender's deception capability: (i) small deception interval, i.e., $\alpha_i = \beta_i = 0.05$; and big deception interval, i.e., $\alpha_i = \beta_i = 0.15$. We only highlight the results with $n = 100$ targets. Our results on different n also exhibit similar trends.

Figure 1(a-h) shows that the defender gains a significantly higher utility for playing deceptively when the attacker either ignores the defender's deception or plays the Maximin attack strategy to counter that deception (`Dec.Ignorant` and `Dec.Maximin` versus `SSE`). On the other hand, if the attacker plays the sophisticated equilibrium strategy, the defender suffers a great loss in utility (`Dec.Equilibrium` versus `SSE`). The defender's benefit against Ignorant and Maximin attacker increases when the size of the deception interval increases ($\alpha_i = \beta_i = 0.05$ versus $\alpha_i = \beta_i = 0.15$). Conversely, the attacker suffers a significant loss in utility on average when he plays either `Dec.Ignorant` or `Dec.Maximin`. Finally, by following `Dec.Equilibrium`, the attacker obtains the same utility as in `SSE`.

In Figure 1(a-b), the defender's utility increases while the attacker's utility decreases gradually as the ratio ($\frac{k}{n}$) increases. This makes sense since the coverage probability of the defender at each target increases when ($\frac{k}{n}$) increases. In addition, the maximin strategy helps the attacker in reducing the impact of the defender's deception compared with `Dec.Ignorant` and this help is substantial when the size of the deception interval is large (Figure 1(c) versus 1(d)).

In Figure 1(e-f), the defender's utility gradually decreases as the covariance value r gets closer to -1. Indeed, when the games become zero-sum ($r = -1$), the attacker would aim to minimize the defender's utility (which is equivalent to maximize the attacker's utility), leading to a roughly lower utility for the defender compared with non-zero-sum games.

Runtime performance. We evaluate the runtime performance of our algorithms to solve `SeGDC` w.r.t different attack strategies. The results are shown in Figure 1(i). We only show the results on small deception intervals $\alpha_i = \beta_i = 0.05$ due to space limit. Note that the runtime of `SSE` and `Dec.Equilibrium` is approximately the same based on Theorem 5 and the results presented in (Korzhyk et al. 2011). Overall, Figure 2 shows that our proposed algorithm can scale up to large games. In particular, it takes all algorithms less than 40 seconds on average to solve 200-target games.

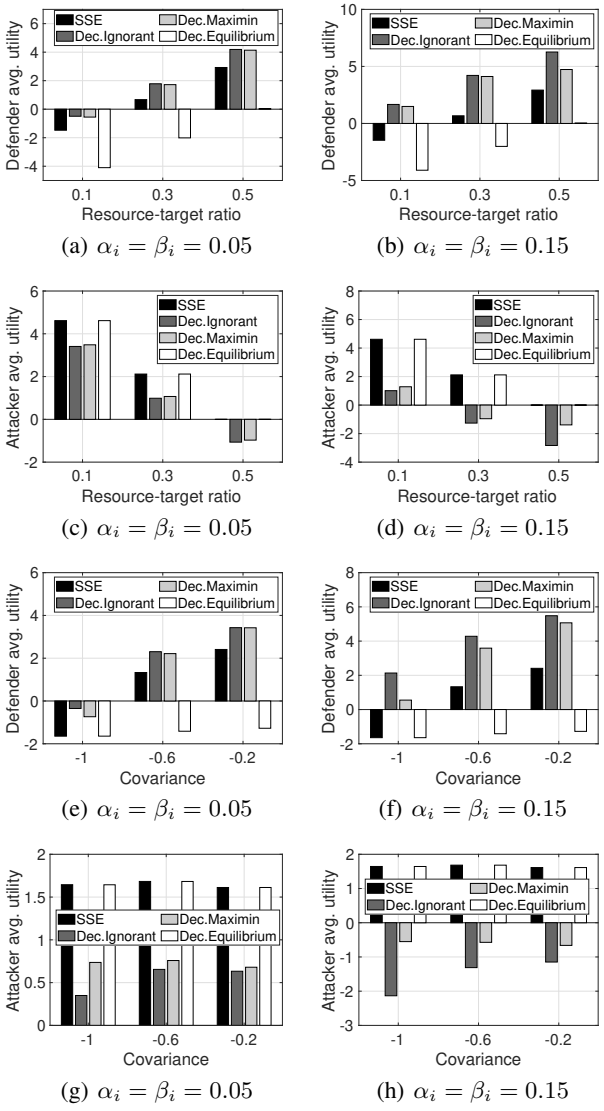

(a) $\alpha_i = \beta_i = 0.05$ (b) $\alpha_i = \beta_i = 0.15$

(c) $\alpha_i = \beta_i = 0.05$ (d) $\alpha_i = \beta_i = 0.15$

(e) $\alpha_i = \beta_i = 0.05$ (f) $\alpha_i = \beta_i = 0.15$

(g) $\alpha_i = \beta_i = 0.05$ (h) $\alpha_i = \beta_i = 0.15$

Figure 1: Solution quality evaluation

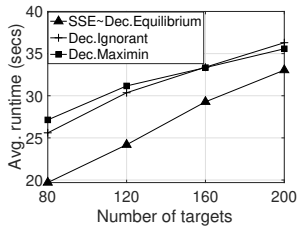

Figure 2: Runtime performance

Summary

This work investigates the complexity and effectiveness of optimal defender deception in three attack models, with an increasing sophistication: (i) Ignorance attack strategy; (ii) Maximin attack strategy; and (iii) Equilibrium attack strategy. In each case, we provide a polynomial-time algorithm to compute an optimal defender deception. We show that the defender's deception becomes strictly less powerful, as the attacker becomes more sophisticated, to the point where deception even leads to a decrease in the defender's utility compared with the no-deception situation. Our results provide formal separations for the effectiveness of patrol deception when facing an attacker of increasing sophistication.

References

Basilico, N.; Gatti, N.; and Amigoni, F. 2009. Leader-follower strategies for robotic patrolling in environments with arbitrary topologies. In *AAMAS*.

Bondi, E.; Oh, H.; Xu, H.; Fang, F.; Dilkina, B.; and Tambe, M. 2020. To Signal or Not To Signal: Exploiting Uncertain Real-Time Information in Signaling Games for Security and Sustainability. In *AAAI*.

Conitzer, V.; and Sandholm, T. 2006. Computing the optimal strategy to commit to. In *Proceedings of the 7th ACM conference on Electronic commerce*, 82–90.

Fang, F.; Nguyen, T. H.; Pickles, R.; Lam, W. Y.; Clements, G. R.; An, B.; Singh, A.; Tambe, M.; and Lemieux, A. 2016. Deploying PAWS: Field optimization of the protection assistant for wildlife security. In *Twenty-Eighth IAAI Conference*.

Fraunholz, D.; Anton, S. D.; Lipps, C.; Reti, D.; Krohmer, D.; Pohl, F.; Tammen, M.; and Schotten, H. D. 2018. Demystifying Deception Technology: A Survey. *arXiv preprint arXiv:1804.06196*.

Gan, J.; Guo, Q.; Tran-Thanh, L.; An, B.; and Wooldridge, M. 2019. Manipulating a Learning Defender and Ways to Counteract. In *Advances in Neural Information Processing Systems*.

Guo, Q.; An, B.; Bosansky, B.; and Kiekintveld, C. 2017. Comparing strategic secrecy and Stackelberg commitment in security games. In *26th International Joint Conference on Artificial Intelligence*.

Korzhyk, D.; Yin, Z.; Kiekintveld, C.; Conitzer, V.; and Tambe, M. 2011. Stackelberg vs. Nash in security games: An extended investigation of interchangeability, equivalence, and uniqueness. *Journal of Artificial Intelligence Research*, 41: 297–327.

Letchford, J.; and Vorobeychik, Y. 2011. Computing Randomized Security Strategies in Networked Domains. In *AARM*.

Nguyen, T.; and Xu, H. 2019. Imitative Attacker Deception in Stackelberg Security Games. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*, 528–534. International Joint Conferences on Artificial Intelligence Organization.

Nguyen, T. H.; Wang, Y.; Sinha, A.; and Wellman, M. P. 2019. Deception in Finitely Repeated Security Games. In *33th AAAI Conference on Artificial Intelligence*.

Rabinovich, Z.; Jiang, A. X.; Jain, M.; and Xu, H. 2015. Information disclosure as a means to security. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*, 645–653. Citeseer.

Shoham, Y.; Leyton-Brown, K.; et al. 2009. Multiagent systems. *Algorithmic, Game-Theoretic, and Logical Foundations*.

Sinha, A.; Fang, F.; An, B.; Kiekintveld, C.; and Tambe, M. 2018. Stackelberg Security Games: Looking Beyond a Decade of Success. In *IJCAI*, 5494–5501.

Tambe, M. 2011. *Security and Game Theory: Algorithms, Deployed Systems, Lessons Learned*. Cambridge University Press. ISBN 9781139503662.

Van Uhm, D. P. 2016. *The illegal wildlife trade: Inside the world of poachers, smugglers and traders*, volume 15. Springer.

Wilcox, R. R. 2003. *Applying contemporary statistical techniques*. Elsevier.

Xu, H. 2016. The mysteries of security games: Equilibrium computation becomes combinatorial algorithm design. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, 497–514.

Xu, H.; Rabinovich, Z.; Dughmi, S.; and Tambe, M. 2015. Exploring Information Asymmetry in Two-Stage Security Games. In *29th AAAI Conference on Artificial Intelligence*, 1057–1063.

Xu, H.; Wang, K.; Vayanos, P.; and Tambe, M. 2018. Strategic coordination of human patrollers and mobile sensors with signaling for security games. In *Thirty-Second AAAI Conference on Artificial Intelligence*.

Zhuang, J.; Bier, V. M.; and Alagoz, O. 2010. Modeling secrecy and Deception in a multi-period attacker-defender signaling game. *European Journal of Operational Research*, 203: 409–418.