Optimal Deceptive Strategies in Security Games: A Preliminary Study

Yue Yin
The Key Lab of Intelligent Information Processing, ICT, CAS
University of Chinese Academy of Sciences
Beijing 100190, China
melody1235813@gmail.com

Bo An
Nanyang Technological University
Singapore 639798
boan@ntu.edu.sg

Yevgeniy Vorobeychik
Vanderbilt University
Nashville, TN 37235
yevgeniy.vorobeychik@vanderbilt.edu

Jun Zhuang
SUNY at Buffalo
Buffalo, NY 14260
jzhuang@buffalo.edu

Abstract

Attacker-defender Stackelberg games have been used in several deployed applications of game theory for infrastructure security. Security resources of the defender are game-theoretically allocated to prevent a strategic attacker from using surveillance to learn and exploit patterns in the allocation. Existing work on security games assumes that the defender honestly displays her real security resources. We introduce a new model in which the defender may use deceptive resources (e.g., a mock camera in the part for deterring potential adversaries, or a hidden camera on the road for detecting overspeed) to mislead the attacker. We provide algorithms for computing the defender's optimal strategy in consideration of deceptions. We also present experimental results evaluating the effectiveness of using deceptive strategies.

Introduction

Attacker-defender Stackelberg games have been used in several deployed applications of game theory for infrastructure security (Tambe 2011; Eric et al. 2012b). In this class of games, the defender first commits to a security strategy, then the attacker learns and responds to the defender's strategy (Paruchuri et al. 2008). Based on the assumption that the attacker responses optimally according to his knowledge of the defender's strategies, a solution to the game yields an optimal randomized strategy for the defender (Conitzer and Sandholm 2006). Applications based on Stackelberg games have been used in real world domains to make recommendations for allocating limited resources for protecting critical infrastructure (Pita et al. 2008; Eric et al. 2012a; Agmon, Urieli, and Stone 2011; Basilico, Gatti, and Amigoni 2009; Fang, Jiang, and Tambe 2013).

Most existing work on security games assumes that the defender honestly displays her security resources (Kiekintveld et al. 2009; Eric et al. 2012a; An et al. 2013). In reality, sometimes a deceptive resource, e.g., a mock camera in the park or a hidden camera on the road for detecting overspeed (Zhuang and Bier 2010), can also be used to deter illegal activities. Using deceptive resources may affect the attacker's knowledge of the defender's strategies, and can be used to potentially improve the payoff of the defender if assigned appropriately. Our goal in this work is to preliminarily investigate models that exploit the advantage of using deceptive resources for the defender.

Using deceptive resources can be considered as the defender holding private information in the game (Rasmusen and Blackwell 1994). Although there has been some research on security games with private information, previous approaches can not be used to solve our problem. This is because most of previous research, including (Wang and Zhuang 2011; Yin and Tambe 2012), assume that only the attacker holds private information, which is in contrast with our case. Tsai et al. (Jain et al. 2010) has focused on a scenario in which the defender used secret resources. However, they assumed that the attacker had a perfect knowledge of the defender's strategies anyway in spite of the deceptions. We will explore how the deceptive actions of the defender may affect the attacker's belief about the defender's strategy, and the attacker's best response. In this paper, we introduce a security game model in which the defenders strategies can include deceptive protections. The deceptive protections succeed with a certain probability. We also provide algorithms for computing the optimal defender strategy when the attacker surveils unlimitedly before attacking and analyze the advantage of using deceptive resources.

Security Games with Deception

In security games, the defender assigns security resources to potential attack targets to protect them. The attacker surveils the defender's actions, then chooses a target to attack. There are n targets $T = \{1, \ldots, n\}$. The defender has η_R real security resources. She can purchase extra $fake\ resources$ at the cost of β_F per fake resource. She can also convert some real resources into $secret\ resources$ at the cost of β_{Se} per secret resource. If the defender assigns a fake resource to a target, she is performing a fake protection. Namely, she pretends to protect a target when she is in fact not protecting it. Similarly, if the defender assigns a secret resource to a target, she is performing a secret protection. Namely, she pretends not to protect a target when she is in fact protecting it. Each deceptive protection fails with a probability of r.[1]

[1]We assume that all deceptive protections fail with the same probability for the ease of analysis.

We assume that the attacker does not reason about the deceptions used by the defender. We also assume that for an attacker who is surveilling the defender's strategy with a fake protection, if the fake protection succeeds, he will observe that the target is protected; If a fake protection fails, he will observe that the target is not protected. For the secret protection, if it succeeds, the attacker will observe no protection; if it fails, the attacker will observe that the target is protected. [2] The defender's budget is Bgt. We refer to this model as SGDB (Security Games with Deceptions and Budget constraints) in the rest of this paper.

If the attacker attacks a target i when i is protected by a real resource or a secret resource, he will achieve a payoff of $U_a^c(i)$, while the defender will achieve a payoff of $U_d^c(i)$. If the attacker attacks when i is protected by a fake resource or is not protected at all, then the attacker's payoff is $U_a^u(i)$ while the defender's payoff is $U_d^u(i)$. We assume that the attacker surveils the defender's actions unlimited times before he attacks, and chooses an optimal target to attack based on his belief about the defender's strategies. Next, we introduce the strategies of the agents and the equilibrium of the game in consideration of deception.

Strategies and Equilibrium

Pure strategies

A pure defender strategy can be defined as $s = \langle s_i, s_i \in \{R, F, Se, N\}\rangle$, with $s_i = R$ representing that target i is protected by a real resource, $s_i = F$ representing that target i is protected by a fake resource, $s_i = Se$ representing secret protection, and $s_i = N$ representing that target i is left alone (i.e., not protected). However, since what matters to the attacker is whether a target is covered or not, but not how a target is covered, a strategy observed by the attacker can be defined as $o = \langle o_i, o_i \in \{0, 1\}\rangle$, where $o_i = 0$ represents that the attacker observes that target i is not covered while $o_i = 1$ represents that the attacker observes that target i is covered.

Let $\gamma(o|s)$ represent the probability that a real defender strategy s is observed as o. For calculating $\gamma(o|s)$, we define an indicator function first.

$$\mathbf{1}_Y(s_i) = \begin{cases} 1, & if \quad s_i = Y, \\ 0, & otherwise. \end{cases} \tag{1}$$

Then we can calculate $\gamma(o|s)$ as follows.

Proposition 1. $\gamma(o|s) = \prod_{i\in T} \Big(\mathbf{1}_{Se}(s_i)(r\cdot o_i + (1-r)(1-o_i)) + \mathbf{1}_F(s_i)((1-r)o_i + r(1-o_i)) + \mathbf{1}_R(s_i)o_i + \mathbf{1}_N(s_i)(1-o_i)\Big)$.

Proof. Let $\gamma(o_i|s_i)$ represent the probability that when the defender plays a strategy s, the protection assigned to target i (namely, s_i) is observed by the attacker as o_i. For the four states of s_i, consider the following situations.

Case 1: $\mathbf{1}_{Se}(s_i) = 1$. With probability r the secret protection will fail, e.g., $o_i = 1$; with probability $1 - r$ the deception will succeed and $o_i = 0$. It follows that $\gamma(o_i|\mathbf{1}_{Se}(s_i) = 1) = r \cdot o_i + (1 - r)(1 - o_i)$.

Case 2: $\mathbf{1}_F(s_i) = 1$. With probability r the fake protection will fail, i.e., $o_i = 0$; with probability $1 - r$ the deception will succeed and $o_i = 1$. It follows that $\gamma(o_i|\mathbf{1}_F(s_i) = 1) = (1 - r)o_i + r(1 - o_i)$.

Case 3: $\mathbf{1}_R(s_i) = 1$. Then $o_i = 1$. It follows that $\gamma(o_i|\mathbf{1}_R(s_i) = 1) = o_i$.

Case 4: $\mathbf{1}_N(s_i) = 1$. Then $o_i = 0$. It follows that $\gamma(o_i|\mathbf{1}_N(s_i) = 1) = (1 - o_i)$.

Therefore, $\gamma(o|s) = \prod_{i\in T} \gamma(o_i|s_i)$. Since $\gamma(o_i|s_i) = \sum_{Y\in\{R,F,Se,N\}} \mathbf{1}_Y(s_i)\gamma(o_i|\mathbf{1}_Y(s_i) = 1)$. We then gain the result in Proposition 1. □

Let $\mathcal{S}$ denote the defender strategy space. When the defender plays $s \in \mathcal{S}$, the attacker will observe it as a certain o such that $\gamma(o|s) > 0$. The attacker's pure strategy can be represented as $\mathbf{a} = \langle a_i : a_i \in \{0, 1\}, \sum_{i\in T} a_i = 1\rangle$ where $a_i = 1$ represents that the attacker attacks target i. As previous research on security game, we restrict attacker strategies to pure strategies.

Mixed strategies

A mixed strategy played by the defender can be defined as $\mathbf{x} = \langle x_s\rangle$, with x_s representing the probability that pure strategy s is used. It can be compactly represented as a coverage vector $\mathbf{c} = \langle c_i^Y : Y \in \{R, F, Se\}\rangle$, where c_i^R is the probability that target i is covered by a real resource, c_i^F is the probability that target i is covered by a fake resource, and c_i^{Se} is the probability that target i is covered by a secret resource. We can compute the compact representation from a mixed strategy as follows.

$$c_i^Y = \sum_{s\in\mathcal{S}} x_s \mathbf{1}_Y(s_i) \quad \forall i \in T, Y \in \{R, F, Se\} \tag{2}$$

Assume that the defender commits to a mixed strategy $\mathbf{x}$, which can be compactly represented as $\mathbf{c}$, and executes a pure strategy s sampled from $\mathbf{x}$ each time. After making a large number of observations, the attacker will observe a mixed strategy $\mathbf{x_o} = \langle x_o\rangle$, with x_o representing the probability that strategy o is observed. This mixed $\mathbf{x_o}$ can be represented compactly as a coverage vector $\mathbf{e} = \langle e_i\rangle$, where e_i is the probability that the attacker observes target i being covered. We define the mapping from the defender's coverage vector $\mathbf{c}$ to the coverage vector $\mathbf{e}$ observed by the attacker as $h(\mathbf{c}) : \mathbf{c} \to \mathbf{e}$.

Proposition 2. *If the defender commits to a coverage vector* $\mathbf{c}$*, the attacker will observe it as a coverage vector* $\mathbf{e}$*, in which*

$$e_i = c_i^R + (1 - r)c_i^F + rc_i^{Se} \quad \forall i \in T \tag{3}$$

[2] In fact, we can also assume that if a fake protection fails, the attacker will observe a fake protection while if a secret protection fails, the attacker will observe a secret protection. This assumption will lead to the same attacker strategy as the previous one, as we show in the appendix.

Proof. Given o, it follows that $\sum_s x_s\gamma(o|s) = x_o$. Given s, it follows that $\sum_o \gamma(o|s) = 1$. Assume that given s, $\mathbf{1}_{Se}(s_j) = 1$ for a target j, it follows that

$$\begin{aligned}\sum_o \gamma(o|s)o_j &= \sum_{o,o_j=1} \gamma(o|s) \\ &= \sum_{o,o_j=1} \prod_{i\in T} \gamma(o_i|s_i) \\ &= \sum_o \gamma(o_j = 1|\mathbf{1}_{Se}(s_j)) \prod_{i\in T, i\neq j} \gamma(o_i|s_i) \\ &= r\sum_o \prod_{i\in T, i\neq j} \gamma(o_i|s_i) \\ &= r\end{aligned}$$

Similarly, if $\mathbf{1}_F(s_j) = 1$, it follows that $\sum_o \gamma(o|s)o_j = 1-r$. If $\mathbf{1}_R(s_j) = 1$, then $\sum_o \gamma(o|s)o_j = 1$. If $\mathbf{1}_N(s_j) = 1$, then $\sum_o \gamma(o|s)o_j = 0$. Thus $\sum_o \gamma(o|s)o_j = (\mathbf{1}_R(s_j) + \mathbf{1}_F(s_j)+\mathbf{1}_{Se}(s_j))\sum_o \gamma(o|s)o_j = \mathbf{1}_R(s_j)+(1-r)\mathbf{1}_F(s_j)+ r\mathbf{1}_{Se}(s_j)$.

Therefore, the relationship between $\mathbf{e}$ and $\mathbf{c}$ can be defined as follows (e_i is the coverage of target i observed by the attacker).

$$\begin{aligned}e_i &= \sum_o x_o o_i \\ &= \sum_o \sum_{s\in\mathcal{S}} x_s\gamma(o|s)o_i \\ &= \sum_s x_s \sum_o \gamma(o|s)o_i \\ &= \sum_s x_s\mathbf{1}_R(s_i) + (1-r)\sum_s x_s\mathbf{1}_F(s_i) + r\sum_s x_s\mathbf{1}_{Se}(s_i) \\ &= c_i^R + (1-r)c_i^F + rc_i^{Se}\end{aligned}$$

□

Stackelberg Equilibrium

As previous work on security game (Kiekintveld et al. 2009), we assume $U_d^c(i) - U_d^u(i) > 0$ and $U_a^u(i) - U_a^c(i) > 0$. For a strategy profile $\langle \mathbf{c}, \mathbf{e} = h(\mathbf{c}), \mathbf{a}\rangle$, the expected utilities for both agents (from their individual perspectives) are given by:

$$U_d(\mathbf{c}, \mathbf{a}) = \sum_{i\in T} a_i U_d(\mathbf{c}, i), \tag{4}$$

$$U_a(\mathbf{e}, \mathbf{a}) = \sum_{i\in T} a_i U_a(\mathbf{e}, i), \tag{5}$$

where $U_d(\mathbf{c}, i) = (c_i^r + c_i^{Se})U_d^c(i) + (1 - c_i^r - c_i^{Se})U_d^u(i)$ and $U_a(\mathbf{e}, i) = e_iU_a^c(i)+(1-e_i)U_a^u(i)$. It is noticeable that unlike security games without deceptions, the attacker's expected utility (from his perspective) now depends on what he observes but not what the defender plays. The attacker's response function is $g(\mathbf{e}) : \mathbf{e} \to \mathbf{a}$. Therefore, the Stackelberg equilibrium in this case can be defined as follows:

1. The defender plays a best-response: $U_d(\mathbf{c}, g(h(\mathbf{c}))) \geq U_d(\mathbf{c}', g(h(\mathbf{c}')))$ for any $\mathbf{c}'$.
2. The attacker plays a best-response: $g(\mathbf{e}) \in F_a(\mathbf{e})$ where $F_a(\mathbf{e}) = \arg\max_{\mathbf{a}} U_a(\mathbf{e}, \mathbf{a})$ is the set of follower best-responses.
3. The attacker breaks ties optimally for the defender: $U_d(\mathbf{c}, g(\mathbf{e})) \geq U_d(\mathbf{c}, \mathbf{a}')$ for any $\mathbf{a}' \in F_a(\mathbf{e})$.

Next, we explore how to compute the optimal defender strategy for a security game considering deceptive resources.

Computing Optimal Defender Strategy

For a security game with deceptive resources, the optimal defender strategy can be calculated by the following MILP (Mixed Integer Linear Program).

$$\mathbf{P1}: \quad \max_{\mathbf{c}} \quad d \tag{6}$$

$$\mathbf{s.t.} \quad a_i \in \{0,1\} \quad \forall i \in T \tag{7}$$

$$\sum_{i\in T} a_i = 1 \tag{8}$$

$$c_i^Y \in [0,1] \quad \forall i \in T, Y \in \{R, F, Se\} \tag{9}$$

$$c_i^R + c_i^F + c_i^{Se} \leq 1 \quad \forall i \in T \tag{10}$$

$$\beta_F\left\lceil \sum_{i\in T} c_i^F \right\rceil + \beta_{Se}\left\lceil \sum_{i\in T} c_i^{Se} \right\rceil \leq Bgt \tag{11}$$

$$\sum_{i\in T} c_i^R + \left\lceil \sum_{i\in T} c_i^{Se} \right\rceil \leq \eta_R \tag{12}$$

$$e_i = c_i^R + (1-r)c_i^F + rc_i^{Se} \quad \forall i \in T \tag{13}$$

$$d - U_d(i, \mathbf{a}) \leq (1-a_i)M \quad \forall i \in T \tag{14}$$

$$0 \leq k - U_a(i, \mathbf{e}) \leq (1-a_i)M \quad \forall i \in T \tag{15}$$

Eqs. (7) - (10) implement the feasibility of coverage and defender strategy. $\lceil x \rceil$ represents the smallest integer no less than x. Thus $\lceil \sum_{i\in T} c_i^F \rceil$ and $\lceil \sum_{i\in T} c_i^{Se} \rceil$ are the smallest number of fake resources and secret resources the defender should use to form the mixed strategy computed above. Eq. (11) constrains the costs on secret resources and fake resources to be within the budget. Eq. (12) ensures that secret resources are converted from real resources. Eq. (13) is used to calculate the attacker's observed coverage vector. In Eqs. (14) and (15), M is a large constant. The two constraints force the attacker to react optimally. Eq. (14), Eq. (15) and the objective together ensures the solutions satisfying the equilibrium we defined before.

A Fast Algorithm Since the run time of MILP increases exponentially with the increase of the scale of the problem, we introduce an algorithm based on ORIGAMI (Kiekintveld et al. 2009) to compute the optimal defender strategy fast. ORIGAMI computes a coverage vector such that the number of targets which are indifferent to the attacker is the largest. Based on the definition of SSE, the attacker breaks ties in favor of the defender. Thus the coverage vector computed by ORIGAMI leads to optimal defender utility. The targets which lead to the same attacker utility are defined as an *attack set*.

Since in our model, the attacker's choice depends on what he observes, namely $\mathbf{e}$, thus an attack set can be represented

as $\Gamma(\mathbf{e})$ with similar definition as ORGAMI uses. We introduce an algorithm called R-ORIGAMI (Revised ORIGAMI) to deal with an SGDB. The inputs of R-ORIGAMI are the number of real resources, the budget, the costs of using a fake or secret resource, and the probability r with which a fake or secret resource may fail. The outputs of R-ORIGAMI are the defender's optimal coverage vector $\mathbf{c}$, the corresponding coverage vector $\mathbf{e}$ observed by the attacker, and the attack set $\Gamma(\mathbf{e})$. The main idea is to calculate $\mathbf{e}$ and $\Gamma(\mathbf{e})$ using ORIGAMI, then choose a $\mathbf{c} = h^{-1}(\mathbf{e})$ and an $\mathbf{a} = g(\mathbf{e})$ to construct an SSE.

To calculate $\mathbf{e}$ and $\Gamma(\mathbf{e})$, we need to input the number of available resources in the attacker's perspective, namely the upper bound of $\sum e_i$, which depends on the number of real resources, fake resources, and secret resources the defender uses. However, given the budget and the costs of using a fake resource and a secret resource respectively, there may be a lot of combinations of fake resources and secret resources. To calculate the number of fake resources and secret resources the defender should use, we first analyze and find out the best combinations. We begin with the following observation.

Proposition 3. *The defender needs at most one secret resource.*

Proof. If the defender does not use any secret resources, all budget can be used to buy fake resources. Let m represent the number of fake resources, then $m = \left\lfloor \frac{Bgt}{\beta_F} \right\rfloor$. $\lfloor x \rfloor$ represents the largest integer no larger than x. Assume that a profile $\langle \mathbf{c}, \mathbf{e}, \mathbf{a} \rangle$ with $a_j = 1$ corresponds to an SSE. If the defender uses one secret resource, then the number of available real resources becomes $\eta_R - 1$, and the number of fake resources will be $m = \left\lfloor \frac{Bgt - \beta_{Se}}{\beta_F} \right\rfloor$. Assume that a profile $\langle \mathbf{c}', \mathbf{e}', \mathbf{a}' \rangle$ with $a'_k = 1$ corresponds to an SSE in which one secret resource is used. We have $U_d(\mathbf{c}, \mathbf{a}) = c_j^R U_d^c(j) + (1 - c_j^R) U_d^u(j), c_j^R \leq e_j$ and $U_d(\mathbf{c}', \mathbf{a}') = (c'^R_k + c'^{Se}_k) U_d^c(j) + (1 - c'^R_k - c'^{Se}_k) U_d^u(j)$. The relationship between $U_d(\mathbf{c}, \mathbf{a})$ and $U_d(\mathbf{c}', \mathbf{a}')$ depends on the game setting, thus using one secret resource may be helpful.

However, if the defender uses two secret resources, the number of real resources becomes $\eta_R - 2$ (assuming $\eta_R \geq 2$), $m = \left\lfloor \frac{Bgt - 2\beta_{Se}}{\beta_F} \right\rfloor$. Assume that a profile $\langle \mathbf{c}'', \mathbf{e}'', \mathbf{a}'' \rangle$ with $a''_l = 1$ corresponds to an SSE in which two secret resources are used. Then $\sum_{i \in T} e''_i \leq \eta_R - 2 + 2r + (1-r)m$. Obviously, $\sum_{i \in T} e''_i < \sum_{i \in T} e'_i$. Therefore, the size of $\Gamma(e'')$ is no larger than the size of $\Gamma(e')$, and for any $i \in \Gamma(e''), e''_i < e'_i$. Since $U_d(\mathbf{c}'', \mathbf{a}'') = (c''^R_l + c''^{Se}_l) U_d^c(l) + (1 - c''^R_l - c''^{Se}_l) U_d^u(l)$ and $c''^R_l + c''^{Se}_l \leq \min\{\frac{e'_l}{r}, 1\}$, we have $U_d(\mathbf{c}'', \mathbf{a}'') \leq U_d(\mathbf{c}', \mathbf{a}') \leq U_d(\mathbf{c}', \mathbf{a}')$. Therefore, using two secret resources will not lead to higher defender utility than using one secret resource. The same reasoning applies to more the two secret resources. Therefore, the defender needs at most one secret resource. □

Thus there are only two reasonable combinations of secret resources and fake resources for the defender. First, converting a real resource into a secret one, using the left budget to buy fake resources. Second, using all budget to buy fake resources. For each combination, the corresponding upper bound of $\sum e_i$ is taken as input of ORIGAMI to calculate $\mathbf{e}$ and $\Gamma(\mathbf{e})$.

Now we need to consider that given $\mathbf{e}$, how to construct $\mathbf{c} = h^{-1}(\mathbf{e})$ and $\mathbf{a} = g(\mathbf{e})$ which will lead to the highest defender utility. Proposition 3 has shown how we should assign secret resources. The following observation shows how to assign fake resources.

Proposition 4. *The defender should not assign any fake protection to the target the attacker will choose to attack.*

Proof. Assume that the attacker will choose to attack target i when i is covered with $c_i^R \geq 0, c_i^{Se} \geq 0, c_i^F > 0$, thus $e_i = c_i^R + r c_i^{Se} + (1-r) c_i^F$. In this case, if the defender remove the fake protections assigned to i, e_i will decrease and the attacker utility of i will increase. Therefore, the attacker will still choose to attack target i. In addition, the defender utility of target i does not change. Thus there is no need to assign fake protections to target i at first. Actually, as long as there are real resources available, the defender could always exchange some fake protections assigned to i with some real protections assigned to other targets while keeping the attack set unchanged, thus increase her utility. □

Based on Proposition 3 and Proposition 4, we also have the following Proposition.

Proposition 5. *Given attacker belief* $\mathbf{e}$*,* $\Gamma(\mathbf{e})$ *and the attacker's choice of target* i*. The defender should set* $c_i^{Se} = \min\{1, \frac{1-e_i}{1-r}\}$ *and* $c_i^R = \max\{0, \frac{e_i - r}{1-r}\}$ *to achieve the highest utility.*

Proof. First, based on Proposition 3 and Proposition 4, it follows that $c_i^R + r c_i^{Se} = e_i$, $c_j^R + (1-r) c_j^F + r c_j^{Se} = e_i (j \in \Gamma(\mathbf{e}), j \neq i)$. Since the defender utility depends on $c_i^R + c_i^{Se}$ and $c_i^R + c_i^{Se} \leq 1$, to maximize her utility, the defender should set $c_i^R = 0$ and $c_i^{Se} = \frac{e_i}{r}$ if $e_i \leq r$, while setting $c_i^{Se} = \frac{1-e_i}{1-r}$ and $c_i^R = \frac{e_i - r}{1-r}$ if $e_i > r$. Thus we gain the results in Proposition 5.

According to the SSE assumption that the attacker breaks ties in favor of the defender, the attacker's choice should be the target which leads to the highest defender utility under $c_i^{Se} + c_i^R$. □

We introduce R-ORIGAMI (Revised ORIGAMI) in Algorithm 1 to solve SGDB. R-ORIGAMI explores the optimal defender utility when the defender uses only real resources (Lines 1 - 2), when the defender uses all her budget to buy fake resources (Lines 3 - 8), and when the defender converts a real resource into a secret one, and uses the left budget to buy fake resources (Lines 9 - 19). Then returns the optimal defender utility in the game (Line 20). Line 2 computes the coverage $\mathbf{c}$, attacker set $\Gamma(\mathbf{c})$, and defender utility U_R using ORIGAMI when there are n resources. In Line 3, n_1 represents the number of resources from the perspective of the attacker when the defender uses all her real resources, and spends all the budget on fake resources. In

Algorithm 1: R-ORIGAMI

1 Let n be the number of real resources;
2 $\mathbf{c}, \Gamma(\mathbf{c}), U_R \leftarrow ORIGAMI(n)$;
3 $n_1 \leftarrow n + (1-r)\left\lfloor \frac{Bgt}{\beta_F} \right\rfloor$;
4 $\mathbf{e}, \Gamma(\mathbf{e}), g(\mathbf{e}), U_{R\&F} \leftarrow ORIGAMI(n_1)$;
5 $\bar{c}^R_{g(\mathbf{e})} = e_{g(\mathbf{e})}$;
6 **for** $i \in \Gamma(\mathbf{e}), i \neq g(\mathbf{e})$ **do**
7 $\quad \bar{c}^R_i + (1-r)\bar{c}^F_i \leftarrow e_i$;
8 $\bar{\mathbf{c}} \leftarrow \langle \bar{c}^R_i, \bar{c}^F_i : i \in \mathcal{T} \rangle$;
9 $n_2 \leftarrow n - 1 + r + (1-r)\left\lfloor \frac{Bgt - \beta_{Se}}{\beta_F} \right\rfloor$;
10 $\mathbf{e}, \Gamma(\mathbf{e}) \leftarrow ORIGAMI(n_2)$;
11 **for** $i \in \Gamma(\mathbf{e})$ **do**
12 $\quad \tilde{e}_i = \min\{\frac{e_i}{r}, 1\}$;
13 $\quad U_i = \tilde{e}_i U^c_d(i) + (1 - \tilde{e}_i U^u_d(i))$;
14 $U_{R\&F\&Se} \leftarrow \max_{i \in \mathcal{T}} U_i$;
15 $g(\mathbf{e}) \leftarrow \arg_i \max_{i \in \mathcal{T}} U_i$;
16 $\tilde{c}^R_{g(\mathbf{e})} = \max\{0, \frac{e_i - r}{1-r}\}, \tilde{c}^{Se}_{g(\mathbf{e})} = \min\{1, \frac{1-e_i}{1-r}\}$;
17 **for** $i \in \mathcal{T}, i \neq g(\mathbf{e})$ **do**
18 $\quad \tilde{c}^R_i + r\tilde{c}^F_i + (1-r)\tilde{c}^{Se}_i \leftarrow e_i$;
19 $\tilde{\mathbf{c}} \leftarrow \langle \tilde{c}^R_i, \tilde{c}^F_i, \tilde{c}^{Se}_i \rangle$;
20 return $\max\{U_R, U_{R\&F}, U_{R\&F\&Se}\}$;

Line 4, $g(\mathbf{e})$ is the attacker's choice against $\mathbf{e}$. Lines 5 - 8 assign real/fake resources based on Proposition 4. $\bar{\mathbf{c}}$ is the best coverage for the defender when she uses only real resources and fake resources. In Line 9, n_2 represents the number of resources from the perspective of the attacker when the defender converts a real resource into a secret one, then spends the left budget on fake resources. Lines 11 - 15 compute the attacker's choice of target $g(\mathbf{e})$ and the defender utility $U_{R\&F\&Se}$. Lines 16 - 19 compute the assignment of real/fake/secret resources based on Proposition 4 and 5. $\tilde{\mathbf{c}}$ is the best coverage for the defender when she uses all three kinds of resources.

Usefulness of Using Deception

In this section, we analyze the usefulness of using deceptions theoretically and evaluate it experimentally. We first explore the circumstances under which using deceptions leads to higher defender utilities than playing honestly. Specifically, assume that the defender's optimal strategy in a security game without deceptions is to play $\mathbf{c} = \langle c_i \rangle$, which leads the attacker to attack target i and brings the defender a highest utility of $U_d(\mathbf{c}, i)$. Keeping the target set and the payoff structure of the game unchanged while allowing the defender to use deceptions within a budget, we can turn the security game into an SGDB. Assume that in the SGDB, the defender's optimal strategy is to play $\mathbf{c}'$ which leads the attacker to attack target j and brings the defender a highest utility of $U_d(\mathbf{c}', j)$. We now discuss when it follows that $U_d(\mathbf{c}', j) > U_d(\mathbf{c}, i)$. We begin with an observation about the usefulness of fake resources.

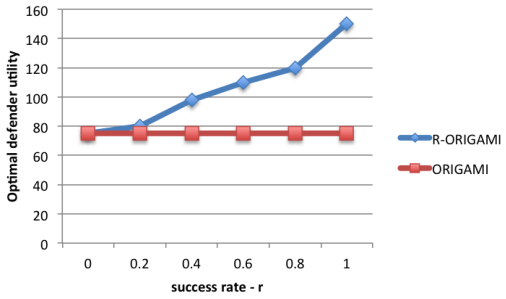

Figure 1: Solution quality against varying success rate

Proposition 6. *Using fake protections as deceptive protections only, $U_d(\mathbf{c}', j) > U_d(\mathbf{c}, i)$ is true as long as $\forall i \in \Gamma(\mathbf{c}), c^R_i < 1$.*

Proof. Adding a fake resource is equivalent to adding $1-r$ real resource in terms of the attacker's observed coverage vector. In the security game without deceptions, the attacker's observed vector is the same as the defender's coverage vector. If for any $i \in \Gamma(\mathbf{c}), c^R_i < 1$, adding real resources will raise the coverage rate of all targets within $\Gamma(\mathbf{c})$. Since the attacker's choice is always covered by real resources, thus the defender utility will increase. □

Assume that the budget could afford to convert a real resource into a secret one and that the defender does not use any fake resources in $\mathbf{c}'$. The following observation shows the usefulness of converting real resources into secret ones.

Proposition 7. *Using secret protections as deceptive protections only, $U_d(\mathbf{c}', j) > U_d(\mathbf{c}, i)$ is true as long as $|\Gamma(\mathbf{c})| > 1$, the attacker's choices in the game without deception and the game with deceptions are the same, namely $i = j$, and in the defender's optimal strategy in the game without deceptions, $c_j < 1$.*

Proof. Converting a real resource into a secret one is equivalent to reducing a total coverage rate of $1-r$ from all targets in the attack set, then adding a coverage rate of $1-r$ to a certain target in the attack set, which is j. If $|\Gamma(\mathbf{c})| > 1$, when the total coverage rate of all targets in the attack set reduces $1-r$, the coverage rate of i reduces less than $1-r$. If $j = i$, adding a coverage rate of $1-r$ will increase the defender utility of j as long as c_j is less than 1. □

We have conducted initial experiments to evaluate the performance of **P1**, R-ORIGAMI and ORIGAMI. Except otherwise specified, there are one real security resource, 2 targets, and the budget is 2. Payoffs are randomly generated. U^c_d and U^u_a are drawn uniformly from the range $[100, 200]$. U^u_d and U^c_a are drawn uniformly from the range $[0, 100]$. The costs of achieving a fake resource (β_F) and converting a real resource into a secret one (β_{Se}) are both 1. The results were averaged over 100 trials.

Figure 1 shows the solution quality of ORIGAMI and R-ORIGAMI when the value of success rate changes. The x-axis indicates the value of $1-r$ while the $y-axis$ indicates

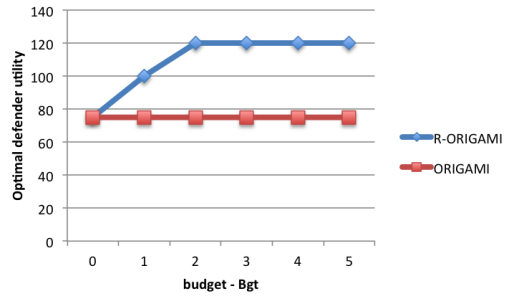

Figure 2: Solution quality against varying budget

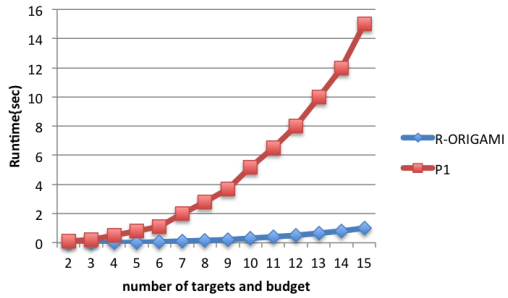

Figure 3: Runtime

the optimal defender utility. As the success rate increases (namely, r decreases), the advantage of R-ORIGAMI over ORIGAMI increases. Figure 2 shows the solution quality of ORIGAMI and R-ORIGAMI when the success rate is fixed at 0.8 (r = 0.2) and the budget varies. The $x-axis$ indicates the value of budgets, the $y-axis$ indicates optimal defender utility. Figure 3 shows the runtime of $\mathbf{P1}$ and R-ORIGAMI. R-ORIGAMI significantly outperforms $\mathbf{P1}$.

Conclusions And Extensions

In this paper, we consider that the defender can perform deceptive protections on targets. Our contributions include: 1) We introduce a model of security games, in which the defender's strategies can include deceptive protections. The deceptive protections succeed with a certain rate. 2) We provide algorithms for computing the optimal defender strategy when the attacker surveils unlimitedly before attacking, and analyze the advantage of using deceptive resources. 3) We conduct some experiments to evaluate the performance of our methods.

Our security game model with deception can be extended in different ways. For example, the attacker may surveil the defender's strategies for a limited number of times as in (An et al. 2013; 2012). An et al (An et al. 2013) has studied the equilibrium when the attacker conducts limited surveillance and the defender has no deceptive resources. They assume that the attacker has a prior over the distribution of defender strategies, surveils the defender's strategies for a limited number of times, then updates the prior based on the observation results. Finally, the attacker chooses the optimal target based on his postior belief of the distribution of defender strategies. If the defender can perform deceptive protections, the attacker's prior belief of the distribution of defender strategies depends on whether the attacker knows the existence of deceptive resources or not. Therefore, the defender needs to consider the distribution of types of the attacker when computing the optimal strategies.

In addition, the assumption made in our model that the attacker does not reason about the deceptions can be relaxed in future work. We can also consider that the budget is not limited, but the defender's objective is to maximize the utility while minimizing the cost. The model can be further expanded by considering robustness issues as in (An et al. 2011b; Yin and Tambe 2012; Pita et al. 2010; Jiang et al. 2013a; 2013b) or the human-agent interaction as in (An et al. 2011a).

Appendix

Assume that if a fake resource fails, the attacker observes a fake protection; if a secret resource fails, the attacker observes a secret protection. Thus the coverage observed by the attacker can be represented as $\mathbf{e}' = \langle e_i^R, e_i^F, e_i^{Se} \rangle$, where e_i^R represents the probability of target i being covered by a real resource observed by the attacker; e_i^F represents the probability of target i being covered by a fake resource; e_i^{Se} represents that of a secret resource. Next, we show that given the defender strategy, this assumption leads to the same attacker strategy as is in our model.

Assume that the defender strategy is to perform coverage $\mathbf{c}$ corresponding to a mixed strategy $\mathbf{x} = \langle x_s \rangle$, while the attacker observes coverage $\mathbf{e}'$. As the proof of Proposition 2, we can prove that the relationship between $\mathbf{e}'$ and $\mathbf{c}$ is as follows.

$$e_i^R = \sum_{s \in \mathcal{S}} x_s \mathbf{1}_R(s_i) + \sum_{s \in \mathcal{S}} (1-r) x_s \mathbf{1}_f(s_i) \quad (16)$$

$$= c_i^R + (1-r) c_i^F \quad (17)$$

$$e_i^F = \sum_{s \in \mathcal{S}} r \times x_s \mathbf{1}_f(s_i) \quad (18)$$

$$= r c_i^F \quad (19)$$

$$e_i^{Se} = \sum_{s \in \mathcal{S}} r \times x_s \mathbf{1}_{Se}(s_i) \quad (20)$$

$$= r c_i^{Se}. \quad (21)$$

Thus the attacker utility of attacking target i from his perspective can be represented as $U_a(\mathbf{e}', i) = (e_i^R + e_i^{Se}) U_d^c(i) + (1 - e_i^R - e_i^{Se}) U_d^u(i) = (c_i^R + (1-r) c_i^F + r c_i^{Se}) U_d^c(i) + (1 - c_i^R - (1-r) c_i^F - r c_i^{Se}) U_d^u(i)$. This is the same as is in our model (In our model, if a fake resource fails, the attacker observes no protection; if a secret resource fails, the attacker observes a real protection). Thus the best attacker response is the same as in our model.

References

Agmon, N.; Urieli, D.; and Stone, P. 2011. Multiagent patrol generalized to complex environmental conditions. In *Pro-*

ceedings of the Twenty-Fifth AAAI Conference on Artificial Intelligence, 1090–1095.

An, B.; Jain, M.; Tambe, M.; and Kiekintveld, C. 2011a. Mixed-initiative optimization in Security Games: A preliminary report. In *AAAI Spring Symposium: Help me help you: Bridging the gaps in human-agent collaboration*, 8–11.

An, B.; Tambe, M.; Ordonez, F.; Shieh, E. A.; and Kiekintveld, C. 2011b. Refinement of Strong Stackelberg Equilibria in Security Games. In *Proceedings of the Twenty-Fifth AAAI Conference on Artificial Intelligence*, 587–593.

An, B.; Kempe, D.; Kiekintveld, C.; Shieh, E.; Singh, S.; Tambe, M.; and Vorobeychik, Y. 2012. Security Games with limited surveillance. In *Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence*, 1241–1248.

An, B.; Brown, M.; Vorobeychik, Y.; and Tambe, M. 2013. Security Games with surveillance cost and optimal timing of attack execution. In *Proceedings of the 12th International Conference on Autonomous Agents and Multi-Agent Systems*, 223–230.

Basilico, N.; Gatti, N.; and Amigoni, F. 2009. Leader-follower strategies for robotic patrolling in environments with arbitrary topologies. In *Proceedings of the 8th International Conference on Autonomous Agents and Multi-Agent Systems*, 57–64.

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.

Eric, S.; An, B.; Yang, R.; Tambe, M.; Baldwin, C.; DiRenzo, J.; Maule, B.; and Meyer, G. 2012a. PROTECT: A deployed game theoretic system to protect the ports of the United States. In *Proceedings of the 11th International Conference on Autonomous Agents and Multi-Agent Systems*, 13–20.

Eric, S.; Bo, A.; Rong, Y.; Milind, T.; Craig, B.; Joseph, D.; Ben, M.; and Garrett, M. 2012b. PROTECT: An application of computational game theory for the security of the ports of the United States. In *Proc. of the 26th AAAI Conference on Artificial Intelligence*, 2173–2179.

Fang, F.; Jiang, A. X.; and Tambe, M. 2013. Optimal patrol strategy for protecting moving targets with multiple mobile resources. In *Proceedings of the 12th International Conference on Autonomous Agents and Multi-Agent Systems*, 957–964.

Jain, M.; Tsai, J.; Pita, J.; Kiekintveld, C.; Rathi, S.; Tambe, M.; and Ordóñez, F. 2010. Software assistants for randomized patrol planning for the LAX Airport Police and the Federal Air Marshal Service. *Interfaces* 40(4):267–290.

Jiang, A. X.; Nguyen, T. H.; Tambe, M.; and Procaccia, A. D. 2013a. Monotonic maximin: A robust Stackelberg solution against boundedly rational followers. In *Decision and Game Theory for Security*. Springer. 119–139.

Jiang, A. X.; Yin, Z.; Zhang, C.; Tambe, M.; and Kraus, S. 2013b. Game-theoretic randomization for security patrolling with dynamic execution uncertainty. In *Proceedings of the 12th International Conference on Autonomous Agents and Multi-Agent Systems*, 207–214.

Kiekintveld, C.; Jain, M.; Tsai, J.; Pita, J.; Ordóñez, F.; and Tambe, M. 2009. Computing optimal randomized resource allocations for massive Security Games. In *Proceedings of the 8th International Conference on Autonomous Agents and Multi-Agent Systems*, 689–696.

Paruchuri, P.; Pearce, J. P.; Marecki, J.; Tambe, M.; Ordóñez, F.; and Kraus, S. 2008. Playing games for security: An efficient exact algorithm for solving Bayesian Stackelberg Games. In *Proceedings of the 7th International Conference on Autonomous Agents and Multi-Agent Systems*, 895–902.

Pita, J.; Jain, M.; Marecki, J.; Ordóñez, F.; Portway, C.; Tambe, M.; Western, C.; Paruchuri, P.; and Kraus, S. 2008. Deployed ARMOR protection: The application of a game theoretic model for security at the Los Angeles International Airport. In *Proceedings of the 7th International Conference on Autonomous Agents and Multi-Agent Systems*, 125–132.

Pita, J.; Jain, M.; Tambe, M.; Ordóñez, F.; and Kraus, S. 2010. Robust solutions to Stackelberg Games: Addressing bounded rationality and limited observations in human cognition. *Artificial Intelligence* 174(15):1142–1171.

Rasmusen, E., and Blackwell, B. 1994. *Games and Information*. Cambridge University Press.

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

Wang, X., and Zhuang, J. 2011. Balancing congestion and security in the presence of strategic applicants with private information. *European Journal of Operational Research* 212(1):100–111.

Yin, Z., and Tambe, M. 2012. A unified method for handling discrete and continuous uncertainty in Bayesian Stackelberg Games. In *Proceedings of the 11th International Conference on Autonomous Agents and Multi-Agent Systems*, 855–862.

Zhuang, J., and Bier, V. M. 2010. Reasons for secrecy and deception in homeland-security resource allocation. *Risk Analysis* 30(12):1737–1743.