

Comparing Strategic Secrecy and Stackelberg Commitment in Security Games

Qingyu Guo¹, Bo An², Branislav Bošanský^{3,4}, Christopher Kiekintveld⁵

¹Joint NTU-UBC Research Centre of Excellence in Active Living for the Elderly, NTU, Singapore

²School of Computer Science and Engineering, Nanyang Technological University, Singapore

³Agent Technology Center, Faculty of Electrical Engineering, Czech Technical University in Prague

⁴Department of Computer Science, Aarhus University

⁵Computer Science Department, University of Texas at El Paso

^{1,2}{qguo005,boan}@ntu.edu.sg,³branislav.bosansky@agents.fel.cvut.cz,⁵cdkiekintveld@utep.edu

Abstract

The *Strong Stackelberg Equilibrium (SSE)* has drawn extensive attention recently in several security domains. However, the SSE concept neglects the advantage of defender’s strategic revelation of her private information, and overestimates the observation ability of the adversaries. In this paper, we overcome these restrictions and analyze the tradeoff between strategic secrecy and commitment in security games. We propose a *Disguised-resource Security Game (DSG)* where the defender strategically disguises some of her resources. We compare strategic information revelation with public commitment and formally show that they have different advantages depending the payoff structure. To compute the *Perfect Bayesian Equilibrium (PBE)*, several novel approaches are provided, including a novel algorithm based on support set enumeration, and an approximation algorithm for ϵ -PBE. Extensive experimental evaluation shows that both strategic secrecy and Stackelberg commitment are critical measures in security domain, and our approaches can efficiently solve PBEs for realistic-sized problems.

1 Introduction

Strong Stackelberg Equilibrium (SSE) has for some time been used to allocate limited security resources to protect targets in many security scenarios including public infrastructures [Kiekintveld *et al.*, 2009; Shieh *et al.*, 2012; Yin *et al.*, 2014; Gan *et al.*, 2015; Zhao *et al.*, 2016] and wildlife [Fang *et al.*, 2016]. In the SSE solution concept, the defender discloses the (possibly randomized) commitment to protect the targets. This commitment is then observed by the attacker who adopts the best response. While in theory it is always favorable for the defender to disclose the commitment, in practice this concept overestimates the surveillance of the attacker [An *et al.*, 2013; Pita *et al.*, 2010] and neglects the fact that certain part of information could have been hidden to the attacker due to imperfect observation [Catchnews, 2016; Gul, 2011]. One example is that the attacker may have been uncertain about the number of the resources of the defender, since the defender can actually use unmarked resources (e.g.,

plainclothes police) to protect the targets. As a consequence, it is no longer clear that disclosing the full commitment is always beneficial for the defender compared to strategically revealing only part of information (termed *strategic secrecy* from now on) – in fact we present an example later in the paper where it is worse for the defender to disclose full commitment compared to strategic secrecy.

We address this discrepancy between the theory and practice and propose a novel class of security games that extends the existing security models in two main aspects: (1) the attacker has an uncertain information about the number of resources of the defender, (2) the defender is allowed to strategically disclose the number of resources.

We model these aspects as *Disguised-resource Security Game (DSG)* and analyze the strategic secrecy where the defender strategically deceives the attacker by disguising some of her resources. The number of the revealed resources is modeled as a *signal* which can only be sent by the defender with enough resources. We compare the expected utility of strategic secrecy with the value of public commitment and formally show that they have different advantages depending on the payoff structure. To be able to evaluate the difference between strategic secrecy and commitment in practice we introduce a collection of novel algorithms to solve for the solution concept based on *Perfect Bayesian Equilibrium (PBE)* [Spence, 1973; Zhuang and Bier, 2011]: (1) we introduce a basic *Mixed-Integer Linear Programming (MILP)* with an exponential number of variables and constraints, (2) an MILP of directly applying compact representation, and (3) a novel approach based on support set enumeration, and (4) an approximation algorithm based on support set enumeration to produce an ϵ -PBE. Finally, we conduct extensive experimental evaluation to show that our algorithms can scale to realistic problems and to examine the fundamental trade-offs between secrecy and public commitment for realistic problems. We conclude that the boundary of such trade-offs is close to zero-sum games which confirms the practical use of plainclothes police due to the approximate zero-sum nature of many security scenarios (e.g., [Banks and Anderson, 2006; Durkota *et al.*, 2015; Haskell *et al.*, 2014; Major, 2002]).

2 Related Work

Previous work on secrecy and deception in security games has failed to address the key dilemma of strategic secrecy




Figure 1: Example (left: strategic secrecy, right: commitment).

and commitment for various reasons. Brown *et al.* [2005] study secrecy in the context of ballistic missile deployment, but assume that the attacker is not aware that the defender can hide resources, so there is no rational possibility for belief update. Other researchers explore signaling games to model a “feint” in homeland security [Hendricks and McAfee, 2006; Oliveros, 2005]. In these games the defender’s resource allocation causes a noisy signal following an uncontrolled signaling technology, in contrast with our model where the defender controls the signals sent. One key feature of the strategic secrecy in our model is that the revealed security resources (signal) are valid information for the attacker, which differs from the cheap talk game [Farrell and Rabin, 1996] where messages are costless and unverifiable.

Recently, there are some literatures studying the information disclosure to persuade the attacker to take the desired action [Rabinovich *et al.*, 2015; Xu *et al.*, 2015] with a strong assumption that the attacker can fully access the defender’s correlated random allocation and signaling scheme by unlimited surveillance, which is unrealistic in the strategic secrecy scenarios. The most closely related model was proposed by Zhuang and Bier [2011], where deception is regarded as a signal to mislead the attacker’s belief about the defender type, while the true defense is treated as a hidden action. Their model assumes that the defender deterministically sends the signal, while we allow the randomized signaling strategy, which is possible and more general in resource allocation domain. Furthermore, they only provide general results for high-level special cases without providing efficient algorithms for realistic problems.

3 Motivating Example

Suppose a small police station (the defender) has three districts to protect, A , B and C . The police station has either one or two patrol units, depending on the day. We call these cases the *weak* and *strong* types, and assume that they are equally likely for the example. An attacker will choose one of the three districts to target, represented by a mixed attacking strategy $\mathbf{a} = \langle a_A, a_B, a_C \rangle$ where a_t denotes the probability of attacking district $t \in \{A, B, C\}$. The police strategy can be compactly represented by a coverage vector $\mathbf{c} = \langle c_A, c_B, c_C \rangle$ such that c_t is the probability that district t is covered by a patrol unit. There are four payoffs associated with each district, $\langle R_t^d, P_t^a, P_t^d, R_t^a \rangle$: if a resource is allocated to attacked district t , then defender receives a reward R_t^d and the attacker receives a penalty P_t^a ; otherwise the payoffs are P_t^d and R_t^a respectively. The game is shown in Figure 1, where A and B are homogeneous districts with the same pay-

offs and $R_A^a = -P_A^d = 12$, and $R_C^a = -P_C^d = 2$, while $R_t^d = P_t^a = 0$ for $t \in \{A, B, C\}$. Consider that the strong type (with 2 resources) can disguise one resource, or choose to commit herself by revealing both resources.

If the defender chooses to reveal the type, the scenario can be captured by the strong Stackelberg equilibrium (SSE), widely adopted in the security game literature, where the attacker can observe the defender strategy with extensive surveillance. We will present the formal equilibrium concept in the next section. For now, it is easy to see that in the equilibrium (shown in the game tree on the right in Figure 1), the weak type will play the coverage vector $\langle 0.5, 0.5, 0 \rangle$ equalizing the attacker’s expected payoff between A and B , resulting in an expected utility of -6 . The strong type will play $\langle \frac{7}{8}, \frac{7}{8}, \frac{1}{4} \rangle$ and receive -1.5 .

It is worthwhile to note that the assumption of the attacker knowing the defender strategy is critical to the advantage of commitment. However, as we discussed in Section 1, the extensive surveillance assumption is actually unrealistic and thus cannot capture the problem here. Even though, the defender can still do better with strategic secrecy, taking into account the limited observation of the attacker. If the defender is allowed to strategically disguise a resource, we have a different game, shown in the game tree on the left in Figure 1. Suppose that the strong type hides one resource with 100% probability. When the attacker observes only 1 patrol unit, he cannot know the type he is facing. Thus, the attacker will update his belief on the defender type and it turns out that he is playing against both types with equal probability according to Bayes’ rule. Now the weak type still plays coverage $\langle 0.5, 0.5, 0 \rangle$, while the strong type plays $\langle 1, 1, 0 \rangle$ and attacker receives an expected payoff of 3 for attacking A or B and plays a mixed strategy $\langle 0.5, 0.5, 0 \rangle$. We can verify that the equilibrium is formed where no player has incentive to deviate. As we will formally present in the next section, such equilibrium is the Perfect Bayesian equilibrium (PBE) in extensive-form games. In the equilibrium, the weak type still receives -6 , while the strong type gets 0, better than -1.5 with commitment. One thing to notice here is that the attacker *does not* know the probability of the strong type disguising one resource, and the reason of updating the posterior belief according to Bayes’ rule is that as long as both players agree on the equilibrium, the attacker can infer such probability and update the belief, which is the key spirit of Nash equilibrium.

4 Disguised-Resource Security Games (DSG)

A DSG extends the structure of a Stackelberg security game [Kiekintveld *et al.*, 2009], but adds a way to model the defender holding private information about the number of resources. The game is played by a defender and an attacker. The defender protects a set of targets T and the attacker chooses a target $t \in T$ to attack. There are multiple defender types and each one has a different number of available resources to protect the targets. We note that this abuses the terminology “type” a bit, since we will assume that all types have the same utility function, but effectively have a different strategy spaces. We use $\theta \in \Theta$ to represent the number of resources available to each defender type (e.g., police teams,

patrol boats). The prior probability distribution over types $p : \Theta \rightarrow [0, 1]$ is known to both players.

In DSG, the defender is allowed to publicly reveal only a subset of her available resources and the number of revealed resources is denoted as a *signal*.¹ W.l.o.g., we assume that the signal is in the set Θ . The remaining resources are disguised, as in the case of a plainclothes police officer or unmarked vehicle. Importantly, in our model the defender *cannot* send invalid signals that claim a greater number of resources than are actually available, so a defender of type θ can only send a signal $s \leq \theta$. The payoffs follow the same definition in Section 3. Consistently with existing work on security games $R_t^d > P_t^d$ and $R_t^a > P_t^a$. The DSG is an extensive-form game which proceeds as illustrated by the motivating example: Nature randomly chooses one type for the defender and let it be the strong type $\theta = 2$. The strong type draws a signal according to the mixed signaling strategy which turns out to be $s = 1$. The attacker cannot distinguish the two decision nodes corresponding to both types sending the same signal $s = 1$, and the set of these decision nodes is called an *information set* denoted by $I(s)$.

Strategies: Let $\mathbf{c} = \langle c_t \rangle$ and $\mathbf{a} = \langle a_t \rangle$ denote the defender's coverage strategy and attacker's mixed attacking strategy as defined in Section 3. Let $\mathbf{o} = \langle o_s \rangle$ denote the mixed signaling strategy such that o_s represents the probability of sending signal s . Let $\Delta_c^\theta = \{\mathbf{c} \in [0, 1]^{|T|} : \|\mathbf{c}\|_1 = \theta\}$ denote the set of coverage strategies available for defender type θ and $\Delta_c = \bigcup_{\theta \in \Theta} \Delta_c^\theta$ be the set of all coverage strategies. Similarly, we denote by $\Delta_o^\theta = \{\mathbf{o} \in [0, 1]^{|S|} : \|\mathbf{o}\|_1 = 1, o_s = 0 \ \forall s > \theta\}$ the set of mixed signaling strategies available for defender type θ and $\Delta_o = \bigcup_{\theta \in \Theta} \Delta_o^\theta$. Let $\Delta_a = \{\mathbf{a} \in [0, 1]^{|T|} : \|\mathbf{a}\|_1 = 1\}$ represent the set of all mixed attacking strategies. Let $\pi_d = \langle \pi_c, \pi_o \rangle$ denote the defender policy where $\pi_c : \Theta \times \Theta \rightarrow \Delta_c$ is the coverage policy such that $\pi_c(\theta, s) \in \Delta_c^\theta$ denotes the coverage strategy adopted by defender type θ conditioned on sending signal s and $\pi_c(t|\theta, s)$ is the corresponding marginal coverage on target t , and $\pi_o : \Theta \rightarrow \Delta_o$ is the signaling policy with $\pi_o(\theta) \in \Delta_o^\theta$ representing the mixed signaling strategy for defender type θ and $\pi_o(s|\theta)$ being the corresponding probability of sending signal s . In particular, we use $\pi_c(\theta) = \langle \pi_c(\theta, s) \rangle$ to denote the coverage policy for defender type θ . Let $\pi_a : \Theta \rightarrow \Delta_a$ denote the attacker policy such that $\pi_a(s) \in \Delta_a$ is the mixed attacking strategy adopted by the attacker observing signal s , where the corresponding probability of attacking target t is denoted by $\pi_a(t|s)$.

Posterior Belief: Let $\Delta_\Theta = \{\langle \delta_\theta \rangle : \|\delta\|_1 = 1\}$ be the set of all possible probability distributions over Θ . We denote by $\mu : \Theta \rightarrow \Delta_\Theta$ the attacker's posterior belief on the defender type conditioned on the received signal. In particular, $\mu(\theta|s)$ denotes the posterior probability of defender type being θ at information set $I(s)$. Apparently, $\mu(\theta|s) = 0$ for $\theta < s$. If $I(s)$ is on the equilibrium path, i.e., s is sent with positive probability ($\sum_{\theta: \theta \geq s} p_\theta \pi_o(s|\theta) > 0$), the belief is determined

by the Bayes' rule, such that:

$$\mu(\theta|s) = p_\theta \pi_o(s|\theta) / \sum_{\theta': \theta' \geq s} p_{\theta'} \pi_o(s|\theta').$$

Otherwise, if $I(s)$ is off the equilibrium path, such as $I(s = 2)$ in the motivating example where no type sends such signal, we adopt the *optimistic conjecture* [Rubinstein, 1985]. That is to say, when the defender acts off the equilibrium strategy, the attacker believes the defender of her weakest type, against which the attacker would gain the most. Intuitively the attacker always prefers to play against a defender type with fewer resources, which is confirmed by Theorem 1. Thus, at information set $I(s)$ which is off equilibrium path, we have: $\mu(s|s) = 1$ and $\mu(\theta|s) = 0$ for all $\theta > s$. Throughout the paper, we assume that posterior belief μ follows Bayes' rule and the optimistic conjecture for information sets on and off equilibrium path respectively.

Utilities: Given the defender coverage strategy $\mathbf{c} \in \Delta_c$ and the mixed attacking strategy $\mathbf{a} \in \Delta_a$, the expected payoffs of both players are defined as follows:

$$\begin{aligned} P_d(\mathbf{c}, \mathbf{a}) &= \sum_{t \in T} a_t c_t (R_t^d - P_t^d) + a_t P_t^d \\ P_a(\mathbf{c}, \mathbf{a}) &= \sum_{t \in T} a_t c_t (P_t^a - R_t^a) + a_t R_t^a. \end{aligned} \quad (1)$$

Given the defender's policy $\pi_d = \langle \pi_c, \pi_o \rangle$ and attacker's policy π_a , the expected utility of the attacker conditioned on receiving signal s , and the expected utility of the defender type θ are defined as follows:

$$\begin{aligned} U_d(\pi_c(\theta), \pi_o(\theta), \pi_a) &= \sum_{s: s \leq \theta} \pi_o(s|\theta) P_d(\pi_c(\theta, s), \pi_a(s)) \\ U_a(\pi_c, \pi_o, \pi_a(s)) &= \sum_{\theta: \theta \geq s} \mu(\theta|s) P_a(\pi_c(\theta, s), \pi_a(s)). \end{aligned}$$

Theorem 1. For two defender types θ and θ' such that $\theta > \theta'$, suppose $\langle \mathbf{c}, \mathbf{a} \rangle$ is a Nash equilibrium between the attacker and defender type θ , then there always exists an NE profile $\langle \mathbf{c}', \mathbf{a}' \rangle$ between the attacker and defender type θ' such that $P_a(\mathbf{c}', \mathbf{a}') \geq P_a(\mathbf{c}, \mathbf{a})$.

Proof. W.l.o.g, assume $c_t < 1$ for each $t \in T$. According to the defender's best response criteria, $c_t > 0$ only for t in the support set of \mathbf{a} . We iteratively construct a coverage vector \mathbf{c}' and a set of targets T' as follows: i) initially, $\mathbf{c}' = \mathbf{c}$, and $T' = \{t \in T : a_t > 0\}$; ii) for each target $t \in T'$, we decrease c'_t a bit by a small amount value δ_t which is proportional to $1/(R_t^a - P_t^a)$ such that $(R_t^a - P_t^a)\delta_t = (R_{t'}^a - P_{t'}^a)\delta_{t'}$ for any $t, t' \in T'$; and iii) if for some target $t \in T'$, $c'_t \leq 0$, we remove t from T' and set c'_t to 0; iv) the procedure terminates as long as $\sum_{t \in T} c'_t = \theta'$. At this point, we obtain a coverage strategy \mathbf{c}' for the defender type θ' covering T' such that targets in T' are all best response targets for the attacker. Let \mathbf{a}' be the mixed attacking strategy such that $a'_t = \lambda/(R_t^d - P_t^d)$ for $t \in T'$ and $a'_t = 0$ otherwise, where $\lambda = 1/\sum_{t \in T'} \frac{1}{R_t^d - P_t^d}$. It can be easily verified that strategy profile $\langle \mathbf{c}', \mathbf{a}' \rangle$ forms an NE, and $P_a(\mathbf{c}', \mathbf{a}') > P_a(\mathbf{c}, \mathbf{a})$. \square

Equilibrium Concepts: Analogous to the equilibrium of extensive-form game with first-mover hidden action-s [Zhuang and Bier, 2011], the solution concept adopted for

¹The intuition of such definition is that the attacker is more likely to observe the number of revealed resources due to the limited observation.

DSG is based on PBE, which is the profile $\langle \pi_d^*, \pi_a^* \rangle$ where:

$$\langle \pi_c^*(\theta), \pi_o^*(\theta) \rangle = \arg \max_{\pi_c(\theta), \pi_o(\theta)} U_d(\pi_c(\theta), \pi_o(\theta), \pi_a^*), \forall \theta$$

$$\pi_a^*(s) = \arg \max_{\pi_a(s)} U_a(\pi_c^*, \pi_o^*, \pi_a(s)), \forall s.$$

Due to the strict requirement in PBE that both players will not play a suboptimal response strategy, the computation of PBE is extremely challenging. Thus, we also consider the ϵ -PBE, an approximation of PBE that allows players to have a small incentive to play strategies other than the one played in the equilibrium. Formally, an ϵ -PBE is a strategy profile $\langle \pi_d^*, \pi_a^* \rangle$ where: i) $U_d(\pi_c^*(\theta), \pi_o^*(\theta), \pi_a^*) \geq U_d(\pi_c(\theta), \pi_o(\theta), \pi_a^*) - \epsilon, \forall \theta \in \Theta, \langle \pi_c(\theta), \pi_o(\theta) \rangle$ and ii) $U_a(\pi_c^*, \pi_o^*, \pi_a^*(s)) \geq U_a(\pi_c^*, \pi_o^*, \pi_a(s)) - \epsilon, \forall s \in \Theta, \pi_a(s)$.

The SSE [Leitmann, 1978] between the defender and attacker is a pair of strategies $\langle \mathbf{c}, f(\mathbf{c}) \rangle$ where: i) $P_d(\mathbf{c}, f(\mathbf{c})) \geq P_d(\mathbf{c}', f(\mathbf{c}')), \forall \mathbf{c}'$, ii) $P_a(\mathbf{c}, f(\mathbf{c})) \geq P_a(\mathbf{c}, \mathbf{a}), \forall \mathbf{a}$; and iii) the attacker breaks ties in favor of defender: $P_d(\mathbf{c}, f(\mathbf{c})) \geq P_d(\mathbf{c}, \mathbf{a})$ for all optimal attacking strategies \mathbf{a} .

Given these formal definitions, we can verify that the game trees in Figure 1 correspond with PBE and SSE respectively, and the defender utility is higher in PBE. However, disguising resources is not always beneficial. Consider the same game in the motivating example. Suppose the payoffs are changed by setting $R_C^d = 8$ and the game is no longer zero-sum. The PBE remains unchanged, as well as the SSE for the weak type. In SSE for the strong type, coverage $\langle \frac{7}{8}, \frac{7}{8}, \frac{1}{4} \rangle$ is still played. However, the attacker will attack C which brings the strong type an expected utility of 0.5, higher than 0 in PBE. We can see that the benefit of PBE is sensitive to the correlation between defender and attacker payoffs and the strategic secrecy is preferred when the game is ‘‘close’’ to zero-sum. Our next section affirms such conjecture with theoretical analysis.

5 PBE versus SSE

For zero-sum DSGs where the players’ payoffs are perfectly correlated, we prove that any PBE gives the defender utility at least as high as SSE (Theorem 2). The intuition for Theorem 2 is that for zero-sum DSGs, the defender cannot benefit from public commitment and SSE reduces to NE, while in PBE the attacker cannot distinguish the defender type, so he may not play the best response to each individual type, and the defender can take the advantage.

Theorem 2. *For a zero-sum DSG, given any PBE $\langle \pi_d^*, \pi_a^* \rangle$ and SSE $\langle \mathbf{c}^\theta, \mathbf{a}^\theta \rangle$ formed by defender type θ and the attacker, we have: $U_d(\pi_c^*(\theta), \pi_o^*(\theta), \pi_a^*) \geq P_d(\mathbf{c}^\theta, \mathbf{a}^\theta), \forall \theta$.*

Proof. SSE is equivalent with NE in zero-sum games. Given that \mathbf{a}^θ is the best response against \mathbf{c}^θ and vice versa, and $\pi_c^*(\theta, s)$ is the best response against $\pi_a^*(s)$, we have:

$$P_d(\mathbf{c}^\theta, \mathbf{a}^\theta) = -P_a(\mathbf{c}^\theta, \mathbf{a}^\theta) \leq -P_a(\mathbf{c}^\theta, \pi_a^*(s))$$

$$= P_d(\mathbf{c}^\theta, \pi_a^*(s)) \leq P_d(\pi_c^*(\theta, s), \pi_a^*(s)).$$

Thus:

$$U_d(\pi_o^*(\theta), \pi_c^*(\theta), \pi_a^*) = \sum_{s \leq \theta} \pi_o^*(s|\theta) P_d(\pi_c^*(\theta, s), \pi_a^*(s))$$

$$\geq P_d(\mathbf{c}^\theta, \mathbf{a}^\theta)$$

□

On the other hand, we analyze a PBE in the special case where all types have a similar number of resources, and show that this PBE has defender utility less than or equal to SSE (Theorem 3). The idea is that when all types have a similar number of resources, it is likely that there exists a set of targets $T' = \{t_1, \dots, t_k\}$, such that for each type θ , the NE coverage strategy has support set T' . In this case, we can show the existence of a PBE where each defender type is playing that NE coverage strategy regardless of signals, and the corresponding defender’s expected utility is equal to her expected utility in NE, which is no higher than that in SSE. Since Theorem 3 has no restriction on the correlation between defender and attacker payoffs, with less correlation the defender benefits more from commitment and the defender utility in SSE can be much higher than that in the PBE in Theorem 3.

Theorem 3. *Let the targets be listed by R_t^a with descending order: $T = \{t_1, \dots, t_{|T|}\}$. If there exists k such that: $\sum_{l=1}^k \frac{R_{t_l}^a - R_{t_{k+1}}^a}{R_{t_l}^a - P_{t_l}^a} \leq \theta \leq \sum_{l=1}^k \frac{R_{t_l}^a - R_{t_{k+1}}^a}{R_{t_l}^a - P_{t_l}^a}$ holds for any type θ , then there exists a PBE $\langle \pi_d^*, \pi_a^* \rangle$ such that $U_d(\pi_c^*(\theta), \pi_o^*(\theta), \pi_a^*) \leq P_d(\mathbf{c}^\theta, \mathbf{a}^\theta)$ for any type, where $\langle \mathbf{c}^\theta, \mathbf{a}^\theta \rangle$ is an SSE between type θ and the attacker.*

Proof. Suppose such a k exists. For any type θ , define u_θ as follows:

$$u_\theta = \frac{\sum_{l=1}^k \frac{R_{t_l}^a}{R_{t_l}^a - P_{t_l}^a} - \theta}{\sum_{l=1}^k \frac{1}{R_{t_l}^a - P_{t_l}^a}}$$

and we have:

$$R_{t_{k+1}}^a \leq u_\theta \leq R_{t_k}^a \quad \forall \theta \in \Theta$$

according to the inequality in the theorem. Consider the strategy profile $\langle \pi_d^*, \pi_a^* \rangle$ where:

$$\pi_c^*(t|\theta, s) = \begin{cases} \frac{R_t^a - u_\theta}{R_t^a - P_t^a}, & t \in \{t_1, \dots, t_k\}; \\ 0, & \text{otherwise.} \end{cases}$$

$$\pi_a^*(t|s) = \begin{cases} \frac{\lambda}{R_t^d - P_t^d}, & t \in \{t_1, \dots, t_k\}; \\ 0, & \text{otherwise.} \end{cases}$$

$$\lambda = 1 / \sum_{l=1}^k \frac{1}{R_{t_l}^d - P_{t_l}^d}.$$

We have: $P_a(\pi_c^*(\theta, s), t) = u_\theta$ for $t \in \{t_1, \dots, t_k\}$. In other words, $\pi_a^*(s)$ is best response against each individual coverage vector $\pi_c^*(\theta, s)$ for all $\theta, s \in \Theta : \theta \geq s$. According to the definition of SSE, we have:

$$P_d(\mathbf{c}^\theta, \mathbf{a}^\theta) \geq P_d(\pi_c^*(\theta, s), \pi_a^*(s)) \quad \forall \theta, s \in \Theta : \theta \geq s$$

$$P_d(\mathbf{c}^\theta, \mathbf{a}^\theta) \geq U_d(\pi_c^*(\theta), \pi_a^*) \quad \forall \theta \in \Theta.$$

□

6 Computing PBE Solutions

We now introduce computation methods for computing PBE. We first try an MILP based on mixed defender strategy representation which is a variant of the sequence-form MILP for extensive-form games. This approach is not scalable due to

the exponential number of pure strategies, even with implementation of constraint-generation approach. To reduce the strategy space we directly apply the compact representation (coverage), and propose another MILP with a polynomial number of variables and constraints. However, the defender's best response criteria turn out to be nontrivial. Although they can be linearly represented based on the *complementary slackness conditions* [Bertsimas and Tsitsiklis, 1997], the auxiliary binary variables and logistic constraints make the MILP not scalable. We omit the formulations and experiments of these (failed) approaches for the ease of reading.

To produce a scalable solution, we further investigate the special structure of PBE. We start with PBEs where the attacker policy is *unbiased* (Definition 1 given later), which makes the defender's best response criteria much easier to represent. We then propose a concise and scalable formulation to compute such PBEs based on support set enumeration. In case no such PBE exists, the formulation is modified to compute the ϵ -PBE instead. The experimental evaluation shows that in almost all cases, our approach can compute a PBE, and a high-quality approximate ϵ -PBE in the remaining cases. We now give a definition of an unbiased attacker strategy, followed by the support set enumeration approach.

Definition 1. A mixed attacking strategy \mathbf{a} with support set T' is called unbiased if $a_t = \lambda_{T'}/(R_t^d - P_t^d)$ for all $t \in T'$ and $a_t = 0$ otherwise, where $\lambda_{T'} = 1/\sum_{t \in T'} \frac{1}{R_t^d - P_t^d}$. The attacker's policy π_a is unbiased if the mixed attack strategy $\pi_a(s)$ is unbiased for each $s \in \Theta$.

There exists one and only one mixed attacking strategy with support set T' , which is unbiased, for any $T' \subseteq T$. Therefore, we denote such strategy with support set T' as $\mathbf{a}^{T'}$. The unbiased attacking strategy is not trivial. In fact, it follows the spirit of NE to make the defender strategy the best response. In particular, the defender's expected payoff against $\mathbf{a}^{T'}$ is:

$$P_d(\mathbf{c}, \mathbf{a}^{T'}) = \lambda_{T'} \sum_{t \in T'} c_t + \lambda_{T'} \sum_{t \in T'} P_t^d / (R_t^d - P_t^d).$$

For defender of type θ , \mathbf{c} is the best response against $\mathbf{a}^{T'}$ if and only if: $\sum_{t \in T'} c_t = \min\{\theta, |T'|\}$, and the corresponding defender's optimal expected payoff is denoted as:

$$P_{\theta T'}^d = \lambda_{T'} \min\{\theta, |T'|\} + \sum_{t \in T'} \lambda_{T'} P_t^d (R_t^d - P_t^d),$$

Support Set Enumeration: Let \mathcal{T} denote the set of all subsets of T . The intuition of support set enumeration is as follows. Suppose there exists a PBE profile where the attacker strategy is unbiased. To compute such a PBE, a naïve way is to consider all possible unbiased attacker policies, which is of size $|\mathcal{T}|^{|\Theta|}$ as there are $|\Theta|$ information sets. For each unbiased attacker's policy π_a , let $T'_s \in \mathcal{T}$ be the support set of $\pi_a(s)$. We can easily verify whether there exists a PBE where the attacker's policy is π_a with linear constraints, since the defender's best response can be easily ensured by:

$$\sum_{t \in T'_s} \pi_c(t|\theta, s) = \min\{\theta, |T'_s|\} \quad \forall \theta, s \in \Theta : \theta \geq s, (2)$$

However, the size of \mathcal{T} is exponential ($2^{|T|}$), which makes it impossible to generate all possible unbiased attacker policies. Fortunately, we do not necessarily need to generate all

of them due to a nice property of the PBE $\langle \pi_d, \pi_a \rangle$ such that there are only limited subsets of T able to serve as the support set of $\pi_a(s)$ no matter if $\pi_a(s)$ is unbiased or not (Lemma 1 & Theorem 4). Thus, we only consider a small subset $\mathcal{T}' \subset \mathcal{T}$, and the property of PBE ensures that \mathcal{T}' is enough to search for a PBE with unbiased attacker strategy. (We will discuss how to generate \mathcal{T}' later.) As such, instead of brute force search, an MILP with no objective function is provided:

$$\sum_{s \in \Theta: s \leq \theta} S_{\theta s} = 1 \quad \forall \theta \quad (3a)$$

$$\sum_{t \in T} \tilde{C}_{\theta st} = \theta S_{\theta s} \quad \forall \theta > s \quad (3b)$$

$$0 \leq \tilde{C}_{\theta st} \leq S_{\theta s} \quad \forall \theta > s, t \quad (3c)$$

$$\sum_{t \in T} \tilde{C}_{sst} = \theta(S_{ss} + 1 - \chi_s) \quad \forall s \quad (3d)$$

$$0 \leq \tilde{C}_{sst} \leq S_{ss} + 1 - \chi_s \quad \forall s, t \quad (3e)$$

$$\chi_s \in \{0, 1\}$$

$$\delta \chi_s \leq \sum_{\theta \in \Theta: \theta \geq s} p_{\theta} S_{\theta s} \leq \chi_s \quad \forall s \quad (3f)$$

$$\phi_{sT'} \in \{0, 1\} \quad \forall s, T' \quad (3g)$$

$$\sum_{T' \in \mathcal{T}'} \phi_{sT'} = 1 \quad \forall s \quad (3h)$$

$$x_{st} + \sum_{\theta \in \Theta: s \leq \theta} p_{\theta} \tilde{C}_{\theta st} P_t^a +$$

$$\sum_{\theta \in \Theta: s \leq \theta} p_{\theta} (S_{\theta s} - \tilde{C}_{\theta st}) R_t^a +$$

$$p_s (1 - \chi_s) R_t^a = v_s^a \quad \forall s, t \quad (3i)$$

$$0 \leq x_{st} \leq (1 - \sum_{T' \in \mathcal{T}': t \in T'} \phi_{sT'}) M \quad \forall s, t \quad (3j)$$

$$\tilde{C}_{\theta st} \leq 1 - \sum_{T': t \notin T', |T'| \geq \theta} \phi_{sT'} \quad \forall \theta, s, t \quad (3k)$$

$$\varphi_{\theta s} \in \{0, 1\}$$

$$0 \leq S_{\theta s} \leq \varphi_{\theta s}$$

$$y_{\theta s} + \sum_{T' \in \mathcal{T}'} P_{\theta T'}^d \phi_{sT'} = v_{\theta}^d \quad \forall \theta, s \quad (3l)$$

$$0 \leq y_{\theta s} \leq (1 - \varphi_{\theta s}) M \quad \forall \theta, s. \quad (3m)$$

In (3), S is the decision variable corresponding with π_o such that $S_{\theta s} = \pi_o(s|\theta)$; \tilde{C} is the decision variable defined as follows: $\tilde{C}_{\theta st} = \pi_o(s|\theta)\pi_c(t|\theta, s)$ if $I(s)$ is on the equilibrium path, otherwise $\tilde{C}_{sst} = \pi_c(t|s, s)$ and $\tilde{C}_{\theta st} = 0$ for any $\theta > s$; Binary variable $\chi_s = 1$ if and only if $I(s)$ is on the equilibrium path; δ in (3f) is a small enough constant as the threshold of probability of sending s , while M is a large enough constant; x_{st} is the slack variable which takes zero when target t is in the support set of $\pi_a(s)$; Variable v_{θ}^d denotes $U_d(\pi_o(\theta), \pi_c(\theta), \pi_a)$; $y_{\theta s}$ is the slack variable which takes zero when s is in the support set of $\pi_o(\theta)$; Binary variable $\phi_{sT'} = 1$ iff T' is the support set of $\pi_a(s)$. (3i) and (3j) correspond to the attacker's best response criteria, including the optimistic conjecture, such that at information set $I(s)$, the expected utility of attacking a target in the support set of $\pi_a(s)$ is the highest among all targets. (3k) ensures that $\pi_c(\theta, s)$ is the best response coverage against $\pi_a(s)$ as required by (2). In particular, given the support set

of $\pi_a(s)$ being T' , we have $\pi_c(t|\theta, s) = 0$ for $t \notin T'$.² Finally, (31)–(3m) ensure that the defender is playing the best response signaling strategy such that $\pi_o(s|\theta) > 0$ only if $P_d(\pi_c(\theta, s), \pi_a(s)) \geq P_d(\pi_c(\theta, s'), \pi_a(s'))$ for any $s' \leq \theta$.

The MILP (3) returns a PBE if and only if there exists one with unbiased attacker's policy π_a whose support sets are in \mathcal{T}' . If no such PBE exists, we slightly modify MILP (3) to compute the ϵ -PBE instead, with the MILP (4), which is the same as MILP (3) except: i) the expected utility of attacking target t in support set of $\pi_a(s)$ is no lower than the best response minus ϵ ; and ii) $\pi_o(s|\theta) > 0$ only if $P_d(\pi_c(\theta, s), \pi_a(s)) \geq P_d(\pi_c(\theta, s'), \pi_a(s')) - \epsilon$ for any $s' \leq \theta$. The feasible solution of MILP (4) is ensured to be an ϵ -PBE. Notice that ϵ in MILP (4) is a constant number instead of a variable since otherwise the formulation becomes non-convex. Thus, to get the best approximation, the binary search on ϵ is conducted.

$$\begin{aligned} & (3a) - (3i), (3k) - (3l) \\ 0 \leq x_{st} \leq & (1 - \sum_{T' \in \mathcal{T}: t \in T'} \phi_{sT'})M + \\ & \epsilon \sum_{\theta: \theta \geq s} p_\theta S_{\theta s} + \epsilon p_s (1 - \chi_s) \quad \forall s, t \\ 0 \leq y_{\theta s} \leq & (1 - \varphi_{\theta s})M + \epsilon \quad \forall \theta, s. \end{aligned} \quad (4)$$

Generating Support Sets: List the targets by R_t^a in descending order: $T = \{t_1, \dots, t_{|T|}\}$. Our next Lemma shows that the support set T' of $\pi_a(s)$ in PBE must contain the first $|T'|$ targets in T .² The intuition is that the defender will always cover the targets in T' with the highest priority, otherwise the attacker's best response is violated.

Lemma 1. *In PBE, the support set of the mixed attack strategy $\pi_a(s)$ at any $I(s)$ has the form: $T' = \{t_1, \dots, t_{|T'|}\}$.*

Proof. It is sufficient to prove that if $t \in T'$, then any target with higher reward is also in the support set. We prove by contradiction. Suppose there exists a target t' with $R_{t'}^a > R_t^a$ which is not in the support set T' . One reasonable assumption adopted here (and throughout the paper) is that in PBE, there exists no defender type which fully covers every target in the support set of the mixed attacking strategy. Thus, we have:

$$\pi_c(t|\theta, s) \geq \pi_c(t'|\theta, s) = 0 \quad \forall \theta, s \in \Theta : \theta \geq s \quad (5)$$

since otherwise the defender type θ can always decrease coverage $\pi_c(t'|\theta, s)$ and increase coverage $\pi_c(t|\theta, s)$ to improve the expected utility, which contradicts to the best response criteria. Let $U_a(\pi_d, s, t)$ denote the attacker's expected utility of attacking target t at information set $I(s)$:

$$U_a(\pi_d, s, t) = \sum_{\theta \in \Theta: \theta \geq s} \mu(\theta|s) [R_t^a - (R_t^a - P_t^a) \pi_c(t|\theta, s)]. \quad (6)$$

According to Eqs.(5) & (6), we have:

$$U_a(\pi_d, s, t) \leq R_t^a < R_{t'}^a = U_a(\pi_d, s, t').$$

Contradict to that t is attacker's best response target. \square

²We assume that in PBE, there exists no defender type which fully covers every target in the support set of mixed attacking strategy which is reasonable for realistic payoffs.

Although Lemma 1 already restricts the number of support sets to $|T|$, we can further eliminate some of them with Theorem 4. The intuition of Theorem 4 is that a defender with more resources can cover more targets while keeping them all the best response targets for the attacker, and k and K are the minimal and maximal numbers of such targets respectively. Therefore, the size of support set of $\pi_a(s)$ against defender of unknown type is within interval $[k, K]$. Notice that when $\theta_{min} \approx \theta_{max}$, we have: $|T'| = |K - k + 1| \ll |T|$.

Theorem 4. *Let k and K be the smallest and the largest values respectively of i such that there exists a type θ satisfying $\sum_{t=1}^i \frac{R_t^a - R_{t_i}^a}{R_t^a - P_t^a} \leq \theta \leq \sum_{t=1}^i \frac{R_t^a - R_{t_{i+1}}^a}{R_t^a - P_t^a}$. In PBE, the support set T' of $\pi_a(s)$ at any $I(s)$ satisfies: $k \leq |T'| \leq K$.*

Proof. Obviously, k corresponds to the smallest i such that:

$$\sum_{t=1}^i \frac{R_t^a - R_{t_i}^a}{R_t^a - P_t^a} \leq \theta_{min} \leq \sum_{t=1}^i \frac{R_t^a - R_{t_{i+1}}^a}{R_t^a - P_t^a}$$

while K corresponds to the largest i such that:

$$\sum_{t=1}^i \frac{R_t^a - R_{t_i}^a}{R_t^a - P_t^a} \leq \theta_{max} \leq \sum_{t=1}^i \frac{R_t^a - R_{t_{i+1}}^a}{R_t^a - P_t^a}$$

Similar to the proof of Theorem 3, there exists a coverage vector \mathbf{c} for type θ_{min} which only covers $\{t_1, \dots, t_k\}$ and makes all of $\{t_1, \dots, t_k\}$ the best response targets for the attacker; Also, there exists a coverage vector \mathbf{c}' for type θ_{max} which only covers $\{t_1, \dots, t_K\}$ and makes all of $\{t_1, \dots, t_K\}$ the best response targets for the attacker.

We prove by contradiction with the above consequences. Suppose there exists a support set T' of $\pi_a(s)$ such that $|T'| < k$, according to Lemma 1, we have: $T' = \{t_1, \dots, t_{|T'|}\}$. Let t_l be an arbitrary target in T' . Since $\{t_1, \dots, t_l, \dots, t_k\}$ are all best response targets for the attacker against \mathbf{c} of θ_{min} :

$$R_{t_k}^a - (R_{t_k}^a - P_{t_k}^a)c_{t_k} = R_{t_l}^a - (R_{t_l}^a - P_{t_l}^a)c_{t_l} \leq R_{t_k}^a \quad (7)$$

In PBE, since $\pi_c(\theta, s)$ is the best response against $\pi_a(s)$ with support T' , where $t_k \notin T'$, we have:

$$\pi_c(t_k|\theta, s) = 0, \quad \forall \theta, s \in \Theta : \theta \geq s$$

given the assumption that in PBE, there exists no defender type which fully covers the targets in the support set of the mixed attacking strategy.

According to Eq.(6), in PBE, the attacker's expected utilities of attacking t_l and t_k are:

$$U_a(\pi_d, s, t_l) = \sum_{\theta \in \Theta: \theta \geq s} \mu(\theta|s) [R_{t_l}^a - (R_{t_l}^a - P_{t_l}^a) \pi_c(t_l|\theta, s)]$$

$$U_a(\pi_d, s, t_k) = R_{t_k}^a.$$

Since $t_l \in T'$ and $t_k \notin T'$, $U_a(\pi_d, s, t_l) \geq U_a(\pi_d, s, t_k)$. According to Eq.(7) and $\sum_{\theta \in \Theta: \theta \geq s} \mu(\theta|s) = 1$, we have:

$$\sum_{\theta \in \Theta: \theta \geq s} \mu(\theta|s) (c_{t_l} - \pi_c(t_l|\theta, s)) \geq 0$$

Sum it up over $l \in \{1, \dots, |T'|\}$, we get:

$$\begin{aligned} & \sum_{\theta \in \Theta: \theta \geq s} \mu(\theta|s) \left(\sum_{l=1}^{|T'|} c_{t_l} - \sum_{l=1}^{|T'|} \pi_c(t_l|\theta, s) \right) \\ &= \sum_{\theta \in \Theta: \theta \geq s} \mu(\theta|s) \left(\sum_{l=1}^{|T'|} c_{t_l} - \theta \right) \geq 0 \end{aligned}$$

which contradicts to the following fact:

$$\begin{aligned} & \sum_{\theta \in \Theta: \theta \geq s} \mu(\theta|s) \left(\sum_{l=1}^{|T'|} c_{t_l} - \theta \right) \\ & \leq \sum_{\theta \in \Theta: \theta \geq s} \mu(\theta|s) (\theta_{min} - \theta) < 0. \end{aligned}$$

Thus, $|T'| \geq k$.

Now assume $|T'| > K$. Notice that $\{t_1, \dots, t_K\}$ are the set of best response targets against coverage \mathbf{c}' of type θ_{max} . Let l be an arbitrary index in $\{1, \dots, K\}$. We have:

$$R_{t_l}^a - (R_{t_l}^a - P_{t_l}^a) c'_{t_l} \geq R_{t_{|T'|}}^a \quad (8)$$

since $c'_{t_{|T'|}} = 0$. In *PBE*, both targets t_l and $t_{|T'|}$ are in the support set of $\pi_a(s)$. According to Eq.(8), we have:

$$\begin{aligned} & \sum_{\theta \in \Theta: \theta \geq s} \mu(\theta|s) [R_{t_l}^a - (R_{t_l}^a - P_{t_l}^a) \pi_c(t_l|\theta, s)] \\ &= \sum_{\theta \in \Theta: \theta \geq s} \mu(\theta|s) [R_{t_{|T'|}}^a - (R_{t_{|T'|}}^a - P_{t_{|T'|}}^a) \pi_c(t_{|T'|}|\theta, s)] \\ & \leq R_{t_{|T'|}}^a \leq R_{t_l}^a - (R_{t_l}^a - P_{t_l}^a) c'_{t_l}. \end{aligned}$$

Since $\sum_{\theta \in \Theta: \theta \geq s} \mu(\theta|s) = 1$, we get:

$$\sum_{\theta \in \Theta: \theta \geq s} \mu(\theta|s) (c'_{t_l} - \pi_c(t_l|\theta, s)) \leq 0.$$

Sum up over all $l \in \{1, \dots, K\}$, we have:

$$\begin{aligned} & \sum_{\theta \in \Theta: \theta \geq s} \mu(\theta|s) \left(\sum_{l=1}^K c'_{t_l} - \sum_{l=1}^K \pi_c(t_l|\theta, s) \right) \\ &= \sum_{\theta \in \Theta: \theta \geq s} \mu(\theta|s) (\theta_{max} - \sum_{l=1}^K \pi_c(t_l|\theta, s)) \leq 0. \end{aligned}$$

However, it contradicts the following fact when $s = \theta_{min}$:

$$\begin{aligned} & \sum_{\theta \in \Theta: \theta \geq s} \mu(\theta|s) (\theta_{max} - \sum_{l=1}^K \pi_c(t_l|\theta, s)) \\ & \geq \sum_{\theta \in \Theta: \theta \geq s} \mu(\theta|s) (\theta_{max} - \theta) > 0. \end{aligned}$$

Thus, $|T'| \leq K$. \square

7 Experimental Evaluation

We performed experiments to evaluate our algorithms, and to gather empirical data on how PBE compares generally to SSE. We use CPLEX for all optimizations on a 64-bit PC with 16 GB RAM and a quad-core 3.4 GHz processor. All values are averaged over 1000 instances expect for the runtime averaged over 100 instances. The game instances are generated as follows unless otherwise specified: each type θ is randomly drawn from $\{\lfloor 0.1|T| \rfloor, \lfloor 0.1|T| \rfloor + 1, \dots, \lfloor 0.4|T| \rfloor\}$. The probability distribution over Θ is randomly generated. The attacker's payoffs R_t^a and P_t^a are randomly drawn from the intervals $[1, 10]$ and $[-10, -1]$ respectively. The defender's payoffs are generated as follows: $R_t^d = \omega(-P_t^a) + (1 - \omega)\tilde{R}^d$ and $P_t^d = \omega(-R_t^a) + (1 - \omega)\tilde{P}^d$, where \tilde{R}^d and \tilde{P}^d are randomly drawn from same intervals as R_t^a and P_t^a respectively. The parameter ω controls correlation between the defender and attacker payoffs, such that when $\omega = 1$, the game is zero-sum, and there is no correlation when $\omega = 0$. The 95% confidence intervals are drawn in all figures which show that all the results are statistically significant.

Scalability & Solution Quality: We test the runtime of our approach on DSG instances with varying numbers of types $|\Theta| \in \{2, 4, 6, 8\}$ and $\omega \in \{0, 1\}$. The results are shown in Figs 2(a)–2(b), which shows that our approach can scale to realistic-sized instances with over 100 targets within minutes for all categories of games. We also test the solution quality of our approach on randomly generated games with $|T| \in \{40, 60, 80, 100\}$ and $|\Theta| \in \{4, 6\}$. $\text{Pr}(\text{PBE})$ denotes the proportion of instances where a PBE is computed. ϵ_{max} is the maximum value of ϵ among returned ϵ -PBEs. $|T'|$ represents the average number of generated support sets per instance. The results for $\omega = 0$ and 1 are given in Tables 2(c) and 2(d) respectively, from which we can see that the PBE is computed for over 99% of all tested instances. Even when the PBE is not returned, ϵ_{max} is significantly small compared with payoffs, showing that our approach can compute solutions with very good quality. We also note that $|T'|$ is much smaller than the number of targets, which empirically supports Theorem 4 to reduce the candidate support sets.

PBE vs. SSE & NE: We now compare the defender utility of PBE with SSE and NE. We test on random game instances with 20 targets, 8 types $\Theta = \{\theta_1 = 1, \dots, \theta_8 = 8\}$, and varying value of $\omega \in \{0.8, 0.85, 0.9, 0.95, 1.0\}$. In reality, different types may be of different importance for the defender. For example, a conservative defender may care more about the utility in the worst case with minimal resources. As such, we list the differences between defender utilities of each individual type in PBE and SSE in Figure 2(e). We also depict the expected defender utilities of PBE, SSE and NE in Figure 2(f). We observe that: i) with increasing ω , the defender benefits more in PBE compared to SSE, and the benefit of strategic secrecy is not limited to zero-sum games, which supports our formal analysis; ii) for the expected utility, the boundary of tradeoffs between strategic secrecy and commitment is within $[0.9, 0.95]$, close to zero-sum games; iii) PBE significantly outperforms NE regardless of the value of ω , supporting the motivation of strategic information revelation; and iv) the benefit of PBE shows a quadratic relationship with

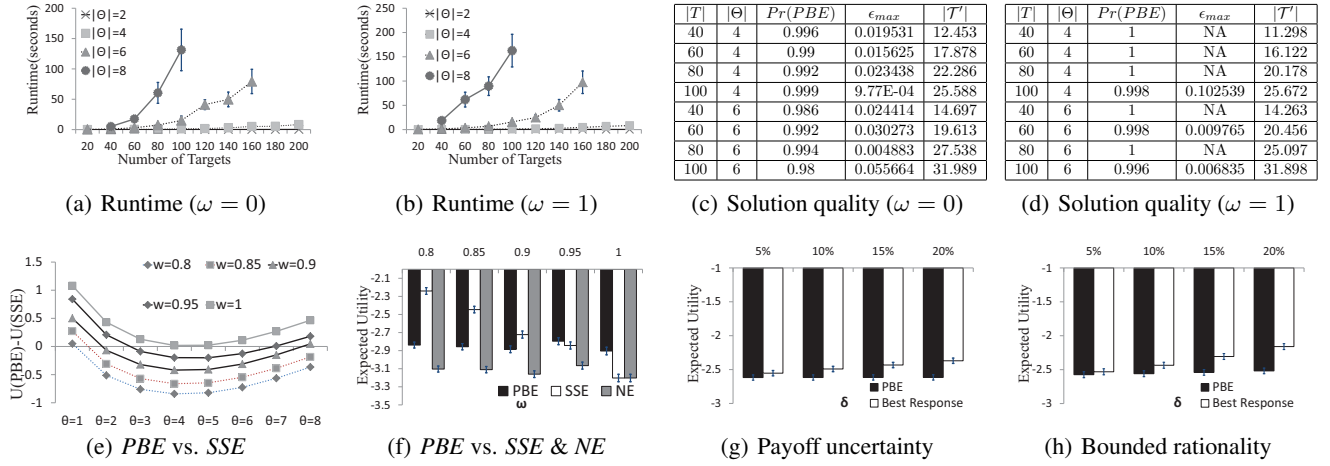


Figure 2: Experimental Evaluation.

the defender types. An intuitive explanation would be that the attacker is playing the best response against the defender of unknown type which can be treated as an “average” type $\bar{\theta}$. Therefore, the type θ closer to $\bar{\theta}$ benefits less from secrecy. Zero-sum games capture the nature of security issues where the attacker’s success indicates the failure of the defender, and the zero-sum approximation is widely adopted in game theoretic analysis in various security domains [Chen, 2007; Durkota *et al.*, 2015; Haskell *et al.*, 2014; Jain *et al.*, 2011; Major, 2002; Nguyen *et al.*, 2009a; 2009b; Wang *et al.*, 2016; Yin and An, 2016; Guo *et al.*, 2016; Yin *et al.*, 2016]. On the other hand, it is also emphasized that the zero-sum model is at best an approximation [Banks and Anderson, 2006]. One interpretation is that although both players are likely to agree on the importance of targets, the costs of conducting an attack or defending a target may be ignored by the opponent. Our results show that the boundary of PBE outperforming SSE is close to zero-sum ($w \approx 0.93$) which, to an extent, explain the coexistence of strategic secrecy and commitment in practice.

Robustness: We analyse the robustness of PBE solution against two major uncertainties on random zero-sum game instances with 30 targets and 6 types. First, the defender and the attacker may have different valuations of targets. Let \tilde{R}_t^* and \tilde{P}_t^* (* represents a or d) denote the payoffs estimated by the attacker. We denote by δ the degree of uncertainty such that $\tilde{R}_t^* \sim R_t^* \cdot [1 - \delta, 1 + \delta]$, $\tilde{P}_t^a = -\tilde{R}_t^d$ and $\tilde{P}_t^d = -\tilde{R}_t^a$. Let $\tilde{\pi}_a$ and $\tilde{\pi}_d$ be policies for attacker and defender correspondingly computed with their own estimations. We compare $U_d(\pi_d, \tilde{\pi}_a)$ with $U_d(\pi_d^*, \tilde{\pi}_a)$ where π_d^* is the best response against $\tilde{\pi}_a$ w.r.t. defender’s estimation of payoffs. Second, we consider the attacker’s bounded rationality, such that with a small probability δ , the attacker randomly chooses one target to attack. The metric for robustness analysis remains the same. The results are depicted in Figures 2(g) & 2(h), which show that the PBE solution is robust enough even with a high degree of uncertainty (20%), which makes it a practical alternative for the defender and also shows that our analysis and explanation of strategic secrecy based on PBE is reasonable.

8 Conclusions

We study a longstanding dilemma in security games: given the theoretical advantages of commitment, why is it that real-world security forces often use secrecy? By introducing the possibility that the defender has valuable private information, we show that there is a fundamental tradeoff between secrecy and commitment. We provide a generalization of security games to capture this, a novel scalable algorithm for computing PBE solutions for these games, and empirical results that demonstrate the effectiveness of our algorithms as well as providing a deeper understanding of the competing advantages of secrecy and commitment. Our theoretical and empirical results show that the boundary of such tradeoffs between secrecy and commitment is close to zero-sum, which is the case for most security domains. We conclude that both secrecy and commitment have a vital role in security policy.

Acknowledgements

This research is supported by the National Research Foundation, Prime Minister’s Office, Singapore under its IDM Futures Funding Initiative, NRF2015 NCR-NCR003-004, the National Science Foundation (NSF) under Grant No. IIS-1253950, and the Czech Science Foundation (grant no. 15-23235S).

References

- [An *et al.*, 2013] Bo An, Matthew Brown, Yevgeniy Vorobeychik, and Milind Tambe. Security games with surveillance cost and optimal timing of attack execution. In *AA-MAS*, pages 223–230, 2013.
- [Banks and Anderson, 2006] David L Banks and Steven Anderson. Combining game theory and risk analysis in counterterrorism: A smallpox example. In *Statistical Methods in Counterterrorism*, pages 9–22. 2006.
- [Bertsimas and Tsitsiklis, 1997] Dimitris Bertsimas and John N Tsitsiklis. *Introduction to Linear Optimization*, volume 6. Athena Scientific Belmont, MA, 1997.

- [Brown *et al.*, 2005] Gerald Brown, Matthew Carlyle, Douglas Diehl, Jeffrey Kline, and Kevin Wood. A two-sided optimization for theater ballistic missile defense. *Operations Research*, 53(5):745–763, 2005.
- [Catchnews, 2016] Catchnews. Pathankot attack: US confirms Pakistan’s involvement in terror strike, 2016.
- [Chen, 2007] Zesheng Chen. *Modeling and defending against internet worm attacks*. PhD thesis, Georgia Institute of Technology, 2007.
- [Durkota *et al.*, 2015] Karel Durkota, Viliam Lisý, Branislav Bosanský, and Christopher Kiekintveld. Approximate solutions for attack graph games with imperfect information. In *GameSec*, pages 228–249, 2015.
- [Fang *et al.*, 2016] Fei Fang, Thanh H Nguyen, Rob Pickles, Wai Y Lam, Gopalasamy R Clements, Bo An, Amandeep Singh, Milind Tambe, and Andrew Lemieux. Deploying PAWS: Field optimization of the protection assistant for wildlife security. In *IAAI*, pages 3966–3973, 2016.
- [Farrell and Rabin, 1996] Joseph Farrell and Matthew Rabin. Cheap talk. *The Journal of Economic Perspectives*, 10(3):103–118, 1996.
- [Gan *et al.*, 2015] Jiarui Gan, Bo An, and Yevgeniy Vorobeychik. Security games with protection externalities. In *AAAI*, pages 914–920, 2015.
- [Gul, 2011] Imtiaz Gul. PNS Mehran attack: Vulnerable, embarrassed and targeted, 2011.
- [Guo *et al.*, 2016] Qingyu Guo, Bo An, Yair Zick, and Chunyan Miao. Optimal interdiction of illegal network flow. In *IJCAI*, pages 2507–2513, 2016.
- [Haskell *et al.*, 2014] William B. Haskell, Debarun Kar, Fei Fang, Milind Tambe, Sam Cheung, and Elizabeth Denicola. Robust protection of fisheries with compass. In *AAAI*, pages 2978–2983, 2014.
- [Hendricks and McAfee, 2006] Kenneth Hendricks and R Preston McAfee. Feints. *Journal of Economics & Management Strategy*, 15(2):431–456, 2006.
- [Jain *et al.*, 2011] Manish Jain, Dmytro Korzhyk, Ondrej Vanek, Vincent Conitzer, Michal Pechoucek, and Milind Tambe. A double oracle algorithm for zero-sum security games on graphs. In *AAMAS*, pages 327–334, 2011.
- [Kiekintveld *et al.*, 2009] Christopher Kiekintveld, Manish Jain, Jason Tsai, James Pita, Fernando Ordóñez, and Milind Tambe. Computing optimal randomized resource allocations for massive security games. In *AAMAS*, pages 689–696, 2009.
- [Leitmann, 1978] George Leitmann. On generalized Stackelberg strategies. *Journal of Optimization Theory and Applications*, 26(4):637–643, 1978.
- [Major, 2002] John A Major. Advanced techniques for modeling terrorism risk. *The Journal of Risk Finance*, 4(1):15–24, 2002.
- [Nguyen *et al.*, 2009a] Kien C Nguyen, Tansu Alpcan, and Tamer Basar. Security games with incomplete information. In *Proceedings of the 2009 IEEE International Conference on Communications*, pages 1–6, 2009.
- [Nguyen *et al.*, 2009b] Kien C Nguyen, Tansu Alpcan, and Tamer Basar. Stochastic games for security in networks with interdependent nodes. In *GameNets*, pages 697–703, 2009.
- [Oliveros, 2005] S Oliveros. Equilibrium bluffs: A model of rational feints. Technical report, Working paper, University of Wisconsin-Madison, Department of Economics, 2005.
- [Pita *et al.*, 2010] James Pita, Manish Jain, Milind Tambe, Fernando Ordóñez, and Sarit Kraus. Robust solutions to Stackelberg games: Addressing bounded rationality and limited observations in human cognition. *Artificial Intelligence*, 174(15):1142–1171, 2010.
- [Rabinovich *et al.*, 2015] Zinovi Rabinovich, Albert Xin Jiang, Manish Jain, and Haifeng Xu. Information disclosure as a means to security. In *AAMAS*, pages 645–653, 2015.
- [Rubinstein, 1985] Ariel Rubinstein. A bargaining model with incomplete information about time preferences. *Econometrica*, 53(5):1151–1172, 1985.
- [Shieh *et al.*, 2012] Eric Anyung Shieh, Bo An, Rong Yang, Milind Tambe, Craig Baldwin, Joseph DiRenzo, Ben Maule, and Garrett Meyer. PROTECT: An application of computational game theory for the security of the ports of the United States. In *AAAI*, pages 2173–2179, 2012.
- [Spence, 1973] Michael Spence. Job market signaling. *The Quarterly Journal of Economics*, pages 355–374, 1973.
- [Wang *et al.*, 2016] Zhen Wang, Yue Yin, and Bo An. Computing optimal monitoring strategy for detecting terrorist plots. In *AAAI*, pages 637–643, 2016.
- [Xu *et al.*, 2015] Haifeng Xu, Zinovi Rabinovich, Shaddin Dughmi, and Milind Tambe. Exploring information asymmetry in two-stage security games. In *AAAI*, pages 1057–1063, 2015.
- [Yin and An, 2016] Yue Yin and Bo An. Efficient resource allocation for protecting coral reef ecosystems. In *IJCAI*, pages 531–537, 2016.
- [Yin *et al.*, 2014] Yue Yin, Bo An, and Manish Jain. Game-theoretic resource allocation for protecting large public events. In *AAAI*, pages 826–834, 2014.
- [Yin *et al.*, 2016] Yue Yin, Yevgeniy Vorobeychik, Bo An, and Noam Hazon. Optimally protecting elections. In *IJCAI*, pages 538–545, 2016.
- [Zhao *et al.*, 2016] Mengchen Zhao, Bo An, and Christopher Kiekintveld. Optimizing personalized email filtering thresholds to mitigate sequential spear phishing attacks. In *AAAI*, pages 658–665, 2016.
- [Zhuang and Bier, 2011] Jun Zhuang and Vicki M Bier. Secrecy and deception at equilibrium, with applications to anti-terrorism resource allocation. *Defence and Peace Economics*, 22(1):43–61, 2011.