The AAAI-17 Workshop on
Computer Poker and Imperfect Information Games
WS-17-06

Strategic Information Revelation and Commitment in Security Games

Qingyu Guo,[1] Bo An,[2] Branislav Bosansky,[3] Christopher Kiekintveld[4]
[1]Joint NTU-UBC Research Centre of Excellence in Active Living for the Elderly, NTU, Singapore
[2]School of Computer Science and Engineering, Nanyang Technological University, Singapore
[3]Department of Computer Science, Faculty of Electrical Engineering, Czech Technical University in Prague, Czech Republic
[4]Computer Science Department, University of Texas at El Paso, USA
[1,2]{qguo005,boan}@ntu.edu.sg,[3]bosansky@agents.fel.cvut.cz,[4]cdkiekintveld@utep.edu

Abstract

The *Strong Stackelberg Equilibrium (SSE)* has drawn extensive attention recently in several security domains, which optimizes the defender's random allocation of limited security resources. 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 basic *MILP* formulations with mixed defender strategy and compact representation, 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 solve *PBE* for realistic-sized problems with good enough and robust solution quality.

Introduction

Consider the problem a police department or homeland security agency faces when deciding how to use limited resources to protect critical infrastructure, patrol transportation systems, or secure a large event such as the Super Bowl. The security force must assess the threat posed by different attack scenarios, and allocate the available security resources to maximize the level of protection provided. *Stackelberg security games* have gained traction as a way to intelligently allocate security resources while keeping the schedule unpredictable. They have been used for protecting public infrastructures (Kiekintveld et al. 2009; Shieh et al. 2012) and wildlife (Fang et al. 2016).

However, there is a dilemma that has not been resolved when we compare the recommendations of Stackelberg security games with the actual practice of security agencies. A key assumption in the Stackelberg model is that the security force will *credibly commit* to following a particular randomized strategy, which is publicly observed by the attacker who chooses a best response to the observed strategy. This is often motivated by the ability of an attacker to use surveillance to learn about the defender strategy before deciding on an attack strategy, though it may not be realistic for the attacker to learn the strategy exactly through such observations (An et al. 2013; Gul 2011; Pita et al. 2010). However, there is a stronger claim: theoretically, the defender will always do at least as well by committing to a mixed strategy publicly as they will do by playing a "secret" strategy. Therefore, according to the theoretical models the defender should not resist surveillance efforts, but instead should actively announce the strategy to make the commitment as credible as possible.

This is at odds with both intuition as well as the actual practice of security agencies who frequently use a combination of a highly visible security presence (e.g., conspicuous uniformed officers and marked vehicles) alongside "plainclothes" security forces who are disguised to fit in with civilians and be difficult to observe. How can we resolve this dilemma, when the current theory suggests that such secrecy is suboptimal? What role is there for hiding security resources, and how does this compare with the advantages of public commitment?

We propose that one resolution to this dilemma lies in considering a factor that is absent from current work on security games: *committing to a strategy may also reveal the defender's private information about the number of resources available.* This frames a tradeoff for the defender: is it more beneficial to keep secret the private information, or to commit to a strategy? We study this question from both theoretical and empirical perspectives. We propose a novel model called a *Disguised-resource Security Game (DSG)* where there are multiple Bayesian types of defenders with different numbers of resources, and the defender can strategically choose whether to disguise resources or not, modeling the real-world decision of whether to make the security forces uniformed or in plainclothes.

We make several key contributions. First, we propose a *Disguised-resource Security Game (DSG)* model to 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. Second, we compare the value of

strategically disguising resources with the value of public commitment, and formally show that they have different advantages depending on the payoff structure. Third, we introduce algorithms to solve for the solution concept based on *Perfect Bayesian Equilibrium (PBE)* (Spence 1973; Zhuang and Bier 2011) of a *DSG*. This is computationally extremely challenging due to the exponential number of pure strategies, so we introduce a basic *Mixed-Integer Linear Programming (MILP)* with an exponential number of variables and constraints, a *MILP* of directly applying compact representation, and a novel approach based on support set enumeration. Based on analysis of the relationship between the support set of *PBE* and payoff structure we can limit the number of support sets and further improve scalability. Fourth, we introduce 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 tradeoffs between secrecy and public commitment for realistic problems. We conclude that the boundary of such tradeoffs is close to zero-sum games, and both strategic secrecy and commitment play a vital role in practice, given the approximate zero-sum nature of homeland security (Banks and Anderson 2006; Durkota et al. 2015; Nguyen, Alpcan, and Basar 2009).

Related Work

Previous work on secrecy and deception in security games has failed to address the key dilemma of strategic secrecy 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. Hespanha *et al.* (2000) study how a defender can manipulate the information available to an attacker, but model assumes Nash equilibrium with no private information held by the defender. 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 since only the defender with enough resources can send the specific signal, which differs from the cheap talk game (Farrell and Rabin 1996) where messages are costless and unverifiable, and any sender can send any message, so that the receiver may ignore them at all.

Recently, there are some literatures studying the information disclosure in security games and the optimal signaling scheme 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 extensive surveillance, which is impractical and not reasonable in situations where the attacker has only limited observation and cannot observe the true defender type or the actual resource allocation, such as 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. They analyze the defender's preference over truthful disclosure, secrecy and deception depending on the costs and private information of both players. 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 general algorithms for realistic problems.

Disguised-Resource Security Games (*DSG*)

We now illustrate our Disguised-Resource Security Games (DSG) model in details. A *DSG* has the same basic structure as 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 $p : \Theta \to [0, 1]$ is known to both players.

We model private information for the defender by allowing the defender to publicly reveal only a subset of her available resources using a *signal*. 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 deceptive signals that claim a greater number of resources than are actually available, so a defender of type θ can only send a signal $s \leq \theta$. There are four payoffs associated with each target, $\langle R_t^d, P_t^a, P_t^d, R_t^a \rangle$: if a resource is allocated to attacked target, 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. Assume $R_t^d > P_t^d$ and $R_t^a > P_t^a$. The overall interaction between the defender and attacker proceeds as follows: The defender moves first, samples a signal s from the mixed signaling strategy and an allocation according to the randomized allocation strategy conditioned on s, and publishes the signal s; The attacker, observes the signal s, infers the posterior distribution of the defender type, and decides the target to attack. Corresponding payoffs are received. The *DSG* is an extensive-form game as illustrated by the example in Section . The attacker cannot distinguish the set of decision nodes where multiple defender types send the same signal s, which is called an *information set* for the attacker, 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 where c_t and a_t represent the probability of target t being covered by a security resource and the proba-

bility of attacking target t respectively. Let $\mathbf{o} = \langle o_s \rangle$ denote the mixed signaling strategy such that o_s represents the probability of sending signal s. Let $\Delta_c^\theta = \{\mathbf{c} \in [0,1]^{|T|} : \sum_{t \in T} c_t = \theta\}$ 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 $\Delta_o^\theta = \{\mathbf{o} \in [0,1]^{|\Theta|} : \sum_{s \in \Theta} o_s = 1, o_s = 0 \quad \forall s > \theta\}$ the set of mixed signaling strategies available for defender type θ and $\Delta_o = \bigcup_{\theta \in \Theta} \Delta_o^\theta$ the set of all mixed signaling strategies. Let $\Delta_a = \{\mathbf{a} \in [0,1]^{|T|} : \sum_{t \in T} a_t = 1\}$ represent the set of all mixed attacking strategies. Let $\pi_d = \langle \pi_c, \pi_o \rangle$ denote the defender's policy where $\pi_c : \Theta \times \Theta \to \Delta_c$ 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 $\pi_o : \Theta \to \Delta_o$ is the signaling policy with $\pi_o(\theta) \in \Delta_o^\theta$ representing the mixed signaling strategy for defender of 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 of type θ. Let $\pi_a : \Theta \to \Delta_a$ denote the attacker's 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 $\Delta_\Theta = \{\langle \delta_\theta \rangle : \sum_{\theta \in \Theta} \delta_\theta = 1\}$ be the set of all possible probability distributions over Θ. We denote by $\mu : \Theta \to \Delta_\Theta$ 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's type being θ after signal s is received. 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:

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

Otherwise, $I(s)$ is off the equilibrium path, and we adopt the *optimistic conjecture* (Rubinstein 1985), such that when the defender acts off the equilibrium strategy, the attacker believes the defender is the weakest type, against which the attacker would gain the most. Intuitively the attacker always prefers to play against a defender type with less resources. Theorem 1 formally illustrates this by showing a procedure to get an *NE* profile $\langle \mathbf{c}', \mathbf{a}' \rangle$ between defender type θ' and the attacker, from the *NE* profile $\langle \mathbf{c}, \mathbf{a} \rangle$ between defender type $\theta > \theta'$ and the attacker, where $\mathbf{c}' \prec \mathbf{c}$. Thus, at information set $I(s)$ which is off equilibrium path, we have: $\mu(s|s) = 1$ and $\mu(\theta|s) = 0$ for all $\theta > s$.

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 $P_a(\mathbf{c}', \mathbf{a}') \geq P_a(\mathbf{c}, \mathbf{a})$*[1].

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:

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

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:

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

$$U_a(\pi_c, \pi_o, \pi_a(s)) = \sum_{\theta:\theta \geq s} \mu(\theta|s) P_a(\pi_c(\theta, s), \pi_a(s)).$$

Equilibrium Concepts: Analogous to the equilibrium of extensive-form game with first-mover hidden actions defined by Zhuang and Bier (2011), the solution concept we use for *DSG* is based on *perfect Bayesian equilibrium (PBE)*, which is the profile $\langle \pi_d^*, \pi_a^* \rangle$ satisfying:

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

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

Due to the strict requirement in *PBE* that both players will not play a suboptimal response strategy, the computation of *PBE* is extremely challenging. Therefore, 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 $\langle \pi_d^*, \pi_a^* \rangle$ satisfying: i) $\forall \theta \in \Theta$ and $\forall \langle \pi_c(\theta), \pi_o(\theta) \rangle$, $U_d(\pi_c^*(\theta), \pi_o^*(\theta), \pi_a^*) \geq U_d(\pi_c(\theta), \pi_o(\theta), \pi_a^*) - \epsilon$, and ii) $\forall s \in \Theta$ and $\forall \pi_a(s)$, $U_a(\pi_c^*, \pi_o^*, \pi_a^*(s)) \geq U_a(\pi_c^*, \pi_o^*, \pi_a(s)) - \epsilon$.

The *Strong Stackelberg equilibrium* (SSE) (Leitmann 1978) between the defender and attacker is a pair of strategies $\langle \mathbf{c}, f(\mathbf{c}) \rangle$ satisfying: i) $P_d(\mathbf{c}, f(\mathbf{c})) \geq P_d(\mathbf{c}', f(\mathbf{c}'))$ for all defender coverage $\mathbf{c}'$, ii) $P_a(\mathbf{c}, f(\mathbf{c})) \geq P_a(\mathbf{c}, \mathbf{a})$ for all attacking strategies $\mathbf{a}$; 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 $\mathbf{a}$.

Motivating Example

Suppose a small police station (the defender) has two districts to protect, A and B. The police station has either one or two patrol units, depending on the day. We call these case the *weak* and *strong* types, and assume that they are equally likely for the example. An attacker will choose one of the two districts to target, represented by a mixed attacking strategy $\mathbf{a} = \langle a_A, a_B \rangle$. The police strategy can be compactly represented by a coverage vector $\mathbf{c} = \langle c_A, c_B \rangle$. Consider that the strong type (with 2 resources) can disguise one resource, or choose to commit to an *SSE* using both resources. The game is shown in Figure 1, where $R_A^a = -P_A^d = 8$

[1] Due to the length limitation, we briefly sketch the idea of proofs in the paper and omit the details for ease of reading.

and $R_B^a = -P_B^d = 6$, while $R_A^d = -P_A^a = 4$ and $R_B^d = -P_B^a = 2$. If the defender chooses to reveal the type and play an *SSE* the weak type will play the coverage $\langle 0.5, 0.5\rangle$ equalizing the expected payoff for the attacker between the two targets, resulting in an expected utility of -2. The strong type will play $\langle 1, 1\rangle$ and the attacker will attack B, so the strong type receives utility 2.

If the defender is allowed to disguise a resource, we have a different game. In the *PBE* of this game, the strong type will hide one resource with 100% probability. When the attacker observes only 1 patrol unit, he is playing against both types with equal probability. In *PBE* the weak type plays coverage $\langle 0.4, 0.6\rangle$, and attacker receives an expected payoff of -0.4 for attacking A or B, while the attacker plays a mixed strategy $\langle 0.4, 0.6\rangle$. Therefore, the weak type still receives an expected utility of -2, while the strong type gets 2.8, resulting in an overall gain for the defender.

However, disguising resources is not always beneficial. Suppose the payoffs are changed to $P_A^d = -4$ and $R_B^d = 2$, so that the game is no longer zero-sum. The defender's behavior in both *PBE* and *SSE* is the same as before. In *PBE*, the attacker will play $\langle 0.5, 0.5\rangle$ when observing 1 defender resource, and the weak type and strong type receive -1 and 3 respectively. However, in *SSE* the attacker will always attack target A and the two types receive 0 and 4 respectively, higher than the expected utilities in *PBE*.

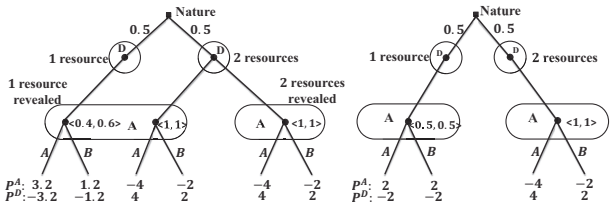

Figure 1: Motivating Example. (left: *PBE*, right: *SSE*)

PBE versus *SSE*

We now compare more formally *PBE* with *SSE* and show that they are beneficial in different situations, and are particularly sensitive to the correlation between defender and attacker payoffs. For zero-sum *DSG*s 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 *DSG*s, the defender cannot benefit from public commitment and *SSE* reduces to *NE*, while in *PBE* the attacker cannot distinguish the defender type, and the mixed attack strategy may not be a best response to the individual coverage of each type, so the defender can take the advantage.

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

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 possible that there exists a set of targets $\{t_1, .., t_k\}$, such that for each type θ, there is an *NE* between defender of type θ and the attacker where the defender covers $\{t_1, .., t_k\}$ while leaving the remaining targets uncovered. In this case, we can show the existence of a *PBE* where each defender type is playing that *NE* coverage regardless of signals, and the attacker is playing the unbiased mixed attack strategy (Definition 1) with support set $\{t_1, .., t_k\}$ in all information sets, which is a best response against the coverage of each individual type. Therefore, the defender's expected utility of any type in such a *PBE* is equal to her expected utility in *NE*, which is less than or equal to that in *SSE*. Since Theorem 3 has no requirement on defender payoff or 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 in the *PBE* in Theorem 3.

Theorem 3. *Suppose* $R_{min}^a > P_{max}^a$ *where* $R_{min}^a = \min_{t\in T} R_t^a$ *and* $P_{max}^a = \max_{t\in T} P_t^a$*. Let the targets be listed by* R_t^a *with descending order:* $T = \{t_1, .., t_{|T|}\}$*. If there exists* k *such that:* $\sum_{l=1}^{k} \frac{R_{t_l}^a - R_{t_k}^a}{R_{t_l}^a - P_{t_l}^a} \leq \theta \leq \sum_{l=1}^{k} \frac{R_{t_l}^a - R_{t_{k+1}}^a}{R_{t_l}^a - P_{t_l}^a}$ *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.*

Computing PBE Solutions

We now introduce computation methods for computing *PBE*. We first try a *MILP* based on mixed defender strategy representation that 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 can 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 non-trivial. Although we can use linear constraints to represent such criteria 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 *PBE*s where the attacker policy is *unbiased* (Definition 1), which makes the defender's best response criteria much easier to represent. We then propose a concise and scalable formulation to compute such *PBE*s 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 for *PBE* and ϵ-*PBE*.

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 nice property of $\mathbf{a}^{T'}$ is that any deviation of defender's marginal coverage between two targets in T' cannot change the defender's expected utility. To show this, according to (1), the defender's expected payoff against $\mathbf{a}^{T'}$ is:

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

which only depends on T' and the total marginal coverage on targets in T'. Therefore, 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 of type θ against $\mathbf{a}^{T'}$ is denoted as:

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

Support Set Enumeration for *PBE*: 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 strategies, which is of size $|\mathcal{T}|^{|\Theta|}$ as there are $|\mathcal{T}|$ possible unbiased attacking strategies at each information set. 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 a set of linear constraints, since the defender's best response criteria, can be easily represented as:

$$\sum\nolimits_{t\in T_s'} \pi_c(t|\theta, s) = \min\{\theta, |T_s'|\} \quad \forall \theta, s \in \Theta : \theta \geq s, \tag{2}$$

However, the size of $\mathcal{T}$ is exponential ($2^{|T|}$), which makes it impossible to generate all possible unbiased attacker strategies. 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 if the attacker penalty is constant value, we can prove 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). Even for the general payoff structure, that property holds for most cases as shown in experimental evaluation. 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.) For this aim, instead of brute force search, we provide an *MILP* with no objective function for arbitrary *PBE* as follows:

$$\sum\nolimits_{s\in\Theta: s\leq\theta} S_{\theta s} = 1 \quad \forall\theta \tag{3a}$$

$$\sum\nolimits_{t\in T} \tilde{C}_{\theta st} = \theta S_{\theta s} \quad \forall \theta > s \tag{3b}$$

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

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

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

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

$$\delta\chi_s \leq \sum\nolimits_{\theta\in\Theta:\theta\geq s} p_\theta S_{\theta s} \leq \chi_s \quad \forall s \tag{3f}$$

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

$$\sum\nolimits_{T'\in\mathcal{T}'} \phi_{sT'} = 1 \quad \forall s \tag{3h}$$

$$x_{st} + \sum\nolimits_{\theta\in\Theta:s\leq\theta} p_\theta \tilde{C}_{\theta st} P_t^a + \sum\nolimits_{\theta\in\Theta:s\leq\theta} p_\theta (S_{\theta s} - \tilde{C}_{\theta st}) R_t^a + p_s(1 - \chi_s) R_t^a = \upsilon_s^a \quad \forall s, t \tag{3i}$$

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

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

$$\tilde{C}_{\theta st} \geq S_{\theta s} + \sum_{T': t\in T', |T'|<\theta} \phi_{sT'} - 1 \quad \forall \theta > s, t \tag{3l}$$

$$\tilde{C}_{sst} \geq S_{ss} - \chi_s + \sum_{T': t\in T', |T'|<s} \phi_{sT'} \quad \forall s, t \tag{3m}$$

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

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

$$y_{\theta s} + \sum\nolimits_{T'\in\mathcal{T}'} P_{\theta T'}^d \phi_{sT'} = \upsilon_\theta^d \quad \forall \theta, s \tag{3n}$$

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

In formulation (3), S is the decision variable representing the signaling policy 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 $I(s)$ is on the equilibrium path, otherwise $\chi_s = 0$; 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 υ_θ^d denotes $U_d(\pi_o(\theta), \pi_c(\theta), \pi_a)$; $y_{\theta s}$ is the slack variable which takes zero when $P_d(\pi_c(\theta, s), \pi_a(s))$ equals υ_θ^d; Binary variable $\phi_{sT'} = 1$ if T' is the support set of $\pi_a(s)$, otherwise $\phi_{sT'} = 0$. (3i) and (3j) correspond to the attacker's best response criteria, taking into account 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)–(3m) ensure 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 $t \notin T'$ if $\theta \leq |T'|$ and $\pi_c(t|\theta, s) = 1$ for $t \in T'$ otherwise. Finally, (3n)–(3o) 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 $s' \leq \theta$.

Support Set Enumeration for ϵ-*PBE*: The *MILP* (3) returns a *PBE* 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 highest expected utility 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. Therefore, to get the best approximation of *MILP* (4), we conduct a binary search on ϵ from initial interval $[0, M]$, where M is a large constant making *MILP* (4) feasible, and obtain the smallest possible ϵ with which *MILP* (4) returns a solution.

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

Generating Support Sets: We now discuss how to generate $\mathcal{T}'$ for the support set enumeration approach. Suppose the attacker penalty is a constant value P and there exists no pair of targets with the same reward for the attacker. We can list the targets by R_t^a in descending order: $T = \{t_1, .., t_{|T|}\}$. Our next Lemma shows that the support set T' of any mixed attacking strategy in *PBE* must contain the first $|T'|$ targets in T. The intuition is that the defender will always cover the targets in T' with the highest priority for best response, and if a target $t \notin T'$ has higher reward than $t' \in T'$, the attacker will gain more by attacking t.

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

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 $\{1, .., k\}$ and $\{1, .., K\}$ correspond to such sets of targets that can be covered by θ_{min} and θ_{max} 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: $|\mathcal{T}'| = |K - k + 1| \ll |T|$.

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

As for *DSG* instances with general payoffs, we can still compute k and K, by replacing the constant penalty P in the inequality of Theorem 4 with individual value P_t^a, as well as the set of candidate support sets $\mathcal{T}'$ which, although it is not guaranteed to include all possible support sets of mixed attack strategy in *PBE*, is enough for the support set enumeration approach to obtain good solutions.

Experimental Evaluation

We performed experiments to test the scalability and quality of our algorithms, and to gather empirical data on how PBE compares generally to SSE. We use CPLEX (version 12.6) 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 scalability analysis where runtime is averaged over 100 instances. The game instances are generated as follows unless otherwise specified: each type θ is randomly generated from $\{\lfloor 0.1|T| \rfloor, \lfloor 0.1|T| \rfloor + 1, .., \lfloor 0.4|T| \rfloor\}$. 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 the standard error is relatively small compared with mean values. Thus, all the results are statistically significant.

Scalability: We test the runtime of our support set enumeration approach on *DSG* game 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). Our approach can scale to realistic-sized instances with 200 targets for $|\Theta| \in \{2, 4\}$, 160 targets for $|\Theta| = 6$ and 100 targets for $|\Theta| = 8$ within minutes, for all categories of games.

Solution Quality: We test the solution quality of support set enumeration on randomly generated games with $|T| \in \{40, 60, 80, 100\}$ and $|\Theta| \in \{4, 6\}$. Pr(*PBE*) denotes the proportion of instances where a *PBE* is computed. ϵ_{max} is the maximum value of ϵ among all ϵ-*PBE*s that are returned by the algorithm. $|\mathcal{T}'|$ 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 a *PBE* is computed for over 99% of all tested instances, and 100% for some trials of zero-sum instances. Even when the *PBE* is not returned, ϵ_{max} is very small compared with the payoff magnitude 10, showing that our approach can compute solutions with very good quality. We also note that $|\mathcal{T}'|$ is much smaller than the number of targets, which provides empirical support for Theorem 4 to dramatically 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 listed by number of resources in ascending order $\Theta = \{\theta_1 = 1, .., \theta_8 = 8\}$, 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, where the number of resources is minimal. For this aim, 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) with varying value of

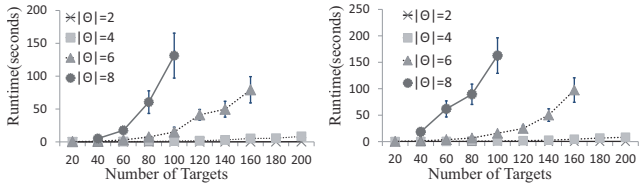

(a) Runtime ($\omega = 0$)

(b) Runtime ($\omega = 1$)

$\|T\|$	$\|\Theta\|$	$Pr(PBE)$	ϵ_{max}	$\|\mathcal{T}'\|$
40	4	0.996	0.019531	12.453
60	4	0.99	0.015625	17.878
80	4	0.992	0.023438	22.286
100	4	0.999	9.77E-04	25.588
40	6	0.986	0.024414	14.697
60	6	0.992	0.030273	19.613
80	6	0.994	0.004883	27.538
100	6	0.98	0.055664	31.989

(c) Solution quality ($\omega = 0$)

$\|T\|$	$\|\Theta\|$	$Pr(PBE)$	ϵ_{max}	$\|\mathcal{T}'\|$
40	4	1	NA	11.298
60	4	1	NA	16.122
80	4	1	NA	20.178
100	4	0.998	0.102539	25.672
40	6	1	NA	14.263
60	6	0.998	0.009765	20.456
80	6	1	NA	25.097
100	6	0.996	0.006835	31.898

(d) Solution quality ($\omega = 1$)

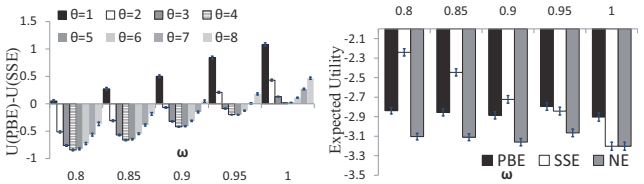

(e) *PBE* vs. *SSE*

(f) *PBE* vs. *SSE* & *NE*

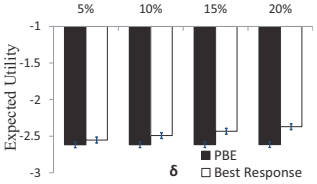

(g) Payoff uncertainty

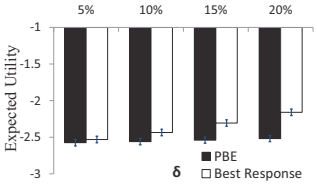

(h) Bounded rationality

Figure 2: Experimental Evaluation.

ω. We observe that: i) with increasing ω, the defender benefits more in *PBE* compared to *SSE*, supporting our formal analysis. For the worst case of type θ_1, even with $\omega = 0.8$, the *PBE* solution outperforms *SSE*, showing that the benefit of strategic secrecy is not limited to zero-sum games; ii) From the perspective of expected defender utility, the boundary of tradeoffs between strategic secrecy and commitment is within $[0.9, 0.95]$, which is 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 the defender types, such that types θ_1 and θ_8 benefit the most from *PBE*, while types θ_4 and θ_5 benefit the least from strategic secrecy. Here we provide an intuitive explanation while leaving the formal analysis to the future work: the defender of an unknown type can be treated as an "average" type $\bar{\theta}$ such that the attacker is playing the best response against $\bar{\theta}$; Therefore, the type θ with larger gap from $\bar{\theta}$ benefits more from secrecy.

Robustness: In practice, the uncertainty of payoff structures always exists and the attacker is not perfectly rational. Therefore, 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 assume that they have the same valuations of targets while not, which causes the uncertainty of payoffs. Let $\tilde{R}_t^*$ and $\tilde{P}_t^*$ ($*$ represents a or d) denote the payoffs estimated by the attacker which still satisfy the zero-sum property. In our experiments, 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 $\tilde{\pi}_a$ be the attacker's policy according to his own estimation of payoffs, and π_d be the defender's policy in her computed *PBE* $\langle \pi_d, \pi_a \rangle$. We compare $U_d(\pi_d, \tilde{\pi}_a)$ with $U_d(\pi_d^*, \tilde{\pi}_a)$ where $U_d(\pi_d, \tilde{\pi}_a)$ denotes the expected defender utility and π_d^* is the actual best response against $\tilde{\pi}_a$ with respect to defender's estimation of payoffs. The result is depicted in Figure 2(g), from which we can observe that: i) with increasing degree of uncertainty, the regret of not playing π_d^* is increasing, and ii) the expected utility $U_d(\pi_d, \tilde{\pi}_a)$ almost remains the same regardless of δ; The reason is that in most instances π_a and $\tilde{\pi}_a$ have the same support sets so that $\tilde{\pi}_a$ is also a best response strategy against π_d. Second, we take into consideration the attacker's bounded rationality, such that with a small probability δ (irrational degree), the attacker randomly chooses one target to attack. The metric for robustness analysis is the same as the first case of payoff uncertainty, and the result is depicted in Figure 2(h), which shows that: i) the regret of not playing the best response strategy is increasing with larger irrational degree δ, and ii) the expected utility $U_d(\pi_d, \tilde{\pi}_a)$ increases with increasing value of δ (the improvement is too subtle to show in the figure directly) since the attacker may choose the targets which are not the best response ones with higher probability. The experimental evaluation shows 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.

Discussion: 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 existing works of game theoretic analysis in various security domains (Chen 2007; Durkota et al. 2015; Nguyen, Alpcan, and Basar 2009). 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, there are some costs of conducting an attack or defending a target which are ignored by the opponent. The ratio between the magnitudes of such cost and the reward of a successful attack decides how close to the zero-sum the game is. Such ratio may vary among different security domains. For example, in physical security, the reward of a successful terrorist attack (9/11 for example) is usually several magnitudes larger than the cost, and the game is more close to zero-sum; While in green security such as illegal fishing and poaching, the ratio is larger and the game is less close to zero-sum. 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.

Conclusion

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 to play in optimal security policy.

Acknowledgements

This research is supported by NRF2015NCR-NCR003-004, the National Research Foundation, Prime Minister's Office, Singapore under its IDM Futures Funding Initiative, the NSF under Grant No. IIS-1253950 and the Czech Science Foundation (grant no. 15-23235S).

References

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, (AAMAS)*, 223–230.

Banks, D. L., and Anderson, S. 2006. Combining game theory and risk analysis in counterterrorism: A smallpox example. In *Statistical methods in counterterrorism*. 9–22.

Bertsimas, D., and Tsitsiklis, J. N. 1997. *Introduction to Linear Optimization*, volume 6. Athena Scientific Belmont, MA.

Brown, G.; Carlyle, M.; Diehl, D.; Kline, J.; and Wood, K. 2005. A two-sided optimization for theater ballistic missile defense. *Operations Research* 53(5):745–763.

Chen, Z. 2007. *Modeling and defending against internet worm attacks*. Ph.D. Dissertation, Georgia Institute of Technology.

Durkota, K.; Lisý, V.; Bosanský, B.; and Kiekintveld, C. 2015. Approximate solutions for attack graph games with imperfect information. In *Proceedings of the 6th International Conference on Decision and Game Theory for Security (GameSec)*, 228–249.

Fang, F.; Nguyen, T. H.; Pickles, R.; Lam, W. Y.; Clements, G. R.; An, B.; Singh, A.; Tambe, M.; and Lemieux, A. 2016. Deploying PAWS: Field optimization of the protection assistant for wildlife security. In *Proceedings of the 28th Innovative Applications of Artificial Intelligence Conference (IAAI)*, 3966–3973.

Farrell, J., and Rabin, M. 1996. Cheap talk. *The Journal of Economic Perspectives* 10(3):103–118.

Gul, I. 2011. PNS Mehran attack: Vulnerable, embarrassed and targeted. http://tribune.com.pk/story/174808/pns-mehran-attack-vulnerable-embarrassed-and-targeted/.

Hendricks, K., and McAfee, R. P. 2006. Feints. *Journal of Economics & Management Strategy* 15(2):431–456.

Hespanha, J. P.; Ateskan, Y.; and Kizilocak, H. 2000. Deception in non-cooperative games with partial information. In *Proceedings of the 2nd DARPA-JFACC Symposium on Advances in Enterprise Control*.

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 Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 689–696.

Leitmann, G. 1978. On generalized Stackelberg strategies. *Journal of Optimization Theory and Applications* 26(4):637–643.

Nguyen, K. C.; Alpcan, T.; and Basar, T. 2009. Security games with incomplete information. In *Proceedings of the 2009 IEEE International Conference on Communications*, 1–6.

Oliveros, S. 2005. Equilibrium bluffs: A model of rational feints. Technical report, Working paper, University of Wisconsin-Madison, Department of Economics.

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.

Rabinovich, Z.; Jiang, A. X.; Jain, M.; and Xu, H. 2015. Information disclosure as a means to security. In *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 645–653.

Rubinstein, A. 1985. A bargaining model with incomplete information about time preferences. *Econometrica* 53(5):1151–1172.

Shieh, E. A.; An, B.; Yang, R.; Tambe, M.; Baldwin, C.; DiRenzo, J.; Maule, B.; and Meyer, G. 2012. PROTECT: An application of computational game theory for the security of the ports of the United States. In *Proceedings of the 26th Conference on Artificial Intelligence (AAAI)*, 2173–2179.

Spence, M. 1973. Job market signaling. *The Quarterly Journal of Economics* 355–374.

Xu, H.; Rabinovich, Z.; Dughmi, S.; and Tambe, M. 2015. Exploring information asymmetry in two-stage security games. In *Proceedings of the 29th Conference on Artificial Intelligence (AAAI)*, 1057–1063.

Zhuang, J., and Bier, V. M. 2011. Secrecy and deception at equilibrium, with applications to anti-terrorism resource allocation. *Defence and Peace Economics* 22(1):43–61.