

Lazy Defenders Are Almost Optimal Against Diligent Attackers

Avrim Blum

Computer Science Department
Carnegie Mellon University
avrim@cs.cmu.edu

Nika Haghtalab

Computer Science Department
Carnegie Mellon University
nika@cmu.edu

Ariel D. Procaccia

Computer Science Department
Carnegie Mellon University
arielpro@cs.cmu.edu

Abstract

Most work building on the Stackelberg security games model assumes that the attacker can perfectly observe the defender’s randomized assignment of resources to targets. This assumption has been challenged by recent papers, which designed tailor-made algorithms that compute optimal defender strategies for security games with limited surveillance. We analytically demonstrate that in zero-sum security games, *lazy* defenders, who simply keep optimizing against perfectly informed attackers, are almost optimal against *diligent* attackers, who go to the effort of gathering a reasonable number of observations. This result implies that, in some realistic situations, limited surveillance may not need to be explicitly addressed.

Introduction

The deployment of game-theoretic algorithms for security (Tambe 2012) is one of computational game theory’s biggest success stories to date. These algorithms are currently in use by major security agencies such as the US Coast Guard, the Federal Air Marshals Service, and the LA International Airport Police.

The basic model involves a *Stackelberg* game with two players, a leader (the defender) and a follower (the attacker). The defender deploys its security resources in a randomized way, by choosing a distribution over assignments of resources to feasible subsets of potential targets (for example, by sending a patrol to a certain area). The attacker then observes this randomized strategy and attacks a target with maximum expected payoff. The defender’s goal is to compute an optimal randomized strategy to commit to.

One of the implicit assumptions underlying this model is that the attacker has accurate information about the defender’s randomized strategy, presumably through surveillance. But in practice we would expect the attacker to obtain only a limited number of observations. This insight motivates the recent work of An et al. (2012; 2013), who design tailor-made algorithms that specifically optimize against an attacker with a

given number of samples of the defender’s randomized strategy.

We, too, are interested in understanding the implications of limited surveillance in security games. But rather than designing new or improved algorithms, as is common in the security games literature, we take the opposite approach by analytically demonstrating that, in realistic situations, limited surveillance may not be a major concern.

Results and Implications

We would like to compare two hypothetical worlds:

1. The defender optimizes against an attacker with unlimited observations (i.e., complete knowledge of the defender’s strategy), but the attacker actually has only k observations. This is, presumably, the current *status quo*.
2. The defender optimizes against an attacker with k observations, and, miraculously, the attacker indeed has exactly k observations.

We assume (without loss of generality) that the defender’s utility function is normalized to be in $[-1, 1]$. Our main result is the following:

Theorem 1 (informal). *For any zero-sum security game with m security resources such that each resource can cover at most d targets, and for any number of samples k , The difference between the defender’s utility in world 2 and world 1 is at most $O\left(\sqrt{\frac{\ln(mdk)}{k}}\right)$.*

The loss decreases rather quickly as the number of samples k grows, and scales gracefully with m and d . Note that the bound does not depend on the overall number of targets, denoted n . In most deployed applications we would expect m and d to be tiny compared to n . For example, there are very few federal air marshals compared to the staggering number of flights, and each can cover only one flight at any given time. We complement this result with a lower bound that shows an almost tight dependence on k .

The zero-sum assumption may seem restrictive, but in practice it is quite mild. Indeed, many deployed security-games-based systems actually use zero-sum payoffs; among others, these include PROTECT (Shieh et al. 2012) and TRUSTS (Yin et al. 2012). We give an example that shows that Theorem 1 does not hold when the game is not zero sum.

Assuming a limited-surveillance situation, moving from world 1 to world 2 requires very significant (deployment and computational) effort, and existing solutions (An et al. 2012; 2013) need new modeling assumptions to come up with an estimate for the attacker’s number of observations k . The operational upshot of our main result is that existing, deployed algorithms — which do, in fact, optimize for an attacker that accurately observes the defender’s mixed strategy, and therefore live in world 1 — may be good enough, despite the limited surveillance concern.

Related Work

Our paper is most closely related to two papers by An et al. (2012; 2013). The first paper (An et al. 2012) proposes a model for security games with limited surveillance. In this model, the attacker has a prior over defender strategies, which is updated as observations are obtained (we use a slightly simpler, but still realistic, belief update model). They present a few theoretical examples of interesting phenomena in this model. Their main contribution is optimization techniques for computing optimal defender strategies in the limited surveillance setting, and experiments that show that these techniques are effective. The games generated in the experiments are not constrained to be zero-sum, so these empirical results are not at odds with our theoretical results. In addition, the defender is assumed to be able to exactly compute the number of observations that the attacker will obtain.

The second paper (An et al. 2013) relaxes the last assumption by giving a Markov Decision Process (MDP) model for the attacker’s problem of obtaining additional (costly) observations. They also focus on optimization techniques for the new model, and on empirical results, which are, similarly, carried out using games that are not zero-sum. Our work complements these two papers by demonstrating that in zero-sum games, customized algorithms for limited surveillance may not be required.

A bit further afield, Yin et al. (2011) propose to model observational uncertainty by assuming that the difference between the defender’s strategy and the attacker’s belief about it is bounded. Pita et al. (2010) model settings with imperfect observability by reasoning about how humans perceive probability distributions. Similarly to the majority of papers on security games, these papers focus on new algorithmic techniques for the proposed extensions. Finally, it is worth mentioning that imperfect observability has also been studied in gen-

eral Stackelberg games (van Damme and Hurkens 1997; Morgan and Vardy 2007)

Preliminaries

A security game is a two-player general-sum game between a *defender* and an *attacker*. In this game, the defender commits to a mixed strategy that allocates a set of resources to defend a set of targets. The attacker observes part or all of the defender’s strategy and attacks a target. The defender and attacker both receive payoffs that depend on the target that was attacked and whether or not it was defended.

Formally, a security game is defined by a 5-tuple $(T, \mathcal{D}, R, A, u)$:

- $T = \{1, \dots, n\}$ is a set of n targets.
- R is a set of m resources.
- $\mathcal{D} \subseteq 2^T$, is a collection of subsets of targets, called *schedules*. Every $D \in \mathcal{D}$ represents a subset of targets that can be simultaneously defended by a resource. The *max-coverage* of \mathcal{D} is the maximum number of targets that any resource can defend simultaneously, i.e. $\max_{D \in \mathcal{D}} |D|$.
- $A : R \rightarrow 2^{\mathcal{D}}$ is a function that takes a resource and returns the set of schedules that can be defended using that resource. An assignment of resources to schedules is *valid* if for every $r \in R$, if r defends D then $D \in A(r)$.
- The *utilities* of target t are given by four functions: If target t is attacked, the defender receives $u_d^c(t)$ when t is covered and $u_d^u(t)$ when it is not covered. Similarly, the attacker receives $u_a^c(t)$ when t is covered and $u_a^u(t)$ when it is not covered. We assume that the attacker prefers it if the attacked target is not covered, and the defender prefers it if the attacked target is covered i.e. $u_d^u(t) \geq u_d^c(t)$ and $u_a^c(t) \leq u_a^u(t)$. Without loss of generality, we assume that the targets are ordered in decreasing order of $u_a^u(t)$. Furthermore, we assume (also without loss of generality) that the utilities are normalized to be in $[-1, 1]$. A security game is *zero-sum* if for any target t , $u_d^c(t) + u_a^c(t) = u_d^u(t) + u_a^u(t) = 0$.

A *pure strategy* of the defender is a valid assignment of resources to schedules. S is a *mixed strategy* (hereinafter, called strategy) if it is a distribution over the pure strategies. Given a defender’s strategy, the *coverage probability* of a target is the probability with which it is defended. Let S be a defender’s strategy, t be the attacker’s choice, and p_t be the coverage probability of target t under S ; then the defender’s and attacker’s utilities are, respectively, as follow:

$$U_d(S, t) = p_t u_d^c(t) + (1 - p_t) u_d^u(t)$$

$$U_a(S, t) = p_t u_a^c(t) + (1 - p_t) u_a^u(t)$$

		Attacker	
		k	∞
Defender	k	$\mathbb{E}_{S \sim (S_k^*)^k} U_d(S_k^*, b(S))$	$U_d(S_k^*, b(S_k^*))$
	∞	$\mathbb{E}_{S \sim (S_\infty^*)^k} U_d(S_\infty^*, b(S))$	$U_d(S_\infty^*, b(S_\infty^*))$

Table 1: Defender’s utilities when he optimizes against an attacker with k or infinitely many samples and the attacker observes k or infinitely many samples.

The attacker’s *best response* to a strategy S is defined by $b(S) = \arg \max_t U_a(S, t)$. We assume (purely for ease of exposition) that ties are broken in favor of targets with lower index and $b(S)$ indicates a single target.

In our model, the attacker may obtain a limited number of observations of the defender’s strategy. The attacker then responds to what he perceives the defender’s strategy to be; we assume that this belief is the empirical distribution over the attacker’s observed samples. For example, if the attacker has observed a certain pure strategy once, and another pure strategy twice, then he will believe that the former pure strategy is played with probability 1/3, and the latter pure strategy is played with probability 2/3.

Given a mixed strategy S and integer k , $S' \sim (S)^k$ is a random variable that indicates (the strategy equal to) the empirical distribution over k random draws from S . We define S_k^* to be the defender’s optimal strategy against an attacker that has k samples. In other words, $S_k^* = \arg \max_S \mathbb{E}_{S' \sim (S)^k} U_d(S, b(S'))$. Similarly, we define S_∞^* to be the defender’s optimal strategy against an attacker who has complete knowledge of the defender’s strategy. In other words, $S_\infty^* = \arg \max_S U_d(S, b(S))$.

Main Results

In this section, we focus on zero-sum security games. We show that S_∞^* is a good replacement for S_k^* in terms of its utility to the defender against an attacker with k observations. That is, we compare $\mathbb{E}_{S \sim (S_k^*)^k} U_d(S_k^*, b(S))$ (“World 2”, light gray cell in Table 1) and $\mathbb{E}_{S \sim (S_\infty^*)^k} U_d(S_\infty^*, b(S))$ (“World 1”, dark gray cell in Table 1), and establish that S_∞^* guarantees an additive gap of $O(\sqrt{\ln(mdk)}/k)$ (Theorem 1). We prove that this bound is relatively tight by showing that S_∞^* cannot guarantee a gap of $o(1/\sqrt{k})$ (Theorem 2) or even a constant gap when m and d are much larger than k (Theorem 3).

Upper Bound

Our main result is:

Theorem 1. *For any zero-sum security game with n targets, m resources, and a set of schedules with max-*

coverage d , and for any k ,

$$\begin{aligned} & \mathbb{E}_{S \sim (S_k^*)^k} U_d(S_k^*, b(S)) - \mathbb{E}_{S \sim (S_\infty^*)^k} U_d(S_\infty^*, b(S)) \\ & \in O\left(\sqrt{\frac{\ln mdk}{k}}\right) \end{aligned}$$

The next lemmas are required for proving Theorem 1. Lemma 1 shows an important relation between the utilities of the defender in various cases (see Table 1).

Lemma 1. *For any zero-sum security game and any integer k , $\mathbb{E}_{S \sim (S_k^*)^k} U_d(S_k^*, b(S)) \geq \mathbb{E}_{S \sim (S_\infty^*)^k} U_d(S_\infty^*, b(S)) \geq U_d(S_\infty^*, b(S_\infty^*))$*

Proof. The first inequality holds by the optimality of S_k^* against an attacker with k observations. The second inequality holds by the optimality of the minimax strategy in zero-sum games. In other words, for all t , $U_a(S_\infty^*, t) \leq U_a(S_\infty^*, b(S_\infty^*))$, so $\mathbb{E}_{S \sim (S_\infty^*)^k} U_a(S_\infty^*, b(S)) \leq U_a(S_\infty^*, b(S_\infty^*))$. However, the attacker’s utility is the negation of the defender’s utility, hence, $\mathbb{E}_{S \sim (S_\infty^*)^k} U_d(S_\infty^*, b(S)) \geq U_d(S_\infty^*, b(S_\infty^*))$. \square

The next two lemmas show that if the observed coverage probability on the attacked target is relatively accurate in S_k^* , then the defender’s utility is close to his utility in S_∞^* .

Lemma 2. *Let S be a defender’s strategy, and let p_i and q_i be the coverage probabilities of target i in S_k^* and S , respectively. Let $t = b(S)$ and assume that $p_t - q_t \leq \epsilon$. Then $U_d(S_k^*, b(S)) \leq U_d(S_\infty^*, b(S_\infty^*)) + 2\epsilon$.*

Proof. Since $t = b(S)$, we have

$$\begin{aligned} U_d(S_k^*, b(S)) &= p_t u_d^c(t) + (1 - p_t) u_d^u(t) \\ &\leq (q_t + \epsilon) u_d^c(t) + (1 - q_t - \epsilon) u_d^u(t) \\ &\leq q_t u_d^c(t) + (1 - q_t) u_d^u(t) + 2\epsilon \\ &\leq U_d(S, b(S)) + 2\epsilon \\ &\leq U_d(S_\infty^*, b(S_\infty^*)) + 2\epsilon, \end{aligned}$$

where the second transition holds because $u_d^c(t) \geq u_d^u(t)$, the third holds because $u_d^u(t), u_d^c(t) \in [-1, 1]$, and the last holds by the optimality of S_∞^* . \square

Lemma 3. *Let the defender have m resources, and a set of schedules with max-coverage d . Let S be a strategy with at most k pure strategies in its support, and let p_i and q_i be the coverage probabilities of target i in S_k^* and S , respectively. If $p_i - q_i \leq \epsilon$ for all $i \leq mdk + 1$, then $U_d(S_k^*, b(S)) \leq U_d(S_\infty^*, b(S_\infty^*)) + 2\epsilon$.*

Proof. Since S has at most k pure strategies in its support, there exists $j \leq mdk + 1$, such that $q_j = 0$. So, for all $i > mdk + 1$, $u_a^u(j) \geq u_a^u(i) \geq (1 - q_i) u_a^u(i) + q_i u_a^c(i)$, where the first inequality uses the fact the targets are sorted in decreasing order of $u_a^u(\cdot)$, and the last

inequality holds because $u_a^u(i) \geq u_a^c(i)$. So, i is not attacked. The rest comes directly from Lemma 2 and the fact that one of the first $mdk + 1$ targets is attacked. \square

Proof of Theorem 1. For $S \sim (S_k^*)^k$, let p_i and q_i be the coverage probabilities of target i in S_k^* and S , respectively. Let $\epsilon = \sqrt{\ln(mdk + 1) / k}$. Let P be the probability that for all $i \leq mdk + 1$, $p_i - q_i \leq \epsilon$. S surely has at most k pure strategies in its support, and therefore by Lemma 3,

$$\begin{aligned} \mathbb{E}_{S \sim (S_k^*)^k} U_d(S_k^*, b(S)) &\leq P \cdot (U(S_\infty^*, b(S_\infty^*)) + 2\epsilon) \\ &\quad + (1 - P) \cdot (U(S_\infty^*, b(S_\infty^*)) + 2) \\ &\leq U(S_\infty^*, b(S_\infty^*)) + 2\epsilon + 2(1 - P). \end{aligned}$$

Using the union bound and Hoeffding's inequality,

$$1 - P \leq \sum_{i \leq mdk + 1} \Pr(p_i - q_i > \epsilon) \leq (mdk + 1)e^{-2\epsilon^2 k}.$$

Hence,

$$\begin{aligned} \mathbb{E}_{S \sim (S_k^*)^k} U_d(S_k^*, b(S)) &\leq U_d(S_\infty^*, b(S_\infty^*)) + 2\epsilon \\ &\quad + 2(mdk + 1)e^{-2\epsilon^2 k} \\ &\leq U_d(S_\infty^*, b(S_\infty^*)) + 2\sqrt{\frac{\ln(mdk + 1)}{k}} + \frac{2}{mdk + 1} \\ &\leq \mathbb{E}_{S \sim (S_\infty^*)^k} U_d(S_\infty^*, b(S)) + 2\sqrt{\frac{\ln(mdk + 1)}{k}} \\ &\quad + \frac{2}{mdk + 1} \\ &= \mathbb{E}_{S \sim (S_\infty^*)^k} U_d(S_\infty^*, b(S)) + O\left(\sqrt{\frac{\ln(mdk)}{k}}\right), \end{aligned}$$

where the penultimate transition holds by Lemma 1. \square

One way to generalize Theorem 1 is to replace S_∞^* with an approximately optimal strategy. This is especially interesting because computing S_∞^* is NP-hard even for zero-sum security games (Korzhyk, Conitzer, and Parr 2010). By slightly modifying the proof of Theorem 1, we can show that if S_∞^* is replaced with an approximately optimal S^* such that $U_d(S^*, b(S^*)) \geq U_d(S_\infty^*, b(S_\infty^*)) - \alpha$, then the same upper bound holds up to an additional term of α .

Lower Bounds

The next result shows that Theorem 1 is almost tight with respect to k .

Theorem 2. *For any m, d , and k , there exists a zero-sum security game with m resources and schedules of max-coverage d , such that*

$$\begin{aligned} \mathbb{E}_{S \sim (S_k^*)^k} U_d(S_k^*, b(S)) - \mathbb{E}_{S \sim (S_\infty^*)^k} U_d(S_\infty^*, b(S)) \\ \in \Theta\left(\frac{1}{\sqrt{k}}\right) \end{aligned}$$

Target	Defender	
i	$u_d^c(\cdot)$	$u_d^u(\cdot)$
1	1	-1
\vdots	\vdots	\vdots
md	1	-1
$md + 1$	ϵ	$-\epsilon$
\vdots	\vdots	\vdots
$2md$	ϵ	$-\epsilon$

Table 2: Utilities associated with each target in Theorem 2. The attacker's utilities are the additive inverse of the defender's utilities.

Proof. Let there be $n = 2md$ targets, m resources, and a schedule set $\mathcal{D} = \{D \subseteq T : |D| = d\}$, such that any resource can defend any schedule. For $i \leq md$, let $u_d^c(i) = -u_d^u(i) = 1$, and for $i > md$, let $u_d^c(i) = -u_d^u(i) = \epsilon$, for an arbitrary $\epsilon < 0.17$ (as seen in Table 2). It is not difficult to see that S_∞^* covers each target with probability 0.5. These probabilities can be imposed using a mixed strategy that with probability 0.5 defends the first md targets, and with probability 0.5 defends the last md targets. In this case, the expected utility of every target is 0 and as a result $\mathbb{E}_{S \sim (S_\infty^*)^k} U_d(S_\infty^*, b(S)) = 0$.

Let the mixed strategy S_k be such that with probability $0.5 + 1 / \sqrt{4(k + 1)}$ it defends the first md targets and with probability $0.5 - 1 / \sqrt{4(k + 1)}$ it defends the last md targets. Clearly, this strategy imposes coverage probabilities of $p_i = 0.5 + 1 / \sqrt{4(k + 1)}$ for $i \leq md$, $p_i = 0.5 - 1 / \sqrt{4(k + 1)}$ for $i > md$.

We are interested in computing the probability that one of the first md targets is attacked. For $S \sim S_k$, let q_i be the observed coverage probability of target i in S . Due to the structure of S_k (defending the first or last half of the targets at any time), $q_i = q_j$ for all $i, j \leq md$ or $i, j > md$. So, the attacker would attack one of the first md targets if $q_i \geq 0.5$ for any (and all) $i > md$.

Let X be a random variable indicating the number of attacker's observations, out of k , in which target $md + 1$ (and any target $i > md$) is defended. Then

$$\begin{aligned} \Pr_{S \sim (S_k)^k} (q_{md+1} \geq \frac{1}{2}) &= \Pr_{X \sim \text{Bin}(k, p_{md+1})} (X \geq \frac{k}{2}) \\ &= \Pr_{X \sim \text{Bin}(k, p_{md+1})} \left(X \geq kp_{md+1} + \frac{k}{\sqrt{4(k+1)}} \right) \\ &= \Pr_{X \sim \text{Bin}(k, p_{md+1})} (X \geq \mathbb{E}[X] + \text{s.d}[X]) \\ &\approx \Pr_{X \sim \mathcal{N}(\mathbb{E}[X], \text{s.d}^2[X])} (X \geq \mathbb{E}[X] + \text{s.d}[X]) \\ &\approx 0.159, \end{aligned}$$

where the third transition holds because in the Bino-

mial distribution, $\mathbb{E}[X] = kp_{md+1}$ and $\text{s.d.}[X] = \sqrt{kp_{md+1}(1-p_{md+1})} = \frac{k}{\sqrt{4(k+1)}}$; and the fourth transition uses a Normal approximation for the Binomial distribution (which is valid because we are interested in a result that is asymptotic in k). Therefore,

$$\begin{aligned} & \mathbb{E}_{S \sim (S_k^*)^k} U_d(S_k^*, b(S)) - \mathbb{E}_{S \sim (S_\infty^*)^k} U_d(S_\infty^*, b(S)) \\ & \geq \Pr\left(q_{md+1} \geq \frac{1}{2}\right) U_d(S_k, 1) \\ & \quad + \Pr\left(q_{md+1} < \frac{1}{2}\right) U_d(S_k, md+1) - 0 \\ & \geq 0.159(2p_1 - 1) + 0.841\epsilon(2p_{md+1} - 1) \\ & \geq \frac{0.159 - 0.841\epsilon}{\sqrt{k+1}} \\ & \in \Theta\left(\frac{1}{\sqrt{k}}\right) \end{aligned}$$

□

The lower bound given in Theorem 2 does not improve (i.e., become larger) with m and d . We next show that when m and d are considerably larger than k , there is a security game in which the approximation gap between playing S_∞^* and S_k^* is a constant. Note that this bound is independent of the number of observations.

Theorem 3. *For any m, d , and k , such that $2md \geq \binom{2k}{k}$, and for any ϵ , there is a zero-sum security game with m resources and max-coverage d , such that regardless of the value of k*

$$\begin{aligned} & \mathbb{E}_{S \sim (S_k^*)^k} U_d(S_k^*, b(S)) - \mathbb{E}_{S \sim (S_\infty^*)^k} U_d(S_\infty^*, b(S)) \\ & \geq \frac{1}{2} - \epsilon \end{aligned}$$

The next lemma is needed for the proof of Theorem 3. This lemma shows that we can cover a set of elements with a relatively small collection of sets, such that each element appears in half of this collection, but if we take only half of the sets, then at least one element does not appear in their union.

Lemma 4. *For any set A and integer k such that $|A| = \binom{2k}{k}$, there exists $\mathcal{D} = \{D_1, \dots, D_{2k}\} \subseteq 2^A$ such that: (1) for all $i \leq 2k$, $|D_i| = |A|/2$; (2) for any $a \in A$, a is in exactly k members of \mathcal{D} ; and (3) for any k elements of \mathcal{D} , there exists $a \in A$ that does not appear in any of them, i.e. if $\mathcal{D}' \subset \mathcal{D}$ and $|\mathcal{D}'| \leq k$, then $\bigcup \mathcal{D}' \neq A$.*

Proof. There are $\binom{2k}{k} = |A|$ subsets of the set $[2k] = \{1, \dots, 2k\}$ with size k , so, there is a bijection between these subsets and elements of A . For any $S \subseteq [2k]$ such that $|S| = k$, let $a_S \in A$ be the member of A that corresponds to S . For all $i \leq 2k$, let $D_i = \{a_S : S \subseteq [2k], |S| = k, i \in S\}$. Define $\mathcal{D} = \{D_i : i \leq 2k\}$. We will show that \mathcal{D} has the required properties.

First, for any D_i , if $a_S \in D_i$ then $a_{S^c} \notin D_i$, where S^c is the complement of S . So $|D_i| = |A|/2$. Second, for a given S and for all $i \in S$, $a_S \in D_i$. Since $|S| = k$, a_S is in exactly k members of \mathcal{D} . Third, for any $\mathcal{D}' \subset \mathcal{D}$, such that $|\mathcal{D}'| \leq k$, $|\{i : D_i \notin \mathcal{D}'\}| \geq k$. Let S' be any k -subset of $\{i : D_i \notin \mathcal{D}'\}$. For all $D_i \in \mathcal{D}'$, $a_{S'} \notin D_i$. Hence, $\bigcup \mathcal{D}' \neq A$. □

Now, we are ready to prove Theorem 3.

Proof of Theorem 3. Let there be $n = \lceil md/\epsilon \rceil$ targets and let any resource be able to defend any d -sized subset of the targets. For any target $i \leq n$, let $u_d^c(i) = 1$ and $u_d^u(i) = 0$. In the best case, S_∞^* defends every target equally with probability md/n . So $\mathbb{E}_{S \sim (S_\infty^*)^k} U_d(S_\infty^*, b(S)) \leq md/n \leq \epsilon$.

Now we define a mixed strategy S^* . Consider the set A of targets $1, \dots, \binom{2k}{k}$. Let $\{D_1, \dots, D_{2k}\}$ be the subsets of targets with the properties mentioned in Lemma 4. For every D_i , let S_i be the pure strategy in which every target in D_i is defended. This is a valid strategy because $|D_i| = |A|/2 \leq md$ (by the theorem's assumption) and any resource can defend any d -sized subset of targets. Define S^* to be the uniform distribution over S_1, \dots, S_{2k} . Since, each target is defended in k out of $2k$ schedules, for all $i \leq \binom{2k}{k}$, the coverage probability of target i is $p_i = 0.5$.

After any k observations, at most k of these strategies are observed. Using Lemma 4, there is at least one target that is not covered in any of these observations, so it is perceived to have the highest expected utility to the attacker. So, the attacker always attacks one of the first $\binom{2k}{k}$ targets. Hence, $\mathbb{E}_{S \sim (S_k)^k} U_d(S_k, b(S)) = 0.5$ and $\mathbb{E}_{S \sim (S_k^*)^k} U_d(S_k^*, b(S)) - \mathbb{E}_{S \sim (S_\infty^*)^k} U_d(S_\infty^*, b(S)) \geq \frac{1}{2} - \epsilon$. □

(Impossible) Extensions

In this section we examine the assumptions and choices made in the statement of Theorem 1, and show that they are necessary. In other words, we give counterexamples to two potential extensions.

Multiplicative vs. Additive Approximation

Theorem 1 asserts that the utility of playing S_∞^* gives a good *additive* approximation to that of S_k^* against an attacker that makes k observations. Another common measure of performance guarantee is *multiplicative* approximation. However, because expected utilities can be positive or negative, a multiplicative approximation is not a suitable measure. Indeed, in the next example, we demonstrate that for any $k < n - 1$ there exists a game with n targets, one resource, and singleton schedules, such that the multiplicative gap between $\mathbb{E}_{S \sim (S_k^*)^k} U_d(S_k^*, b(S))$ and $\mathbb{E}_{S \sim (S_\infty^*)^k} U_d(S_\infty^*, b(S))$ is infinitely large, although these values are additively close.

Targets	Defender's Utility	
	$u_d^c(\cdot)$	$u_d^u(\cdot)$
i	1	$-\frac{1}{n-1}$
\vdots	\vdots	\vdots
n	1	$-\frac{1}{n-1}$

Table 3: Utilities associated with each target in Example 1. The attacker's utilities are the zero-sum complement of the defender's utilities.

Example 1. Let there be n targets and one resource that can defend any single target. For all $i \leq n$, let $u_d^c(i) = 1$ and $u_d^u(i) = -\frac{1}{n-1}$ and let the attacker's utilities be the additive inverse of the defender's (see Table 3). Let P_i represent the coverage probability of i in S_∞^* . Then for all $i \leq n$, $P_i = \frac{1}{n}$ and the utility of every target is 0. So, $\mathbb{E}_{S \sim (S_\infty^*)^k} U_d(S_\infty^*, S) = 0$.

Let S_k be the strategy with $p_i = \frac{1}{k+1}$ for all $i \leq k+1$, and $p_i = 0$ otherwise. Note that in any k observations, at least one of the first $k+1$ targets is never defended. This target has higher perceived utility to the attacker than any target $i > k+1$. Then,

$$\begin{aligned} \mathbb{E}_{S \sim (S_k^*)^k} U_d(S_k^*, b(S)) &\geq \mathbb{E}_{S \sim (S_k)^k} U_d(S_k, b(S)) \\ &\geq \min_{i \leq k+1} p_i u_d^c(i) + (1 - p_i) u_d^u(i) \\ &\geq \frac{1}{k+1} - \frac{k}{(k+1)(n-1)} \\ &> 0. \end{aligned}$$

It follows that

$$\frac{\mathbb{E}_{S \sim (S_k^*)^k} U_d(S_k^*, b(S))}{\mathbb{E}_{S \sim (S_\infty^*)^k} U_d(S_\infty^*, b(S))} = \infty,$$

but using Theorem 1, the additive gap is $O\left(\sqrt{\frac{\ln k}{k}}\right)$.

General-Sum vs. Zero-Sum

Unlike the zero-sum setting, S_∞^* is not always a good alternative to S_k^* in the general-sum games (in terms of its utility to the defender). In the next example, we construct a (general-sum) security game where $\mathbb{E}_{S \sim (S_k^*)^k} U_d(S_k^*, b(S)) - \mathbb{E}_{S \sim (S_\infty^*)^k} U_d(S_\infty^*, b(S))$ is arbitrarily close to $\frac{1}{2}$, even for large k .

Example 2. Consider n targets and one resource that can defend any single target. Let $u_d^c(1) = u_a^u(1) = u_a^c(2) = 1$, $u_d^u(2) = -1$, and set other utilities to 0 (see Table 4). Let P_i be the coverage probability of target i in S_∞^* . Then, $P_1 = P_2 = 0.5$ and $P_i = 0$ for all $i > 2$.

For ease of exposition, assume k is odd. Let X be a random variable for the number of observations (out of

Targets	Defender		Attacker	
	$u_d^c(\cdot)$	$u_d^u(\cdot)$	$u_a^c(\cdot)$	$u_a^u(\cdot)$
1	1	0	0	1
2	0	-1	0	1
3	0	0	0	0
\vdots	\vdots	\vdots	\vdots	\vdots
n	0	0	0	0

Table 4: This table shows the utilities associated with each target in Example 2.

k) in which the first target is defended. Since, k is odd, $\Pr_{S \sim (S_\infty^*)^k}(X < k/2) = \Pr_{S \sim (S_\infty^*)^k}(X > k/2) = 0.5$. Therefore, $\mathbb{E}_{S \sim (S_\infty^*)^k} U_d(S_\infty^*, b(S)) = 0$.

On the other hand, let S_k be the strategy in which $p_1 = 1 - p_2 = 0.5 + \epsilon$ and $p_i = 0$ for all $i > 2$ and a small $\epsilon > 0$. Using the Central Limit Theorem,

$$\lim_{k \rightarrow \infty} \Pr_{S \sim (S_k)^k} \left(X > \frac{k}{2} \right) = 1.$$

Therefore, for large enough k that is also odd, the attacker will attack target 1 with high probability. Because $\epsilon > 0$ is arbitrarily small, $\mathbb{E}_{S \sim (S_k^*)^k} U_d(S_k^*, b(S)) \approx 1/2$, and hence the gap is arbitrarily close to $1/2$.

Discussion

Limited surveillance is not the only way in which the Stackelberg security games model (Kiekintveld et al. 2009) is unrealistic. Motivated by the practical applications of security games, researchers have pointed out the model's shortcomings, and have proposed custom-made algorithms to alleviate them, in a slew of papers. For example, there may be different types of attackers (Paruchuri et al. 2008; Jain, Kiekintveld, and Tambe 2011), but even taking that into account, the defender's estimate of the attacker's payoffs may be inaccurate (Kiekintveld, Marecki, and Tambe 2010; 2011). And instead of choosing a best response to the attacker's strategy, the attacker may play suboptimally (Pita et al. 2010; Yang, Ordóñez, and Tambe 2012; Jiang et al. 2013). Most papers tend to focus on algorithmic solutions for one shortcoming, and some papers (especially those that are more theoretically inclined) abstract all of these issues away (Korzhyk, Conitzer, and Parr 2010; Korzhyk et al. 2011).

This is why we believe that our work offers a new perspective on security games research. Our conceptual approach is unusual, in seeking to understand in which cases we do *not* need new algorithms, thereby helping focus the community's attention on the most pressing algorithmic problems. Going forward, we imagine that a similar approach would be useful in all aspects of security games research.

References

- An, B.; Kempe, D.; Kiekintveld, C.; Shieh, E.; Singh, S. P.; Tambe, M.; and Vorobeychik, Y. 2012. Security games with limited surveillance. In *Proceedings of the 26th AAAI Conference on Artificial Intelligence (AAAI)*, 1242–1248.
- An, B.; Brown, M.; Vorobeychik, Y.; and Tambe, M. 2013. Security games with surveillance cost and optimal timing of attack execution. In *Proceedings of the 12th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 223–230.
- Jain, M.; Kiekintveld, C.; and Tambe, M. 2011. Quality-bounded solutions for finite Bayesian Stackelberg games: scaling up. In *Proceedings of the 10th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 997–1004.
- Jiang, A. X.; Nguyen, T. H.; Tambe, M.; and Procaccia, A. D. 2013. Monotonic maximin: A robust Stackelberg solution against boundedly rational followers. In *Proceedings of the 4th Conference on Decision and Game Theory for Security (GameSec)*, 119–139.
- Kiekintveld, C.; Jain, M.; Tsai, J.; Pita, J.; Ordóñez, F.; and Tambe, M. 2009. Computing optimal randomized resource allocations for massive security games. In *Proceedings of the 8th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 689–696.
- Kiekintveld, C.; Marecki, J.; and Tambe, M. 2010. Robust Bayesian methods for Stackelberg security games. In *Proceedings of the 9th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 1467–1468.
- Kiekintveld, C.; Marecki, J.; and Tambe, M. 2011. Approximation methods for infinite Bayesian Stackelberg games: Modeling distributional payoff uncertainty. In *Proceedings of the 10th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 1005–1012.
- Korzhyk, D.; Yin, Z.; Kiekintveld, C.; Conitzer, V.; and Tambe, M. 2011. Stackelberg vs. Nash in security games: An extended investigation of interchangeability, equivalence, and uniqueness. *Journal of Artificial Intelligence Research* 41:297–327.
- Korzhyk, D.; Conitzer, V.; and Parr, R. 2010. Complexity of computing optimal Stackelberg strategies in security resource allocation games. In *Proceedings of the 24th AAAI Conference on Artificial Intelligence (AAAI)*, 805–810.
- Morgan, J., and Vardy, F. 2007. The value of commitment in contests and tournaments when observation is costly. *Games and Economic Behavior* 60(2):326–338.
- Paruchuri, P.; Pearce, J. P.; Marecki, J.; Tambe, M.; Ordóñez, F. F.; and Kraus, S. 2008. Playing games for security: An efficient exact algorithm for solving Bayesian Stackelberg games. In *Proceedings of the 7th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 895–902.
- 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.
- Shieh, E.; An, B.; Yang, R.; Tambe, M.; Baldwin, C.; DiRenzo, J.; Maule, B.; and Meyer, G. 2012. PROTECT: A deployed game theoretic system to protect the ports of the United States. In *Proceedings of the 11th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 13–20.
- Tambe, M. 2012. *Security and Game Theory: Algorithms, Deployed Systems, Lessons Learned*. Cambridge University Press.
- van Damme, E., and Hurkens, S. 1997. Games with imperfectly observable commitment. *Games and Economic Behavior* 21(1–2):282–308.
- Yang, R.; Ordóñez, F.; and Tambe, M. 2012. Computing optimal strategy against quantal response in security games. In *Proceedings of the 11th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 847–854.
- Yin, Z.; Jain, M.; Tambe, M.; and Ordóñez, F. 2011. Risk-averse strategies for security games with execution and observational uncertainty. In *Proceedings of the 25th AAAI Conference on Artificial Intelligence (AAAI)*, 758–763.
- Yin, Z.; Jiang, A. X.; Tambe, M.; Kiekintveld, C.; Leyton-Brown, K.; Sandholm, T.; and Sullivan, J. P. 2012. TRUSTS: Scheduling randomized patrols for fare inspection in transit systems using game theory. *AI Magazine* 33(4):59–72.