

Defender (Mis)coordination in Security Games*

Albert Xin Jiang

University of Southern California
jiangx@usc.edu

Ariel D. Procaccia

Carnegie Mellon University
arielpro@cs.cmu.edu

Yundi Qian

University of Southern California
yundi.qian@usc.edu

Nisarg Shah

Carnegie Mellon University
nkshah@cs.cmu.edu

Milind Tambe

University of Southern California
tambe@usc.edu

Abstract

We study security games with multiple defenders. To achieve maximum security, defenders must perfectly synchronize their randomized allocations of resources. However, in real-life scenarios (such as protection of the port of Boston) this is not the case. Our goal is to quantify the loss incurred by miscoordination between defenders, both theoretically and empirically. We introduce two notions that capture this loss under different assumptions: the price of miscoordination, and the price of sequential commitment. Generally speaking, our theoretical bounds indicate that the loss may be extremely high in the worst case, while our simulations establish a smaller yet significant loss in practice.

1 Introduction

Security games, a special class of Stackelberg games, have been deployed as decision aids to schedule limited security resources at critical infrastructure sites including key airports and ports in the United States (e.g., Los Angeles, Boston, New York) as well as in protecting transportation infrastructure (international flights and metro trains) [Tambe, 2011; Shieh *et al.*, 2012]. These deployments have fueled the security games research area, with focus on efficient computation for scale-up and handling the significant uncertainty in this domain [Korzhyk *et al.*, 2011a; Basilico *et al.*, 2009; Yin *et al.*, 2010; Jain *et al.*, 2010].

Previous research in security games has assumed a single defender agency with control over all the security resources even if there were multiple attacker types or multiple coordinated attackers [Paruchuri *et al.*, 2008; Korzhyk *et al.*, 2011b]. Yet at most major critical infrastructure sites, e.g., port of New York or the Los Angeles International Airport, multiple defender agencies including city police, state and federal law enforcement agencies are responsible for security [Panel, 2011]; these agencies each have control over their own security resources and schedules.

In theory, achieving perfect coordination is simple: defenders can pool their resources together, and employ a single centralized algorithm to allocate resources. Our work is

motivated by the basic observation that this ideal is far from reality; in practice we observe limited coordination between defenders, which we believe can lead to significantly poorer performance. For example, it is well known that at major ports such as the port of New York, the US Coast Guard and different police departments each schedule their own boat patrols and security resources independently; there is no centralized scheduling. Our research question is therefore: *how much do defenders lose due to lack of coordination?* While the ideal coordination mechanism may not be feasible from a policy perspective, our goal is to use it as a gold standard in order to inform policy makers and ultimately bring about a higher level of coordination.

To that end, the paper offers four key contributions. First, we introduce and analyze the *price of miscoordination* (PoM) in simultaneous move settings and provide bounds on the worst case loss under different assumptions of target values. Second, we introduce the *price of sequential commitment* (PoSC) when defenders move sequentially. We illustrate that even in situations with low PoM, PoSC may be unbounded. Third, we introduce techniques to compute PoM and PoSC. Finally, we present experimental results based on realistic port scenarios illustrating a smaller yet significant loss in practice.

2 Background on Security Games

A security game is a 5-tuple (T, S, R, A, U) . T is the set of targets; we denote $|T| = n$. A *schedule* is a subset of the target set T ; $S \subseteq 2^T$ denotes the collection of feasible schedules. R is the set of resources that belong to the defender; we denote $|R| = m$. Every resource has a collection of schedules (a subset of S) to which it can be assigned; $A : R \rightarrow 2^S$ is the function that maps a resource to its collection of possible schedules. In our theoretical results we make the standard assumption that for any resource, any subset of a feasible schedule is itself a feasible schedule [Yin *et al.*, 2010]. When a resource is assigned to a schedule, we say that all the targets in the schedule are *covered* by the resource.

The payoffs are given by four different utility functions. If target t is attacked, the defender's utility is $U_d^c(t)$ if t was covered, and $U_d^u(t)$ if t was not covered. Similarly, the attacker's utility is $U_a^c(t)$ if t was covered, and $U_a^u(t)$ if t was not covered. We assume that $U_d^c(t) \geq U_d^u(t)$ and $U_a^c(t) \leq U_a^u(t)$. Note that it makes no difference to the play-

*This research is supported by US Coast Guard grant HSHQDC-10-D-00019 and MURI grant W911NF-11-1-0332.

ers' utilities whether a target is covered by one resource or by more than one resources.

A pure strategy of the defender is an assignment of resources to feasible schedules; a defender may opt to employ a *mixed* strategy, which selects at random from pure strategies according to a distribution. The solution of the security game is the optimal *Stackelberg* strategy for the defender. Given the defender's mixed strategy that randomly allocates resources to schedules, the attacker selects a best response by choosing to attack a target that maximizes its expected utility. The defender chooses its mixed strategy to maximize its own utility given that the attacker best responds.

Given the emphasis on explicit game-theoretic models and equilibrium analysis, our paper also complements research on strategies for multi-robot patrol [Agmon, 2010; Agmon *et al.*, 2009].

3 The Price of Miscoordination

We extend the basic model to the setting where multiple defenders, with their disjoint sets of resources, defend a set of targets against a single attacker. In particular, we analyze the need for the defenders to coordinate their moves, and the incurred loss when coordination is lacking.

3.1 Our Model

Let D be the set of defenders; we denote $|D| = d$. Let the set of resources R be partitioned into $\{R_i\}_{i \in D}$ where R_i is the set of resources owned by defender $i \in D$. We take an optimistic point of view by assuming that all defenders are interested in overall security, hence they are all endowed with the same utility functions U_d^c and U_d^u as before.

The maximum utility is achieved by the defenders when they pool their resources together, and schedule the whole set R of resources as if they were owned by a single defender. We call this strategy profile the *optimal correlated profile* (OCP) as the mixed strategies of individual defenders are correlated. In other words, the exact realizations of assignments of the resources of different defenders may be dependent on each other. This represents the scenario with full *coordination*.

Consider an alternative scenario, where the defenders choose their mixed strategies but these mixed strategies are uncorrelated. In other words, the instantiations of the various mixed strategies are independent of each other. We define the *optimal uncorrelated profile* (OUP) as the profile of uncorrelated mixed strategies for the defenders that yields maximum utility among all profiles of uncorrelated mixed strategies. Clearly, the utility to the defenders under OUP is no greater, and may be strictly smaller, than the utility under OCP. We define the supremum (over a given class of security games) of the ratio of the utility under OCP to that under OUP as the *price of miscoordination* (PoM).

Example 1. Let there be two defenders, defender 1 with resource r_1 and defender 2 with resource r_2 . Let $T = \{t_1, t_2, t_3\}$, with $U_d^u(t) = U_a^c(t) = 0$ and $U_d^c(t) = U_a^u(t) = 1$ for all $t \in T$. Resource r_1 can cover either target t_1 or target t_2 , and resource r_2 can cover either target t_2 or target t_3 . The optimal correlated profile (OCP) uniformly randomizes between assigning the resources to $\{t_1, t_2\}$, $\{t_1, t_3\}$, and

$\{t_2, t_3\}$. Each target is covered with probability $2/3$, hence the defenders are guaranteed utility $2/3$. In contrast, under the optimal uncorrelated profile (OUP) each defender covers its own target (t_1 or t_3) with probability $(\sqrt{5}-1)/2 = 0.618$, and the shared target t_2 with the complement probability (which equalizes the coverage probability of t_2 with t_1 and t_3). The defenders' utility is therefore 0.618, and the ratio of the utilities under OCP and OUP is 1.078, which is therefore a lower bound for the PoM.

The PoM is related to the notion of *mediation value* [Ashlagi *et al.*, 2008; Bradonjic *et al.*, 2009] (which in turn is inspired by the *price of anarchy* [Koutsoupias and Papadimitriou, 1999; Roughgarden and Tardos, 2002]). Briefly, the mediation value in a game is the ratio between the maximum social welfare (i.e., sum of utilities) in any correlated equilibrium and the maximum social welfare in any mixed-strategy Nash equilibrium. If we define an artificial game between the defenders, where the utility function of each defender is the same as the common utility function, then the OCP is the welfare-maximizing correlated equilibrium and the OUP is the welfare-maximizing Nash equilibrium. Hence, the PoM coincides with the mediation value in this game.

To guarantee that the PoM is meaningful, we assume hereinafter that the defender utilities are non-negative. General security games may have negative utilities; however, by shifting all defender utilities by the minimum defender utility we can obtain a security game with non-negative defender utilities, which is equivalent in terms of its optimal strategies for the defenders and the attacker. Moreover, the PoM of the shifted game has a natural interpretation in the original game: informally, it tells us what fraction of the *gap* between the worst possible outcome and the best possible correlated outcome is due to miscoordination.

3.2 Bounds on the PoM

Our first result shows that in some games coordination is crucial, as the price of miscoordination may be arbitrarily large.

Theorem 1. *The PoM is unbounded in general security games.*

Proof. Let there be two defenders, defender 1 with resource r_1 and defender 2 with resource r_2 . Let $T = \{t_1, t_2, t_3\}$. Resource r_1 can cover either target t_1 or target t_2 , and resource r_2 can cover either target t_2 or target t_3 . In terms of structure, this is identical to the game of Example 1. However, the utilities are defined differently. For the defenders, targets t_1 and t_3 are identical. If either target is attacked, it gives the defenders utility x if covered and 1 if uncovered, where $x > 1$. Target t_2 , if attacked, gives the defenders utility 0 whether it was covered or not. Formally, $U_d^c(t_1) = U_d^c(t_3) = x$, $U_d^u(t_1) = U_d^u(t_3) = 1$ and $U_d^c(t_2) = U_d^u(t_2) = 0$. For the attacker, targets t_1 and t_3 give utility 1 if uncovered and 0 if covered. Target t_2 gives the attacker utility z if uncovered and 0 if covered. Formally, $U_a^c(t) = 0$ for all $t \in T$, $U_a^u(t_1) = U_a^u(t_3) = 1$ and $U_a^u(t_2) = z$.

Irrespective of the values of x and z , the defenders never let the attacker attack target t_2 as it would give them the worst possible utility (zero utility). Now fix x and increase z . As

$z \rightarrow \infty$, the utility of the attacker for a successful attack against target t_2 increases, hence it must be covered by the defenders with probability 1 in the limit to avoid such an attack. Thus, as $z \rightarrow \infty$, the OCP for the defenders converges to assigning $r_1 \rightarrow \{t_1\}$ and $r_2 \rightarrow \{t_2\}$ with probability $1/2$, and $r_1 \rightarrow \{t_2\}$ and $r_2 \rightarrow \{t_3\}$ with probability $1/2$. This way, target t_2 is fully covered while both targets t_1 and t_3 are covered with probability $1/2$, giving the defenders utility $1/2 + x/2$.

In the limit of the uncorrelated case, at least one defender must cover target t_2 with probability 1. The other defender can only cover one of targets t_1 and t_3 . The attacker attacks the other target, resulting in a utility of 1 for the defenders. Thus, taking $z \rightarrow \infty$ in this security game implies that the ratio of utility under OCP and OUP is at least $(1/2 + x/2)/1$. As x can be arbitrarily large, the PoM is unbounded. \square

The proof of Theorem 1 relies on an extreme asymmetry in the utilities for different targets. It is natural to ask, what is the PoM when the targets are *identical*? Specifically, the utility functions of the defenders satisfy $U_d^c(t) = x$ and $U_d^u(t) = 0$ for all targets $t \in T$ and some $x \geq 0$, where the utility for targets being uncovered is zero since we shift the utility functions to make the worst-case utility zero. Similarly, the utility functions of the attacker satisfy $U_a^c(t) = 0$ and $U_a^u(t) = x'$ for all targets $t \in T$ and some $x' \geq 0$. We believe that although this is a special case, it is useful for practitioners to explore the extreme points of the problem. It is also worth noting that the case of identical targets is still rich, even from a complexity-theoretic point of view: it immediately follows from existing proofs [Korzhyk *et al.*, 2010, Theorem 5] that computing the optimal defender strategy for this case is \mathcal{NP} -hard, even for one defender and schedules of size at most 3.

It is easy to check that in the case of identical targets, the attacker would attack the target that is covered with minimum probability in order to maximize his own utility. Therefore, the utility of the defenders is proportional to the minimum probability with which any target is covered. Let c_t denote the coverage probability of target $t \in T$ in a strategy profile, and let $c_{\min} = \min_{t \in T} c_t$ denote the minimum coverage probability. Then we are interested in the ratio $c_{\min}(OCP)/c_{\min}(OUP)$. We show that in this case the PoM is upper-bounded by a small constant.

Theorem 2. *The PoM of security games with identical targets is at most $\frac{e}{e-1} \approx 1.582$.*

Proof. Given a security game with identical targets and d defenders, we claim that there exists an OCP where no target is covered by more than one defenders simultaneously in any pure strategy realization of the OCP. Indeed, if a target is covered by more than one defender in a realization, we remove resources of all but one defender assigned to that target in that realization (the assumption that any subset of a feasible schedule is a feasible schedule plays a key role here).

Next, consider any such OCP. Let $c_{t,i}$ denote the probability with which defender i covers target t in the OCP. Since a target is not covered by more than one defender at a time, the coverage probability of target t is $\sum_{i=1}^d c_{t,i}$, so $c_{\min}(OCP) = \min_{t \in T} \sum_{i=1}^d c_{t,i}$.

Now consider the *marginal uncorrelated profile* (MUP) where each defender i schedules its resources as they were scheduled in the OCP, but the mixed strategies of different defenders are now uncorrelated.¹ In other words, each defender follows the mixed strategy that is obtained as the marginal of the OCP. In this case, target t is still covered by defender i with probability $c_{t,i}$. However, a target may be covered by more than one defender simultaneously, so the probability of target t being covered overall is $1 - \prod_{i=1}^d (1 - c_{t,i})$. Thus,

$$c_{\min}(OUP) \geq c_{\min}(MUP) = \min_{t \in T} \left[1 - \prod_{i=1}^d (1 - c_{t,i}) \right],$$

and it follows that

$$PoM \leq \frac{\min_{t \in T} \sum_{i=1}^d c_{t,i}}{\min_{t \in T} \left(1 - \prod_{i=1}^d (1 - c_{t,i}) \right)}.$$

We claim that for all $t \in T$,

$$1 - \prod_{i=1}^d (1 - c_{t,i}) \geq \left(1 - \frac{1}{e} \right) \cdot \sum_{i=1}^d c_{t,i}. \quad (1)$$

Indeed, if $\sum_{i=1}^d c_{t,i} = 0$, then Equation (1) follows trivially. Hence, assume that $\sum_{i=1}^d c_{t,i} > 0$. Now,

$$\frac{1 - \prod_{i=1}^d (1 - c_{t,i})}{\sum_{i=1}^d c_{t,i}} \geq \frac{1 - e^{-\sum_{i=1}^d c_{t,i}}}{\sum_{i=1}^d c_{t,i}} \geq 1 - \frac{1}{e},$$

where the first transition holds because $1 - x \leq e^{-x}$, and the last transition is true since $f(x) = (1 - e^{-x})/x$ is a decreasing function in $(0, 1]$, so $f(x) \geq f(1)$ for all $x \in (0, 1]$. Also, $\sum_{i=1}^d c_{t,i}$ is the coverage probability of target t in the OCP, so $0 < \sum_{i=1}^d c_{t,i} \leq 1$. We have thus established (1). We can now conclude that

$$\begin{aligned} PoM &\leq \frac{\min_{t \in T} \sum_{i=1}^d c_{t,i}}{\min_{t \in T} \left[1 - \prod_{i=1}^d (1 - c_{t,i}) \right]} \\ &= \max_{t \in T} \left[\frac{\min_{t' \in T} \sum_{i=1}^d c_{t',i}}{1 - \prod_{i=1}^d (1 - c_{t,i})} \right] \\ &\leq \max_{t \in T} \frac{\sum_{i=1}^d c_{t,i}}{1 - \prod_{i=1}^d (1 - c_{t,i})} \leq \frac{e}{e-1}, \end{aligned}$$

where the last transition is due to Equation (1). \square

Is it possible that the upper bound for identical targets is in fact much closer to 1? Our next theorem answers this question in the negative.

Theorem 3. *The PoM of security games with identical targets is at least $4/3$.*

¹The mixed strategy assignments of various resources of the same defender are still correlated as before.

Proof. Consider the following security game with identical targets. There are d defenders, each defender i has a single resource r_i . There are $d + 1$ identical targets, t_1, \dots, t_{d+1} . For $1 \leq i \leq d$, resource r_i can cover either target t_i or target t_{i+1} .

Consider the following $d + 1$ pure strategies where strategy i leaves only target i uncovered. This is uniquely achieved by assigning resource r_j to target t_j for $j < i$ and resource r_j to target t_{j+1} for $j \geq i$. Since at least one target must be uncovered in any pure strategy, the OCP uniformly randomizes over these $d + 1$ pure strategies to achieve the optimal minimum coverage probability $c_{\min}(OCP) = 1 - 1/(d + 1)$.

We next prove that $c_{\min}(OUP) < 0.75$. Suppose for contradiction that $c_{\min}(OUP) \geq 0.75$, so all targets are covered with probability at least 0.75. First, we prove by induction that for all $1 \leq i \leq d$, resource r_i covers target t_i with probability at least 0.5. For the base case of resource r_1 , this is obvious since target t_1 is covered with probability at least 0.75 and r_1 is the only resource that can cover it. Suppose it is true for resource r_k . Observe that resource r_k covers target t_k with probability at least 0.5, hence it covers target t_{k+1} with probability at most 0.5. If resource r_{k+1} covers target t_{k+1} with probability less than 0.5, then due to the lack of correlation between the assignments of resources r_k and r_{k+1} , the probability of coverage of t_{k+1} would be less than $0.5 + 0.5 - 0.5 \cdot 0.5 = 0.75$, which contradicts the assumption. Hence, r_{k+1} must cover t_{k+1} with probability at least 0.5. Therefore, the induction hypothesis holds. In particular, resource r_d covers target t_d with probability at least 0.5, and hence target t_{d+1} with probability at most 0.5. It follows that the coverage probability of target t_{d+1} is at most 0.5, which yields a contradiction. Hence, $c_{\min}(OUP) < 0.75$.

We conclude that $PoM \geq (1 - 1/(d + 1))/0.75$. Taking $d \rightarrow \infty$, we get that $PoM \geq 4/3$. \square

4 The Price of Sequential Commitment

The PoM is optimistic in a sense, because it compares the optimal correlated utility with the optimal uncorrelated utility. Even achieving the optimal uncorrelated utility requires some level of coordination among the defenders: they are not pooling their resources together, but they are coordinating their mixed strategies. Below we consider a more pessimistic model of defender commitment, and show that in this model the loss is unbounded even when targets are identical.

4.1 Our Model

We assume that the defenders commit *sequentially*. The first defender commits to a mixed strategy that optimizes the joint utility function in the absence of the other defenders. Subsequently, each defender chooses a mixed strategy that maximizes the joint utility function given the strategies of the earlier defenders, in the absence of the later defenders; note that there may be more than one optimal strategy. The mixed strategies of various defenders are still uncorrelated.²

In such a model, we can now consider the loss in utility due to sequential commitment compared to the OCP. We de-

²Our results hold even if the defenders commit to mixed strategies that are correlated with the strategies of the earlier defenders.

fine two prices: the price under the best order of commitment (PoSC_b) and the price under the worst order of commitment (PoSC_w). Formally, PoSC_b (resp. PoSC_w) is the supremum (over all security games) of the ratio of the utility under the OCP to the maximum (resp. minimum) utility over all orders of sequential commitment.³ It is easy to check that the PoM is a lower bound for the PoSC_b, which in turn is a lower bound for the PoSC_w.

4.2 Bounds on the PoSC

While Theorem 2 shows that the PoM in security games with identical targets is upper-bounded by a small constant, we show that this is not the case even for PoSC_b in security games with identical targets.

Theorem 4. *The PoSC_b (even) in security games with identical targets is unbounded.*

Proof. Consider a security game with two defenders: defender 1 owns resource r_1 and defender 2 owns resource r_2 . The target set $T = T_1 \cup T_2$ consists of $2 \cdot k$ identical targets, where $|T_1| = |T_2| = k$. Resource r_1 can either cover all the targets in T_1 simultaneously or any single target in T_2 , and resource r_2 can either cover all the targets in T_2 simultaneously or any single target in T_1 .

The defender who commits first has $k + 1$ feasible schedules for its resource. To maximize the minimum coverage probability, it uniformly randomizes over the feasible schedules to cover each target with identical probability $1/(k + 1)$. Due to the uniform coverage by the first defender, the defender committing later also uniformly randomizes over its $k + 1$ feasible schedules. Hence, each target is covered with probability at most $2/(k + 1)$. Note that the optimal strategies are unique. In contrast, we have $c_{\min}(OCP) = 1$ because defenders 1 and 2 can respectively cover targets in T_1 and T_2 simultaneously. Thus, the price under the best order of commitment in this particular security game is at least $(k + 1)/2$. Hence, PoSC_b is $\Omega(k)$, and k can be arbitrarily large. \square

The proof of Theorem 4 uses a very simple construction: only two symmetric defenders and identical targets. However, to obtain the lower bound the defenders need to be able to cover an increasingly larger number of targets simultaneously. It turns out that we can obtain a matching upper bound that scales linearly with the maximum number of targets any defender can cover simultaneously; call this parameter the *max-simultaneous-coverage*. In the security game constructed in the proof of Theorem 4, the max-simultaneous-coverage is k . We additionally assume that every defender can cover every target — otherwise an optimal strategy for a defender who commits first and cannot cover all targets would be to not cover anything, as any strategy would yield a utility of 0. We call this property *complete individual coverage*.

Theorem 5. *Denote the max-simultaneous-coverage by k . Then the PoSC_w is $O(k)$ in security games with identical targets and complete individual coverage.*

³Our lower bound on PoSC_b in Theorem 4 (resp. upper bound on PoSC_w in Theorem 5) works even with the best (resp. worst) tie-breaking among the set of all optimal strategies at each step.

Proof. Let d be the number of defenders and n be the number of targets. Since any defender can cover at most k targets simultaneously, the total number of targets covered simultaneously in any pure strategy is at most $d \cdot k$. It follows that the minimum coverage probability in any mixed strategy (and hence in the OCP) is at most $(d \cdot k)/n$. Moreover, the minimum coverage probability cannot be more than 1. Hence, $c_{\min}(OCP) \leq \min(d \cdot k/n, 1)$.

Next, consider sequential commitment with the worst order. The first defender commits to a mixed strategy that maximizes the minimum coverage probability in absence of the other defenders; if this probability is p_1 , then it is easy to see that the worst choice of optimal strategy for the defender is the strategy that covers *every* target with probability exactly p_1 . Such a mixed strategy is feasible since one can reduce the coverage probability of a target as much as required by removing it from the schedules in some of the pure strategy assignments (recall that any subset of a schedule is a feasible schedule). Inductively, each successive defender also commits to a mixed strategy that covers all targets with identical probability, maximizing this probability. Denote the identical coverage probability by defender i as p_i . Due to complete individual coverage, $p_i \geq 1/n$ for all i as the defenders can achieve uniform coverage probability of at least $1/n$ by uniformly randomizing over pure strategies that cover different targets. Hence, every target is covered with probability at least $1 - \prod_{i=1}^d (1 - p_i) \geq 1 - (1 - 1/n)^d \geq 1 - e^{-d/n}$. Thus,

$$\begin{aligned} PoSC_w &\leq \frac{\min(d \cdot k/n, 1)}{1 - e^{-d/n}} \leq k \cdot \frac{\min(d/n, 1)}{1 - e^{-d/n}} \\ &\leq k \cdot \frac{1}{1 - e^{-1}} = O(k), \end{aligned}$$

where the third transition holds because the function $f(x) = \min(x, 1)/(1 - e^{-x})$ achieves its maximum at $x = 1$. \square

From the proof of Theorem 4 and the statement of Theorem 5, we obtain a complete picture regarding the price of sequential commitment.

Corollary 1. *Denote the max-simultaneous-coverage by k . Both $PoSC_b$ and $PoSC_w$ are $\Theta(k)$ in security games with identical targets and complete individual coverage.*

Note that $PoSC_b \geq PoM$, and Theorem 1 shows that the PoM can be unbounded even with $k = 1$ if targets are non-identical. However, that proof crucially uses a game that does not have the complete individual coverage property. We can actually construct an example with $k = 1$ and complete individual coverage where the price under the best order scales with the utilities, showing that $PoSC_b$ is unbounded for non-identical targets even with $k = 1$.

5 Experimental Results

We now turn to an empirical investigation of the PoM and PoSC; our goal is to quantify the loss caused by defenders' miscoordination in realistic security games.

We compute the PoM and PoSC in these games as follows. The OCP can be computed by solving the traditional security

game (with a single defender) where all the resources are assumed to belong to one defender. Computing PoSC is also easy; the optimal sequential strategies can be computed by rolling the coverage of earlier defenders into the utility functions, and at each step solving a traditional security game with a single defender using existing computational tools [Tambe, 2011]. The computation of the OUP is trickier; we formulate a nonlinear program whose solution is the OUP, and solve it using the YALMIP toolbox of MATLAB [Lofberg, 2004]; we omit the nontrivial details due to lack of space. We worry about the computational efficiency (or lack thereof) of our algorithms only insofar as it restricts our simulations, because — while our simulations can inform policy decisions — these are not computations that we expect defenders to perform on a regular basis.

For realistic security games, we use the port patrolling problem where a patrolling boat (resource) goes around and checks multiple targets a day [Shieh *et al.*, 2012]. Specifically, we use the actual map of Boston Port shown in Figure 1(a). The figure also shows the time required for the boat to move between various nodes. We assume that the boat requires two time units at each target it visits for the check procedure. Due to the extreme computational burden, we restrict ourselves to the case of two defenders, each with one boat. Figure 1(b) shows the locations of homebases of the two defenders (marked in black). The other nodes are considered potential targets. In any schedule, a boat starts from its homebase, visits some targets and finally returns back to its homebase. All visited targets are considered to be covered by the boat in that schedule. The boat must return within a patrolling time limit.




Fig. 1: Boston Port, homebases, and targets.

We consider two distributions of target valuations, both of which create non zero-sum security games with non-identical targets. The importance of non-zero-sum security games in realistic domains has been emphasized in the risk analysis literature [Powell, 2007]. Denote by $U[a, b]$ the uniform distribution over all integers between a and b .

- *Homogeneous distribution:* Under this distribution, every target's utility to the defenders is in $U[20, 39]$ if covered, and in $U[0, 9]$ if uncovered (and vice versa for the attacker). Thus, there is homogeneity in target valuations although the actual realizations may create non-identical targets.
- *Heterogeneous distribution:* This distribution is related to the target valuation of the worst case instance of Theorem 1, and captures realistic scenarios where some tar-

gets are far more valuable than others. For the *gray* targets (which are in some sense shared between the two defenders), the utility to the defenders is in $U[6, 10]$ if covered, and in $U[1, 5]$ if uncovered. The utility to the attacker for these targets is 0 if covered and in $U[30, 59]$ if uncovered. These targets are highly valuable; the defenders do not want them to be attacked whether or not they are covered, and the attacker has high utility for a successful attack. For the remaining targets, the utility to the defenders is in $U[30, 59]$ if covered, and in $U[10, 15]$ if uncovered. The utility to the attacker is 0 if covered, and in $U[1, 5]$ if uncovered. These represent targets which the attacker is almost indifferent about, and the defenders do not have drastically low utility even in case of a successful attack.

Note that these are shifted utility functions as explained in Section 3.1. Figures 2 and 3 show the PoM and PoSC for the homogeneous and heterogeneous distributions, respectively.⁴ All values are averaged over 100 random trials on target utilities. In both graphs, the x-axis shows the patrolling time limit for both defenders. Legends $PoSC_b$ and $PoSC_w$ show the average loss (under the 100 random trials) where the average is taken over the best and the worst sequence respectively in *every instance*, whereas legend $PoSC_{01}$ (resp. $PoSC_{10}$) shows the average loss under the fixed commitment sequence where defender 0 (resp. defender 1) commits first.




Fig. 2: PoM and PoSC under the Homogeneous Distribution.

Interestingly, the PoM in the homogeneous case (Figure 2) is almost 1 (the worst patrol times give an average PoM of 1.038), but the PoSC_w can be significant (as high as 1.749). As expected, the loss due to miscoordination is much higher in the heterogeneous case (Figure 3), with the PoM reaching 1.434, and the PoSC_w achieving a whopping 3.347.

Corollary 1 shows that both the PoSC_b and PoSC_w are $\Theta(k)$ where k is the max-simultaneous-coverage parameter, which reduces to the maximum schedule size (over all resources) when every defender has one resource. The maximum schedule size is a monotonic function of the patrolling time limit. Hence, the graphs show various losses as functions of the maximum schedule size. While the theoretical

⁴Technically the figures show average ratios rather than the PoM and PoSC themselves, which are defined as supremums.




Fig. 3: PoM and PoSC under the Heterogeneous Distribution.

results predict a monotonic increment in PoSC_b and PoSC_w in the worst-case, they seem to have an inverse U shape relationship with the maximum schedule size in our empirical analysis. The difference arises because the lower bound of Theorem 4 uses an example with an increasing number of targets while the targets are fixed in our case. Hence at a threshold patrolling time limit, the defenders can fully protect all targets in the OCP, and increasing it further maintains the utility in the OCP while increasing the utility in the uncorrelated case for both PoSC and PoM — this results in the downward curves when time limits are large. The low PoM and PoSC for low patrol time limits is due to the small overlap between the schedules of the two defenders, which entails less need for coordination among their strategies.

Another interesting observation is that in our particular graph structure, the commitment sequence where defender 1 commits first generally outperforms the sequence where defender 0 commits first. Further, always using the former sequence helps avoid the extremely high PoSC in the disastrous special cases. Thus, the commitment sequence (even if fixed) has great influence on the PoSC.

6 Discussion

Our theoretical results suggest that the loss due to miscoordination can be extremely high or relatively low, depending on the assumptions that are made. In general though, the PoM (which in the worst case is arbitrarily high) can be significant (more than 30% loss) in realistic simulations. The PoSC is consistently significant, with a loss of as much as 70% in our simulations. We view the PoSC as a better proxy of reality, and in fact one can argue that even the PoSC is optimistic. We therefore interpret our results as suggesting a need for greater coordination among defenders in critical infrastructure sites.

As we hinted in our introduction, the challenge for future work is not purely computational. Indeed, the simplest and most effective solution is scheduling the joint pool of resources, but it seems unrealistic to expect organizations as independent as, e.g., the New York Police Department and the Coast Guard to adopt such a policy. Hence, to design realistic coordination mechanisms, it is necessary to determine what feasible form of coordination can overcome most of the loss while keeping policy makers in the loop.

References

- [Agmon *et al.*, 2009] Noa Agmon, Sarit Kraus, Gal A Kaminka, and Vladimir Sadov. Adversarial uncertainty in multi-robot patrol. In *IJCAI*, pages 1811–1817, 2009.
- [Agmon, 2010] Noa Agmon. On events in multi-robot patrol in adversarial environments. In *AAMAS*, pages 591–598, 2010.
- [Ashlagi *et al.*, 2008] I. Ashlagi, D. Monderer, and M. Tennenholtz. On the value of correlation. *Journal of Artificial Intelligence Research*, 33(1):575–613, 2008.
- [Basilico *et al.*, 2009] N. Basilico, N. Gatti, and F. Amigoni. Leader-follower strategies for robotic patrolling in environments with arbitrary topologies. In *AAMAS*, 2009.
- [Bradonjic *et al.*, 2009] M. Bradonjic, G. Ercal-Ozkaya, A. Meyerson, and A. Roytman. On the price of mediation. In *EC*, pages 315–324, 2009.
- [Jain *et al.*, 2010] M. Jain, E. Kardes, C. Kiekintveld, F. Ordonez, and M. Tambe. Security games with arbitrary schedules: A branch and price approach. In *AAAI*, 2010.
- [Korzhyk *et al.*, 2010] D. Korzhyk, V. Conitzer, and R. Parr. Complexity of computing optimal Stackelberg strategies in security resource allocation games. In *Proc. of The 24th AAAI Conference on Artificial Intelligence*, pages 805–810, 2010.
- [Korzhyk *et al.*, 2011a] D. Korzhyk, V. Conitzer, and R. Parr. Solving stackelberg games with uncertain observability. In *AAMAS*, 2011.
- [Korzhyk *et al.*, 2011b] Dmytro Korzhyk, Vincent Conitzer, and Ronald Parr. Security games with multiple attacker resources. In *IJCAI*, 2011.
- [Koutsoupias and Papadimitriou, 1999] E. Koutsoupias and C. Papadimitriou. Worst-case equilibria. In *STACS*, pages 404–413, 1999.
- [Lofberg, 2004] J. Lofberg. Yalmip: A toolbox for modeling and optimization in matlab. In *CACSD*, pages 284–289. IEEE, 2004.
- [Panel, 2011] Los Angeles Mayor’s Blue Ribbon Panel. Report of the Mayor’s Blue Ribbon Panel on airport security, 2011.
- [Paruchuri *et al.*, 2008] P. Paruchuri, J. P. Pearce, J. Marecki, M. Tambe, F. Ordonez, and S. Kraus. Playing games for security: An efficient exact algorithm for solving bayesian stackelberg games. In *AAMAS*, 2008.
- [Powell, 2007] R. Powell. Defending against terrorist attacks with limited resources. *American Political Science Review*, 101(3):527–541, 2007.
- [Roughgarden and Tardos, 2002] T. Roughgarden and É. Tardos. How bad is selfish routing? *Journal of the ACM*, 49(2):236–259, 2002.
- [Shieh *et al.*, 2012] E. Shieh, B. An, R. Yang, M. Tambe, C. Baldwin, J. DiRenzo, B. Maule, and G. Meyer. Protect: A deployed game theoretic system to protect the ports of the united states. In *AAMAS*, 2012.
- [Tambe, 2011] M. Tambe. *Security and Game Theory: Algorithms, Deployed Systems, Lessons Learned*. Cambridge University Press, New York, NY, 2011.
- [Yin *et al.*, 2010] Z. Yin, D. Korzhyk, C. Kiekintveld, V. Conitzer, and M. Tambe. Stackelberg vs. nash in security games: Interchangeability, equivalence, and uniqueness. In *AAMAS*, 2010.