

Equilibrium Analysis of Multi-Defender Security Games

Jian Lou and Yevgeniy Vorobeychik

Electrical Engineering and Computer Science
Vanderbilt University

{jian.lou,yevgeniy.vorobeychik}@vanderbilt.edu

Abstract

Stackelberg game models of security have received much attention, with a number of approaches for computing Stackelberg equilibria in games with a single defender protecting a collection of targets. In contrast, multi-defender security games have received significantly less attention, particularly when each defender protects more than a single target. We fill this gap by considering a multi-defender security game, with a focus on theoretical characterizations of equilibria and the price of anarchy. We present the analysis of three models of increasing generality, two in which each defender protects multiple targets. In all models, we find that the defenders often have the incentive to over-protect the targets, at times significantly. Additionally, in the simpler models, we find that the price of anarchy is unbounded, linearly increasing both in the number of defenders and the number of targets per defender. Surprisingly, when we consider a more general model, this results obtains only in a “corner” case in the space of parameters; in most cases, however, the price of anarchy converges to a constant when the number of defenders increases.

1 Introduction

With terrorism and cyber threats ever on people’s minds, developing and refining our understanding of both security threats and responses to these is both an important research area, and of great practical value. Game theory has come to play an important role in the security domain, with considerable modeling and algorithmic advances, as well as actual deployment of security systems in practice that are based on such models and algorithms, including LAX Airport [Jain *et al.*, 2008; Pita *et al.*, 2009], US Coast Guard [Shieh *et al.*, 2012], and the Federal Air Marshals Service [Jain *et al.*, 2010a; 2010b; Kiekintveld *et al.*, 2009], among others.

A popular game theoretic model of security that has received much attention both in the research and in practice is a Stackelberg game between a single defender and a single attacker, in which the defender commits to a randomized strategy, while the attacker, upon learning this strategy, chooses an optimal target or a subset of targets to attack [Conitzer

and Sandholm, 2006]. In most of the associated literature, it is assumed that a single defender is responsible for all the targets that need protection, and that she has control over all of the security resources. However, there are many domains in which there are multiple defender agencies who are in charge of different subsets of all targets. While sometimes such agencies can be aligned to follow the same set of goals, in general different defender entities exhibit at least some disparities in goals. In particular, a defender is typically responsible (financially, politically, or legally) for targets in their direct charge, rather than other targets that may have social importance. This is certainly the case for the private sector, where different corporations secure their own resources without necessarily much concern for those of others, but is also common for the public sector, with different government agencies held accountable for their own assets, and not for those of others. In such *non-cooperative* security scenarios, the typical single-defender Stackelberg game model is clearly inadequate. Instead, we must consider the consequences of *strategic interactions* among *multiple defenders*, each charged with protecting their assets from common adversaries. An important consideration in such games is the *negative externalities* that security decisions impose on others: specifically, when a defender chooses a high level of security investment, budget-constrained attackers are more likely to choose others to attack. The resulting dynamics is likely to lead to over-investment in security, a phenomenon observed in several related efforts [Bachrach *et al.*, 2013].

We consider a problem with multiple defenders protecting a collection of homogeneous targets. Each defender chooses a probability distribution over protection levels for all targets in their charge. A single attacker then best responds to the defenders’ action by attacking the target with the lowest probability to be protected, breaking ties uniformly at random. Our analysis is focused on three models of such multi-defender games, with defenders acting non-cooperatively in all of these. We show that a Nash equilibrium among defenders in this two-stage game model need not always exist, even when the defenders utilize randomized strategies (i.e., probability distributions over target protection levels); this is distinct from a model in which the attacker moves simultaneously with the defenders, where a mixed strategy equilibrium is guaranteed to exist. When an equilibrium does exist, we show that the defenders protect all of their targets

with probability 1 in all three models, whereas the socially optimal protection levels are generally significantly lower. When no equilibrium exists, we characterize the best approximate Nash equilibrium (that is, one in which defenders have the least gain from deviation), showing that over-investment is substantial in this case as well. Our *price of anarchy* (PoA) analysis, which relies on the unique equilibrium when it exists, and the approximate equilibrium otherwise, demonstrates a surprising finding: whereas PoA is unbounded in the simpler models, increasing linearly with the number of defenders, the more general model shows this to be an atypical special case achieved when several parameters are exactly zero. More generally, PoA tends to a constant as the number of defenders increases.

Related Work Although most work in Stackelberg models of security concerns computing a Stackelberg equilibrium in a single-defender single-attacker scenario, there are several exceptions. [Jiang *et al.*, 2013] consider (mis)-coordination in cases where there are multiple defenders who are responsible for different sets of targets and share the common utility function over all targets. In this work, the defenders are fundamentally cooperative (sharing identical goals), however, making it distinct from our contribution. [Bachrach *et al.*, 2013] examined non-cooperative security games among many defenders, in a two-stage model, but imposed strong assumptions on the model structure, and only considered one-dimensional continuous “security investment” strategies for the defender (departing significantly from the typical structure of Stackelberg security games, in which defensive strategies are discrete protection choices). [Smith *et al.*, 2014] extend the standard computational Stackelberg game framework to analyze games with multiple defenders, but offer no theoretical analysis. [Chan *et al.*, 2012] propose interdependent defense (IDD) games, to study aspects of the interdependence of risk and security in a natural extension of interdependent security (IDS) games previously proposed by Heal and Kunreuther [Heal and Kunreuther, 2002] to consider attackers as explicit players in the game. In IDD games, unlike our setting, defenders and the attacker move simultaneously. To our knowledge, no previous work considers a theoretical analysis of multi-defender games with defenders protecting multiple targets, making previous literature on the subject qualitatively distinct from the typical practical considerations of single-defender settings, where resource allocation among multiple targets is a fundamental concern.

2 Modeling Multi-Defender Security Games

Our modeling effort proceeds in three steps, each generalizing the previous. As we see below, each generalization step reveals new and surprising insights about the multi-defender security setting, allowing us to appreciate the fundamental incentive forces.

2.1 The Baseline Model

We start with a model which most reflects the related literature: in particular, this model involves n defenders and a single attacker, with each defender engaged in protecting a single target. Each target has the same value to the defender

$v > 0$. We suppose that the defender has two discrete choices: to protect the target, or not. In addition, the defender can randomly choose among these; our focus is on these *coverage probabilities* (i.e., the probability of protecting, or covering, the target), which we denote by s_i for a given defender i . The attacker is strategic and could observe the defenders’ strategies to choose a target so as to maximize the damage. We assume that attacker is indifferent among the targets, and attacks the target with the lowest coverage probability, breaking ties uniformly at random. In a given scenario, for all defenders, the attacker’s strategy is a vector of probabilities $P = \langle p_1, p_2, \dots, p_n \rangle$, where p_i is the probability of attacking a target i , with $\sum_{i=1}^n p_i = 1$.

We assume that if the attacker chooses to attack a target corresponding to defender i and defender i chooses to protect the target, then the utility of the defender i is 0, and if the attacker attacks the target but it is not protected, then the utility of the defender is $-v$. If a defender chooses to cover a target, it will incur a cost $c > 0$. Additionally, we assume that the defender gets a utility of zero whenever another defender’s target is attacked. We can thus define the expected utility of a defender i as

$$u_i = p_i u_i^a + (1 - p_i) u_i^u,$$

where u_i^a is the utility of i if it is attacked, and u_i^u is the utility of i if it is not attacked. By the assumptions above,

$$u_i^a = -(1 - s_i)v - s_i c = -v + s_i(v - c)$$

$$u_i^u = -s_i c.$$

2.2 The Multi-Target Model

Our key conceptual departure from related work is in allowing each defender to protect multiple targets, aligning it better with practical security domains. Specifically, suppose that there are n defenders, each protecting $k \geq 1$ targets. Then the strategy of defender i will be a vector $\langle s_{i1}, s_{i2}, \dots, s_{ik} \rangle$. The strategy profile of the attacker can then be described as a matrix of probabilities,

$$\begin{pmatrix} p_{11} & p_{12} & \dots & p_{1k} \\ p_{21} & p_{22} & \dots & p_{2k} \\ \vdots & \vdots & \ddots & \vdots \\ p_{n1} & p_{n2} & \dots & p_{nk} \end{pmatrix}$$

in which $\sum_{i=1}^n \sum_{j=1}^k p_{ij} = 1$ and $p_{ij} \geq 0$ for each i and j . The expected utility of a defender i in this model is

$$u_i = \sum_{j=1}^k p_{ij} u_{ij}^a + (1 - p_{ij}) u_{ij}^u,$$

where u_{ij}^a is the utility of target j to defender i if it is attacked, and u_{ij}^u is the utility of target j to i if it is not attacked. Using the notation introduced earlier, we have

$$u_{ij}^a = -(1 - s_{ij})v - s_{ij} c = -v + s_{ij}(v - c)$$

$$u_{ij}^u = -s_{ij} c.$$

2.3 The General Model

Finally, we analyze the most general version of the model considered. We assume that if the attacker chooses to attack a target i and the defender i chooses to protect the target, then the utility of the defender i is U_c , and if the attacker attacks the target but it is not protected, then the utility of the defender is U_u . It is reasonable to assume that $U_c \geq U_u$. If the target of defender i is not attacked, then we assume that the utility of defender i is $T \geq U_c$. Other assumptions are same as those in the multi-target model. In the general model, therefore,

$$u_{ij}^a = s_{ij}U_c + (1 - s_{ij})U_u - s_{ij}c = (U_c - U_u - c)s_{ij} + U_u$$

$$u_{ij}^u = T - s_{ij}c.$$

2.4 Solution Concepts

We consider a class of solutions to multi-defender security games where the defenders' simultaneously commit to a probability distribution s over their pure strategies (corresponding to protection decisions for each target), and the attacker subsequently chooses an optimal target to attack, breaking ties uniformly at random. Since the attacker's behavior is straightforward (given s), we will focus on the simultaneous-move game among the defenders, and Nash equilibria thereof (giving rise to subgame perfect equilibria of the two-stage game of interest). As we demonstrate below, Nash equilibria are not guaranteed to exist, in which case we focus on ϵ -equilibria, in which no player gains more than ϵ by deviating; in particular, we will consider ϵ -equilibria with the smallest attainable ϵ . In fact, we could see Nash equilibrium as a special case of ϵ -equilibrium in which $\epsilon = 0$.

To measure how the efficiency of the game degrades due to selfish behavior of its defenders, we consider *Utilitarian Social Welfare* and (ϵ)-*Price of Anarchy* in our paper. *Utilitarian Social Welfare* means the sum of all defenders' payoffs. For the smallest attainable ϵ , we define ϵ -Price of Anarchy (ϵ -PoA) as follows:

$$\epsilon\text{-PoA} = \frac{SW_O}{SW_E}$$

where SW_O means the optimal (utilitarian) social welfare we can get from the game, SW_E means the worst-case (utilitarian) social welfare in ϵ -equilibrium. An underlying assumption of this definition is that the value of SW_O and SW_E are both positive. If they are both negative, then ϵ -PoA will be the reciprocal of above equation. We should also note that the ordinary *Price of Anarchy* can be seen as a special case of ϵ -Price of Anarchy in which $\epsilon = 0$.

3 The Baseline Model

Our first result presents necessary and sufficient conditions for the existence of a Nash equilibrium in the baseline model, and characterizes it when it does exist. We omit the proofs of Baseline Model (this section) and Multi-target Model (next section) due to space constraints, and they can be found in extended version of the paper on the website.

Theorem 1. *In the Baseline model, Nash equilibrium exists if and only if $v \geq c$. In this equilibrium all targets are protected with probability 1.*

Thus, if a Nash equilibrium does exist, it is unique, with all defenders always protecting their target. But what if the equilibrium does not exist? Next, we characterize the (unique) ϵ -equilibrium with the minimal ϵ that arises in such a case. We will use this approximate equilibrium strategy profile as a *prediction* of the defenders' strategies.

Theorem 2. *In the Baseline model, if $v < c$, the optimal ϵ -equilibrium is for all defenders to cover their target with probability $\frac{v}{c}$. The corresponding ϵ is $\frac{v(c-v)}{cn}$.*

Armed with a complete characterization of predictions of strategic behavior among the defenders, we can now consider how this behavior related to socially optimal protection decisions. Since the solutions are unique, there is no distinction between the notions of *price of anarchy* and *price of stability*; we term the ratio of socially optimal welfare to welfare in equilibrium as the price of anarchy for convenience.

First, we characterize the socially optimal outcome.

Theorem 3. *In the Baseline model, the optimal social welfare SW_O is*

$$SW_O = \begin{cases} -cn, & \text{if } v \geq cn; \\ -v, & \text{if } v < cn. \end{cases}$$

From this result, it is already clear that defenders systematically over-invest in security, except when values of the targets are quite high. This stems from the fact that the attacker creates a *negative externality* of protection: if a defender protects his target with higher probability than others, the attacker will have an incentive to attack another defender. In such a case, we can expect a "dynamic" adjustment process with defenders increasing their security investment well beyond what is socially optimal. To see just how much the defenders lose in the process, we now characterize the price of anarchy of our game.

If $v \geq c$, it is one and only one Nash equilibrium when all defenders have the coverage probability 1 for their targets. And the corresponding social welfare is

$$SW_E = -cn$$

Because it is the only Nash equilibrium, we could get the *Price of Anarchy* as follows:

$$PoA = \begin{cases} 1, & \text{if } v \geq cn; \\ \frac{nc}{v}, & \text{if } c < v < cn. \end{cases}$$

Figure 1 shows the relationship among Price of Anarchy, number of defenders, and ratio of cost c and value v . From the figure we could find that when number of defenders and ratio of c and v are small enough (e.g. $n \leq 5$ and $\frac{c}{v} = 0.2$), the price of anarchy is close to 1. Otherwise, the price of anarchy is unbounded, growing linearly with n .

If $v < c$, there is no Nash equilibrium. However, we could get the optimal ϵ -equilibrium when all defenders have the same coverage probability $\frac{v}{c}$ for their targets. The corresponding Social Welfare is

$$SW_E = (v - cn)\frac{v}{c} - v$$




Figure 1: Price of Anarchy when $v \geq c$

Similarly, we could get the $\frac{v(c-v)}{cn}$ -Price of Anarchy as follows,

$$\frac{v(c-v)}{cn} \cdot PoA = \frac{cn + c - v}{c},$$

which is, again, linear in n .

4 The Multi-Target Model

Armed with observations from the model with a single target for each defender, we now extend the model to a case not as yet considered in the literature in a theoretical light: each defender protects a set of k targets. This gives rise to a combinatorial set of possible decisions for each defender, so that even computing a best response is not necessarily easy. Remarkably, we are able to characterize equilibria and approximate equilibria in this setting as well.

Our first result is almost a mirror-image of the corresponding result in the baseline model: when a Nash equilibrium exists, all defenders protect all of their targets with probability 1.

Theorem 4. *In the Multi-Target model, Nash equilibrium exists if and only if $v \geq kc$. In this equilibrium all targets are protected with probability 1.*

Next, we consider scenarios when $v < kc$, in which there is no Nash equilibrium. Our next result characterizes optimal (lowest- ϵ) approximate equilibria.

Theorem 5. *In the Multi-Target model, if $v < kc$, then in the optimal ϵ -equilibrium all targets are protected with probability $\frac{v}{kc}$. The corresponding ϵ is $\frac{v(kc-v)}{cnk}$.*

Thus, as n increases, the optimal approximate equilibrium approaches a Nash equilibrium. Figure 2 illustrates the relationship between ϵ and the number of targets each defender protects when $v = 10$ and $c = 1$. In this figure, $\epsilon = 0$ when $k \leq 10$, which means that an exact Nash equilibrium exists; ϵ increases with k when $k > 10$, but at a decreasing rate, converging to $\frac{v}{n}$ when $k \rightarrow \infty$.

Finally, we characterize socially optimal welfare, and, subsequently, put everything together in describing the price of anarchy.




Figure 2: ϵ value when $v = 10, c = 1$

Theorem 6. *In the Multi-Target model, the optimal social welfare SW_O is*

$$SW_O = \begin{cases} -cnk, & \text{if } v \geq cnk; \\ -v, & \text{if } v < cnk. \end{cases}$$

Thus, just as in the baseline case, the defenders will generally over-invest in security.

If $v \geq kc$, there is a unique Nash equilibrium with all targets protected with probability 1. The corresponding social welfare is

$$SW_E = -cnk$$

Because it is the only Nash equilibrium, the *Price of Anarchy* is

$$PoA = \begin{cases} 1, & \text{if } v \geq cnk; \\ \frac{cnk}{v}, & \text{if } ck \leq v < cnk. \end{cases}$$

If $v < kc$, there is no Nash equilibrium. However, we could get the optimal approximate equilibrium when all defenders have the same coverage probability $\frac{v}{kc}$ for their targets. The corresponding Social Welfare is

$$SW_E = (v - cnk) \frac{v}{kc} - v$$

And the $\frac{v(kc-v)}{cnk}$ -Price of Anarchy is

$$\frac{v(kc-v)}{cnk} \cdot PoA = n + 1 - \frac{v}{kc}$$

Clearly, in either case, and just as in the baseline model, the price of anarchy is unbounded, growing linearly with n .

We now consider how PoA changes as a function of k , i.e. the number of targets each defender has. When $k \leq \frac{v}{cn}$, a Nash equilibrium exists and the PoA is 1; when $\frac{v}{cn} < k \leq \frac{v}{c}$, PoA increases linearly in k with the slope $\frac{nc}{v}$. However, when $k > \frac{v}{c}$, a Nash equilibrium does not exist and the approximate PoA is $n + 1 - \frac{v}{kc}$, which increases very slowly with k , and is bounded by $n + 1$ when $k \rightarrow \infty$. Figure 3 illustrates the relationship between (approximate) Price of Anarchy and k for $n = 2$. When k is very small, PoA = 1. For intermediate k , PoA increases linearly, and when k is sufficiently large, Nash equilibrium no longer exists, and ϵ -PoA increases quite slowly, converging to 3 when $k \rightarrow \infty$.




Figure 3: (Approximate) Price of Anarchy when $\frac{c}{v} = 0.1$

5 The General Model

Both the baseline and the multi-target models made rather strong assumptions about the structure of the utility functions of the players. In the general model, we relax these assumptions, allowing for arbitrary utilities for the players when the target is attacked or not, and when it is protected or not (when attacked). Quite surprisingly, our findings here are *qualitatively different*: the special case of the baseline and multi-target models turns out to be an exception, rather than the rule when more general models are considered.

Just as before, we start by characterizing Nash and approximate Nash equilibria.

Theorem 7. *In the General model, Nash equilibrium exists if and only if $U_c - U_u \geq kc - \frac{(n-1)(T-U_c)}{n}$. In this equilibrium all targets are protected with probability 1.*

Proof. We firstly claim that Nash equilibrium can appear *only if* coverage probability of all of targets t_{ij} are identical. Otherwise, there will be a target t_{ik} which has the possibility of 0 to be attacked, and defender i has incentive to decrease s_{ik} . To find Nash equilibrium, we need only consider scenarios in which all targets have the same coverage probability.

When all targets have the same possibility s to be protected. Then for each defender, her expected utility is

$$u = \frac{(U_c - U_u - nkc)s + U_u + (nk - 1)T}{n}$$

If $s < 1$, then some defender i could slightly increase s to $s + \delta$ (δ is a very small positive real number) for all of her targets to make sure them not be attacked and get the value $u' = kT - k(s + \delta)c$,

$$u' - u = \frac{(U_c - U_u)(1 - s) + (T - U_c) - nkc\delta}{n}$$

As $U_c \geq U_u$, $T \geq U_c$, and δ can be very small, $u' - u > 0$ when $s < 1$. We could know that the defender has incentive to improve s when $s < 1$. So the Nash equilibrium can appear *only if* $s_{ij} = 1$ for all defenders' targets s_{ij} . When all targets have the same possibility $s = 1$ to be covered, for each defender, her expected utility is

$$u = \frac{U_c - nkc + (nk - 1)T}{n}$$

We claim that, if a defender i has incentive to deviate, the “optimal” deviation could appear *only if* defender i has the same protection probability s' for all her targets. Otherwise, for some target t_{ik} which has the probability 0 to be attacked, she could always decrease s' to get higher utility. If probabilities of targets protected by defender i are all s' ($0 \leq s' < 1$), then her expected utility is $u' = (U_c - U_u - c)s' + U_u + (k - 1)(T - s'c)$, and

$$u' - u = (U_c - U_u - kc)(s' - 1) + \frac{(n - 1)(U_c - T)}{n}$$

- 1) If $U_c - U_u \geq kc$, then $u' - u \leq 0$, we could know that it is a Nash equilibrium when all targets have the probability 1 to be protected.
- 2) If $U_c - U_u < kc$, the maximal value of $u' - u$ corresponds to $s' = 0$,

$$\max_{0 \leq s' < 1} u' - u = -(U_c - U_u - kc) - \frac{(n - 1)(T - U_c)}{n}$$

If $kc - \frac{(n-1)(T-U_c)}{n} \leq U_c - U_u < kc$, $u' - u \leq 0$, it is a Nash equilibrium; otherwise, it is not.

To sum up, Nash equilibrium exists *if and only if* $U_c - U_u \geq kc - \frac{(n-1)(T-U_c)}{n}$, and the equilibrium corresponds to all targets having probability 1 of being protected. \square

Next, we characterize the optimal approximate equilibrium when no Nash equilibrium exists.

Theorem 8. *In the General model, in the optimal ϵ -equilibrium all targets are protected with probability $\frac{T-U_u}{kc}$. The corresponding ϵ is $\frac{(T-U_u)(kc-U_c+U_u)}{cnk}$.*

Proof sketch. When all targets have the same probability s of being protected, the expected utility of each defender is

$$u = \frac{(U_c - U_u - nkc)s + U_u + (nk - 1)T}{n}$$

Assume $0 \leq s < 1$. If some defender i slightly increases s to $s + \delta_{ij}$ for target t_{ij} , then she would obtain utility $u' = \sum_{j=1}^k T - (s + \delta_{ij})c$,

$$\begin{aligned} u' - u &= \frac{T - (U_c - U_u)s - U_u}{n} - \sum_{j=1}^k \delta_{ij}c \\ &\leq \frac{T - (U_c - U_u)s - U_u}{n} \end{aligned} \quad (1)$$

Then we consider scenarios in which a defender i could get higher utility by decreasing protection probability. We claim that the “optimal” deviation could appear *only if* defender i has the same protection probability s' for all of her targets. Otherwise, for some target t_{ik} which has the probability 0 to be attacked, she could always decrease coverage probability of t_{ik} to get higher utility.

Then we need only consider cases in which a defender deviates by decreasing probabilities of all her targets to $s - \delta$. And her utility will be $u'' = (U_c - U_u - kc)(s - \delta) + U_u + (k - 1)T$,

As $U_c - U_u < kc$, when $\delta = s$ (the maximal value of δ), we could get maximal value of $u'' - u$,

$$\max_{0 < \delta \leq s} u'' - u = \frac{T - (U_c - U_u)s - U_u}{nk} + kcs + U_u - T \quad (2)$$

By comparing the value of equation (1) and equation (2), we could get different value of ϵ for ϵ -equilibrium as shown below:

$$\epsilon = \begin{cases} \frac{T - (U_c - U_u)s - U_u}{nk}, & \text{if } 0 \leq s \leq \frac{T - U_u}{kc}; \\ \frac{T - (U_c - U_u)s - U_u}{n} + kcs + U_u - T, & \text{if } \frac{T - U_u}{kc} < s \leq 1. \end{cases}$$

When $s = \frac{T - U_u}{kc}$, we could get the minimal $\epsilon = \frac{(T - U_u)(kc - U_c + U_u)}{cnk}$.

We claim that the $\frac{(T - U_u)(kc - U_c + U_u)}{cnk}$ -equilibrium can appear *only if* all targets have the same coverage probability s . We prove this by contradiction. Suppose that targets have different coverage probabilities. This gives rise to two cases: 1) Each defender uses an identical coverage probability for each target she owns (these may differ between defenders); 2) There exists a defender, who has a different probability to protect her own targets.

In case 1), there exist β defenders ($1 \leq \beta < n$) who have the same minimal probability s' to protect all of their targets. The expected utility for each defender among these β defenders is:

$$u = \frac{(U_c - U_u - k\beta c)s' + U_u + (k\beta - 1)T}{\beta}$$

When $\frac{T - U_u}{kc} < s' \leq 1$, some defender i among these β defenders could decrease probability of all her targets to 0 to get value $u_1 = U_u + (k - 1)T$,

$$\begin{aligned} u_1 - u &= \frac{T - (U_c - U_u)s' - U_u}{\beta} + kcs' + U_u - T \\ &> \frac{T - (U_c - U_u)s' - U_u}{n} + kcs' + U_u - T \end{aligned}$$

When $0 \leq s' \leq \frac{T - U_u}{kc}$, some defender i among these β defenders could slightly increase probabilities of all her targets to $s' + \delta_3$ to get the utility $u_2 = kT - k(s' + \delta_3)c$

$$\begin{aligned} u_2 - u &= \frac{T - (U_c - U_u)s' - U_u - k\beta c\delta_3}{\beta} \\ &> \frac{T - (U_c - U_u)s' - U_u}{n} \end{aligned}$$

The above inequality holds because δ_3 can be arbitrarily small. Thus, no profile in case 1) can be a $\frac{(T - U_u)(kc - U_c + U_u)}{cnk}$ -equilibrium.

In case 2), for defenders who have different coverage probability for their own targets, they could always increase payoff by decreasing some of their targets' probability to get a corresponding profile in case 1). Then we could know that any profile in case 2) cannot be a $\frac{(T - U_u)(kc - U_c + U_u)}{cnk}$ -equilibrium. \square

As the final step towards characterizing the Price of Anarchy, we derive optimal social welfare in this model.

Theorem 9. *In the General model, the optimal social welfare SW_O is*

$$SW_O = \begin{cases} U_c - nkc + (nk - 1)T, & \text{if } U_c - U_u \geq nkc; \\ U_u + (n - 1)T, & \text{if } U_c - U_u < nkc. \end{cases}$$

Proof sketch. We firstly claim that we could get optimal social welfare *only if* all targets have the same probability s to be protected their. Otherwise, some target t_{ij} has the probability of 0 to be attacked. Then we could decrease s_{ij} to get a better social welfare. Consequently, we need only to consider an optimal identical coverage probability s to obtain optimal social welfare, which can be done in a relatively straightforward way. \square

If $U_c - U_u \geq kc - \frac{(n-1)(T - U_c)}{n}$, the Nash equilibrium is unique, with all targets protected with probability 1. The corresponding social welfare in equilibrium is

$$SW_E = U_c - nkc + (nk - 1)T.$$

So far we have not yet added any constrains to value of T , U_c , and U_u (except that $T \geq U_c \geq U_u$). In order to make *Price of Anarchy* well-defined, we need to add constraints that values of T , U_c , and U_u are all non-positive (just as in the previous two models) or all non-negative. To be consistent with previous models, we add constraints that U_c , U_u and T are all non-positive (little changes if all are non-negative).

In the case of a unique Nash equilibrium, the price of anarchy is

$$PoA = \begin{cases} 1, & \text{if } U_c - U_u \geq nkc; \\ \frac{U_c - U_u - nkc}{U_u + (nk - 1)T} + 1, & \text{if } kc - \frac{(n-1)(T - U_c)}{n} \leq U_c - U_u < nkc. \end{cases}$$

If $U_c - U_u < kc - \frac{(n-1)(T - U_c)}{n}$, there is no Nash equilibrium. However, we could get the optimal approximate Nash equilibrium when all defenders use the same coverage probability $\frac{T - U_u}{kc}$ for all targets. The corresponding Social Welfare is

$$SW_E = (U_c - U_u - nkc) \frac{T - U_u}{kc} + U_u + (nk - 1)T,$$

and the $\frac{(T - U_u)(kc - U_c + U_u)}{cnk}$ -Price of Anarchy is

$$\begin{aligned} &\frac{(T - U_u)(kc - U_c + U_u)}{cnk} \cdot PoA \\ &= \frac{(U_c - U_u - nkc)(T - U_u)}{kcU_u + (nk - 1)kcT} + 1 \end{aligned}$$

We now analyze the relationship between (ϵ -)PoA and the values of n and k . Here are the key differences from Multi-Target Model. First we consider (ϵ -)PoA as the function of n . If $T = 0$, then we could find that the result is same as that in Multi-Target Model and (ϵ -)PoA linearly increases in n , which is unbounded. However, if $T \neq 0$, PoA and ϵ -PoA are increasing in n , and as $n \rightarrow \infty$, approach $1 - \frac{c}{T}$ and $1 + \frac{U_u - T}{kT}$, respectively. In other words, PoA (exact and approximate) is bounded by a constant, for a constant k !

Consider now approximate price of anarchy as a function of k . If $T = 0$, it is bounded by $n + 1$. However, if $T \neq 0$,




Figure 4: (Approximate) Price of Anarchy when $c = 1$, $T = -1$, $U_c = -2$ and $U_u = -10$

when $kc - \frac{(n-1)(T-U_c)}{n} \leq U_c - U_u$, it is an increasing function of k . When $kc - \frac{(n-1)(T-U_c)}{n} > U_c - U_u$, it may at first increase or decrease in k , depending on the values of the model parameters. However, when k is large enough, price of anarchy will invariably be decreasing in k , and as $k \rightarrow \infty$, ϵ -PoA $\rightarrow 1$. Figure 4 provides an example of the relationship between ϵ -PoA and k . Observe that all the curves begin to decrease when $k > 10$; they all approach 1 as $k \rightarrow \infty$. Thus, price of anarchy in the general model is only unbounded in the special case when $T = 0$, whereas when $T \neq 0$, price of anarchy is always bounded by a constant. This observation is particularly surprising and significant considering the fact that the baseline and simplified multi-target models are quite natural, and seemingly innocuous, restrictions of the general case.

6 Conclusion

We examined a non-cooperative multi-defender security game in which defenders may protect multiple targets, offering complete characterization of Nash and approximate equilibria, socially optimal solutions, and price of anarchy. Our results show that defenders generally over-protect the targets in this model, but different modeling assumptions give rise to qualitatively different outcomes: a simpler model gives rise to an unbounded price of anarchy, whereas a more general model sees PoA converge to a constant when the number of defenders increases.

There are a number of directions for further work. Firstly, in our work we assume that strategy spaces of defenders are symmetric, and the impact of asymmetry is far from clear. Secondly, defenders in our model always have incentive to over-invest for their targets. Some public policies such as taxation and penalties may be helpful to improve the overall efficiency. Finally, it would be important to consider how interdependence among targets affects the levels of security investment and the price of anarchy.

Acknowledgments

This work was partially supported by the Air Force Research Laboratory under award FA8785-14-2-0180, National Sci-

ence Foundation under award CNS-1238959, and Sandia National Laboratories.

References

- [Bachrach *et al.*, 2013] Yoram Bachrach, Moez Draief, and Sanjeev Goyal. Contagion and observability in security domains. Allerton Conference, 2013.
- [Chan *et al.*, 2012] Hau Chan, Michael Ceyko, and Luis E. Ortiz. Interdependent defense games: Modeling interdependent security under deliberate attacks. In Nando de Freitas and Kevin P. Murphy, editors, *UAI*, pages 152–162. AUAI Press, 2012.
- [Conitzer and Sandholm, 2006] Vincent Conitzer and Tuomas Sandholm. Computing the optimal strategy to commit to. In *Proceedings of the 7th ACM Conference on Electronic Commerce, EC '06*, pages 82–90, New York, NY, USA, 2006. ACM.
- [Heal and Kunreuther, 2002] Geoffrey Heal and Howard Kunreuther. Interdependent security. *Journal of Risk and Uncertainty*, 26:231–249, 2002.
- [Jain *et al.*, 2008] Manish Jain, James Pita, Milind Tambe, Fernando Ordóñez, Praveen Paruchuri, and Sarit Kraus. Bayesian stackelberg games and their application for security at los angeles international airport. *SIGecom Exch.*, 7(2):10:1–10:3, June 2008.
- [Jain *et al.*, 2010a] Manish Jain, Erim Kardes, Christopher Kiekintveld, Fernando Ordez, and Milind Tambe. Security games with arbitrary schedules: A branch and price approach. In Maria Fox and David Poole, editors, *AAAI*. AAAI Press, 2010.
- [Jain *et al.*, 2010b] Manish Jain, Jason Tsai, James Pita, Christopher Kiekintveld, Shyamsunder Rathi, Milind Tambe, and Fernando Ordóñez. Software assistants for randomized patrol planning for the lax airport police and the federal air marshal service. *Interfaces*, 40(4):267–290, July 2010.
- [Jiang *et al.*, 2013] Albert Xin Jiang, Ariel D. Procaccia, Yundi Qian, Nisarg Shah, and Milind Tambe. Defender (mis)coordination in security games. In *Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence, IJCAI '13*, pages 220–226. AAAI Press, 2013.
- [Kiekintveld *et al.*, 2009] Christopher Kiekintveld, Manish Jain, Jason Tsai, James Pita, Fernando Ordóñez, and Milind Tambe. Computing optimal randomized resource allocations for massive security games. In *Proceedings of The 8th International Conference on Autonomous Agents and Multiagent Systems - Volume 1, AAMAS '09*, pages 689–696, Richland, SC, 2009. International Foundation for Autonomous Agents and Multiagent Systems.
- [Pita *et al.*, 2009] James Pita, Manish Jain, Fernando Ordóñez, Christopher Portway, Milind Tambe, Craig Western, Praveen Paruchuri, and Sarit Kraus. Using game theory for los angeles airport security. *AI Magazine*, 30(1):43–57, 2009.
- [Shieh *et al.*, 2012] Eric Shieh, Bo An, Rong Yang, Milind Tambe, Craig Baldwin, Joseph DiRenzo, Ben Maule, and Garrett Meyer. 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 Multiagent Systems - Volume 1, AAMAS '12*, pages 13–20, Richland, SC, 2012. International Foundation for Autonomous Agents and Multiagent Systems.
- [Smith *et al.*, 2014] Andrew Smith, Yevgeniy Vorobeychik, and Joshua Letchford. Multi-defender security games on networks. *ACM SIGMETRICS Performance Evaluation Review*, 41(4):4–7, 2014.