# Multidefender Security Games

Jian Lou, Andrew M. Smith, *Student Member, IEEE,* and Yevgeniy Vorobeychik, *Member, IEEE*

***Abstract*—Current Stackelberg security game models primarily focus on isolated systems where only one defender is present, despite being part of a more complex system with multiple players. However, many real systems such as transportation networks and the power grid exhibit interdependencies between targets and, consequently, between decision makers jointly charged with protecting them. To understand such multidefender strategic interactions present in security, we investigate security games with multiple defenders. Unlike most prior analysis, we focus on situations in which each defender must protect multiple targets, so even a single defenders best response decision is, in general, non-trivial. We start with an analytical investigation in a special case of multidefender security games with independent targets, offer an equilibrium and price-of-anarchy analysis, and show that defenders have an incentive to over-protect the targets. Considering interdependencies among targets, we develop a novel mixed-integer linear programming formulation to compute a defenders best response, and approximate Nash equilibria of the game using this formulation. We apply this approach towards computational strategic analysis of several network models representing interdependencies, including real-world power networks. Our analysis shows how network structure and the probability of failure spread determine the propensity of defenders to over- or under-invest in security.**

## I. Introduction

Security, physical and cyber, has come to the forefront of national attention, particularly after 9/11. Among the variety of approaches that are used to tackle security problems, from risk analysis to red teaming, the Stackelberg Security game model [1] has had a significant impact, with tools based on game theoretic analysis having been deployed in LAX airport to schedule canine patrols [2], by Federal Air Marshall Service (FAMS) to schedule the air marshals [3], and by the US Coast Guard to schedule boat patrols [4].

A crucial assumption that all these efforts have in common is 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. In practice, numerous parties are responsible for security; indeed, the fact that the basic framework has been deployed by different entities and agencies makes this manifest already. If security decisions made by different parties were entirely independent, both from the defender's and the attacker's perspective, a single-defender model would be entirely satisfactory. However, the assets protected by different entities are typically interdependent, or, more generally, have value to others who are not involved in security decisions. Additionally, attackers, insofar as they may target different sectors under the charge of different defenders, are resource constrained, implicitly coupling otherwise independent targets.

An important motivating application for our multidefender security game is security and reliability in the power grid. Independent System Operators (ISOs) and profit-driven independent utility operators are largely responsible for operating and controlling subsystems of the entire grid [5]. These operators are held responsible for the reliability of their system, and thus have independent, and possibly even competing, goals with neighboring ISOs. As such, their security decisions are made independently, despite the interdependencies present between subsystems. As a result of this organization, cascading failures in the power grid can present a great threat to the entire system, particularly when subsystems are under attack.

In this paper, we investigate a Stackelberg game model, in which there are multiple defenders charged with protecting disjoint subsets of targets, which may be interdependent (for example, failures at one target may cascade to another). First, we examine a case where the values of the targets are *independent* and *homogeneous* among the defenders, and provide equilibrium and price-of-anarchy analysis. Specifically, 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, 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) and corresponding (approximate) Price of Anarchy, showing that over-investment is substantial in this case as well.

For the general case in which targets are interdependent, we propose a novel mixed-integer linear programming approach for computing a defender's best response, combined with a novel heuristic method to approximate equilibrium behavior. Interdependent multidefender games feature two competing externalities of security decisions: a *positive* externality, where greater security implies reduced contagion risk to other defenders, and a *negative* externality, which arises because high security investment by one defender incentivizes the attacker to attack someone else's assets. We study the impact of competing externality effects of defense on the resulting Nash equilibrium outcomes as a function of network topology (using both synthetic and real networks), interdependent risk, and the level of system decentralization. One of our key findings is that the impact of system decentralization on security and welfare can be non-monotonic when systems are highly

J. Lou and Y. Vorobeychik are with Vanderbilt University.

A. M. Smith is with Sandia National Laboratories and University of California, Davis.

interdependent: high levels of decentralization can yield near-optimal outcomes, even as moderate decentralization results in significant underinvestment. With weak interdependencies, on the other hand, an increasingly decentralized system tends more strongly to over-invest in security.

**Related Work**

Among the earliest multidefender models is in the literature on *interdependent security games* [7], in which interactions among multiple defenders are modeled as an $n$-player, 2-action game, where a player decides whether to invest in security; however, no attacker is considered. More recently, time-dependent scenarios where coordination of defender resources amongst multiple defenders is assumed have been studied using Markov decision processes [8]. Since total cooperation is assumed, this model effectively reduces to a single defender game in which the defender controls all resources. A recent extension, *interdependent defense games* [9], does consider an attacker who acts *simultaneously* with the defenders, rather than after observing the joint defense configuration, as in our model. Interdependent defense games have also been studied in the context of traffic infrastructure defense [10]. Two recent efforts studying multidefender games explicitly model interdependence among targets through a probabilistic contagion process [11], [12]. Like our paper, they consider attackers who observe the joint defense prior to making a decision, but each defender is restricted to secure a single node, and the strategy space is assumed to be continuous. [13] is, to our knowledge, the only other attempt to study strategic settings related to security in which each player's decision space is combinatorial. However, this work does not consider a strategic attacker.

## II. Multidefender Model

In the multidefender security game model, there are a collection of defenders $N = \{1, 2, 3, \ldots, n\}$, and a single attacker. A collection of targets $T$ will be protected by these defenders. Each defender $i$ is in charge of a set of targets $T_i$, such that $T_i \subseteq T$. We assume $T_i \cap T_{i'} = \emptyset$ when $i \neq i'$, and $\cup_{i \in N} T_i = T$.

**Strategies** Suppose that each defender $i$ can choose from a finite set $O = \{o_1, o_2, \ldots, o_{|O|}\}$ of security configurations for each target $t \in T_i$. A *pure strategy of defender* $i$ is $\mathbf{o}_i =< o_{i,t_{i1}}, o_{i,t_{i2}}, \ldots, o_{i,t_{ik}} \ldots, o_{i,t_{i|T_i|}} >$, in which $t_{ik}$ is the $k_{th}$ target of defender $i$, and $o_{i,j}$ (here $j = t_{i1}, t_{i2}, \ldots$, etc.) means defender $i$'s security configuration on target $j$ s.t. $j \in T_i$. We assume the attacker is resource constrained and can only attack one target in the game. That is, a pure strategy of the attacker is $j$, s.t. $j \in T$.

A *Mixed Strategy of a defender* $i$ is a matrix

$$\mathbf{q}_i = \begin{pmatrix} q^{o_1}_{i,t_{i1}} & q^{o_1}_{i,t_{i2}} & \cdots & q^{o_1}_{i,t_{i|T_i|}} \\ q^{o_2}_{i,t_{i1}} & q^{o_2}_{i,t_{i2}} & \cdots & q^{o_2}_{i,t_{i|T_i|}} \\ \vdots & \vdots & \ddots & \vdots \\ q^{o_{|O|}}_{i,t_{i1}} & q^{o_{|O|}}_{i,t_{i2}} & \cdots & q^{o_{|O|}}_{i,t_{i|T_i|}} \end{pmatrix}$$

In which, $q^o_{i,j}$ (here $o = o_1, o_2, \ldots, o_{|O|}$ and $j = t_{i1}, t_{i2}, \ldots, t_{i|T_i|}$) is the probability that the defender $i$ chooses $o$ at target $j$, and $\sum_{o \in O} q^o_{i,j} = 1$.

In our model, we assume a single strategic attacker that observes the defenders' coverage probabilities and chooses a target that maximizes its utility. A mixed strategy of the attacker can be denoted by $\mathbf{p} =< p_{t_1}, p_{t_2}, \ldots, p_{t_k} \ldots, p_{t_{|T|}} >$, in which, $t_k$ is the $k_{th}$ target in target set $T$, and $p_j$ (here $j = t_1, t_2, \ldots, t_{|T|}$) is the probability of attacking target $j \in T$.

We let $\mathbf{q} =< \mathbf{q}_1, \mathbf{q}_2, \ldots, \mathbf{q}_n >$ denote the strategy profile of the defenders, and $(\mathbf{q}, \mathbf{p})$ denote the strategy profile of the defenders and the attacker.

**Payoffs** A configuration $o \in O$ for target $j \in T_i$ incurs a cost $c^o_j$ to the defender $i$. If the attacker attacks a target $j \in T$ while configuration $o$ is in place, the expected value to a defender $i$ is denoted by $U^o_{i,j}$, while the attacker's value is $V^o_j$. We assume in this model that each player's utility depends only on the target attacked and its security configuration [3], [14].

**Solution Concepts** Traditionally, in single defender Stackelberg security games, the solution concept used is Strong Stackelberg Equilibrium (SSE). A SSE is characterized by an assumption that the attacker breaks ties in defender's favor. However the notion of "breaking ties in defender's favor" is no longer well defined when there are multiple defenders, as we must specify *which* defender will receive the favor. In our paper, we adopt a natural tie-breaking rule in which the attacker chooses a target uniformly at random from the set of all best responses. We call the corresponding solution concept (which is a refinement of the subgame perfect equilibrium of our game) the *Average-case Stackelberg Equilibrium (ASE)*.

**Definition 1.** *(Average-case Stackelberg Equilibrium) A strategy profile $(\mathbf{q}, \mathbf{p})$ is ASE if each defender's strategy is a best response, taking other defenders' strategies as given and assuming that the attacker will always play a best-response strategy, breaking ties uniformly at random if there are multiple best-response strategies.*

As we demonstrate below, ASE is not guaranteed to exist, in which case we focus on $\epsilon$-ASE (a refinement of $\epsilon$-equilibrium), in which no defender gains more than $\epsilon$ by deviating; in particular, we will consider $\epsilon$-ASE with the smallest attainable $\epsilon$.

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

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

where $SW_O$ is the optimal (utilitarian) social welfare that can be obtained (i.e., if there was a single defender), and $\epsilon$-$SW_E$ is the worst-case (utilitarian) social welfare in $\epsilon$-ASE. An underlying assumption of this definition is that the value of $SW_O$ and $\epsilon$-$SW_E$ are both positive. If they are both negative, then $\epsilon$-PoA will be the reciprocal of above equation. Note that the ordinary *Price of Anarchy* is a special case of $\epsilon$-Price of Anarchy with $\epsilon = 0$.

## III. Equilibrium Analysis of Independent Multidefender Security Games

Before investigating the multidefender security game generally, we first consider a special case of the model to reveal some useful insights. We consider scenarios in which the values of the targets are *independent* and *homogeneous* among the defenders. Our equilibrium and Price of Anarchy analysis will show that a Nash equilibrium among defenders in the Stackelberg game model (equivalently, ASE) [1] need not always exist, even when the defenders utilize randomized strategies (i.e., probability distributions over target protection levels). For cases when there is no Nash equilibrium, we make use of approximate Nash (ASE) equilibrium and the associated ($\epsilon$)-Price of Anarchy.

### A. Problem Setting

In this model, we assume all targets in $T$ are homogeneous, and each target has the same value to the defender. In the game model, each defender protects $k$ targets, i.e. $|T_1| = |T_2| = \ldots = |T_n| = k$. The security configuration space is $O = \{0, 1\}$, i.e., the defender's decision is binary. For example, 1 can correspond to the decision to protect an asset, while the configuration 0 would leave the asset unprotected. The *pure strategy* of defender $i$ is $\mathbf{o}_i =< o_{i,t_{i1}}, o_{i,t_{i2}}, \ldots, o_{i,t_{ik}} >$, in which $o_{i,j}$ (here $j = t_{i1}, t_{i2}, \ldots$) is a binary value. The mixed strategy of a defender $i$ is $\mathbf{q}_i =< q_{i,t_{i1}}, q_{i,t_{i2}}, \ldots q_{i,t_{ik}} >$, in which $q_{i,j}$ is the probability of protecting target $j$ for defender $i$ (coverage probability). The cost to defend each target is denoted by $c$.

If the attacker chooses to attack a target controlled by defender $i$ and the defender chooses to protect the target, we define the value of the target to defender $i$ to be $U^c$, and if the attacker attacks the target but it is not protected, then the value of the target to 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, the value of the target for defender $i$ is $\Omega \geq U^c$. In this setting, we assume that the attacker aims to maximize expected damage to the defender, so that the attacker's utility is $-U^u$, $-U^c$, and $-\Omega$ for the three outcomes above, respectively. Since these values are uniform across targets, equivalently the attacker attacks a target with lowest coverage probability (breaking ties uniformly at random).

### B. Equilibrium Analysis and Price of Anarchy

Our first result presents necessary and sufficient conditions for the existence of a Nash equilibrium among defenders (ASE) in the independent multidefender setting, and characterizes it when it does exist.

**Theorem 1.** *In the* Independent Multidefender setting*, Nash equilibrium among defenders (ASE) exists* if and only if $U^c - U^u \geq kc - \frac{(n-1)(\Omega - U^c)}{n}$. *In this equilibrium all targets are protected with probability 1.*

[1]If we treat attacker as externality, we could see an ASE as a Nash equilibrium among defenders. For ease of exposition, we will also use "Nash equilibrium among defenders" to denote ASE in the paper.

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

**Theorem 2.** *In* Independent Multidefender setting*, in the optimal $\epsilon$-equilibrium ($\epsilon$-ASE) all targets are protected with probability $\frac{\Omega - U^u}{kc}$. The corresponding $\epsilon$ is $\frac{(\Omega - U^u)(kc - U^c + U^u)}{cnk}$.*

Armed with a complete characterization of predictions of strategic behavior among the defenders, we can now consider how this behavior is 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.

**Theorem 3.** *In the* Independent Multidefender setting*, the optimal social welfare $SW_O$ is*

$$SW_O = \begin{cases} U^c - nkc + (nk-1)\Omega, & \text{if } U^c - U^u \geq nkc; \\ U^u + (n-1)\Omega, & \text{if } U^c - U^u < nkc. \end{cases}$$

*Proof sketch.* First, we claim that we could get optimal social welfare *only if* all targets have the same coverage probability $q$. Otherwise, some target $j$, which is influenced by defender $i$ has probability 0 of being attacked, and we can decrease $q_{i,j}$ to improve social welfare. Consequently, we need only to consider an optimal symmetric coverage probability $q$ to maximize social welfare, which can be done in a relatively straightforward way. □

If $U^c - U^u \geq kc - \frac{(n-1)(\Omega - U^c)}{n}$, the Nash equilibrium is unique, with all targets protected with probability 1. The corresponding social welfare is

$$SW_E = U^c - nkc + (nk-1)\Omega.$$

So far we have not yet added any constrains to value of $\Omega$, $U^c$, and $U^u$ (except that $\Omega \geq U^c \geq U^u$). In order to make *Price of Anarchy* well-defined, we need to add constraints that values of $\Omega$, $U^c$, and $U^u$ are all non-positive or all non-negative. We add constraints that $U^c$, $U^u$ and $\Omega$ are all non-positive (little changes if all are non-negative).

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

$$PoA = \begin{cases} 1, & \text{if } U^c - U^u \geq nkc; \\ \frac{U^c - U^u - nkc}{U^u + (nk-1)\Omega} + 1, & \text{if } kc - \frac{(n-1)(\Omega - U^c)}{n} \leq \\ & U^c - U^u < nkc. \end{cases}$$

If $U^c - U^u < kc - \frac{(n-1)(\Omega - U^c)}{n}$, there is no Nash equilibrium. The Social Welfare in the optimal approximate equilibrium is

$$\epsilon\text{-}SW_E = (U^c - U^u - nkc)\frac{\Omega - U^u}{kc} + U^u + (nk-1)\Omega,$$

and the $\frac{(\Omega - U^u)(kc - U^c + U^u)}{cnk}$-Price of Anarchy is $\frac{(U^c - U^u - nkc)(\Omega - U^u)}{kcU^u + (nk-1)kc\Omega} + 1$.

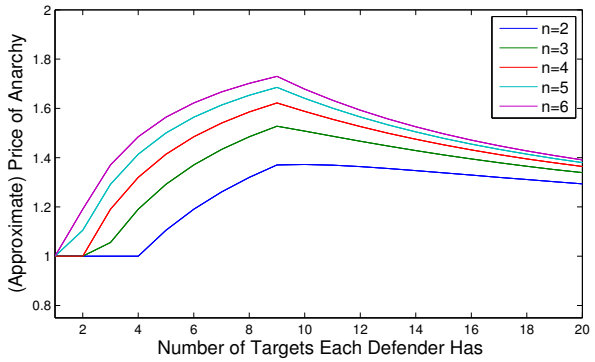


Fig. 1. (Approximate) Price of Anarchy when $c = 1, \Omega = -1, U^c = -2$ and $U^u = -10$

From this result, it is already clear that defenders systematically over-invest in security. 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.

We now analyze the relationship between ($\epsilon$-)PoA and the values of $n$ and $k$. First we consider ($\epsilon$-)PoA as the function of $n$. If $\Omega = 0$, ($\epsilon$-)PoA linearly increases in $n$, and is therefore unbounded. However, if $\Omega \neq 0$, while PoA and $\epsilon$-PoA are increasing in $n$, as $n \to \infty$, they approach $1 - \frac{c}{\Omega}$ and $1 + \frac{U^u - \Omega}{k\Omega}$, 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 $\Omega = 0$, it is bounded by $n+1$. However, if $\Omega \neq 0$, when $kc - \frac{(n-1)(\Omega - U^c)}{n} \leq U^c - U^u$, it is an increasing function of $k$. When $kc - \frac{(n-1)(\Omega - U^c)}{n} > U^c - U^u$, it may at first increase or decrease in $k$, depending on the the values of the model parameters. However, when $k$ is large enough, price of anarchy will invariably be decreasing in $k$, and as $k \to \infty$, $\epsilon$-PoA $\to 1$. Figure 1 provides an example of the relationship between $\epsilon$-PoA and $k$. Observe that all the curves begin to decrease when $k > 10$, and they all approach 1 as $k \to \infty$. Thus, price of anarchy in the independent multidefender setting is only unbounded in the special case when $\Omega = 0$, whereas when $\Omega \neq 0$, price of anarchy is always bounded by a constant. This observation is particularly surprising considering the fact that the $\Omega = 0$ is a natural and seemingly innocuous restriction of the general case.

## IV. Computing Average-case Stackelberg Equilibrium

We now develop and analyze a computational framework for approximating Nash equilibria in multidefender security games in the general case (which can include *interdependencies* between targets). A crucial step in computing (or approximating) a Nash equilibrium of a game is to consider the problem of computing a best response for an arbitrary player (in our case, defender, since the attacker's best response is straightforward). Next, we develop a novel mixed-integer linear programming formulation for computing ASE best response, and then propose an effective heuristic method for approximating ASE in multidefender games.

### A. Computing Defender Best Response: A Mixed-Integer Linear Programming Formulation

While ASE seems a natural alternative to SSE in multiplayer security games, we are not aware of any proposals for computing it. Below, in equations 1-11, we present the first (to our knowledge) mixed-integer linear programming formulation for computing ASE. The solution to the MILP below is a best response for an arbitrary defender $i$ when the strategies of all other players, $q_{-i}$, are fixed.

$$\max_{a,q_i,s,u,v} u - \sum_{j \in T_i} \sum_{o \in O} c_j^o q_{i,j}^o \tag{1}$$

s.t.

$$0 \leq q_{i,j}^o \leq 1 \qquad \forall\ j \in T_i,\ \forall o \tag{2}$$

$$\sum_{o \in O} q_{i,j}^o = 1 \qquad \forall j \in T_i \tag{3}$$

$$a_j \in \{0,1\} \qquad \forall\ j \in T \tag{4}$$

$$\sum_{j \in T} a_j \geq 1 \tag{5}$$

$$0 \leq v - \sum_o q_{i,j}^o V_j^o \leq (1 - a_j)M \qquad \forall\ j \in T_i \tag{6}$$

$$0 \leq v - \sum_o q_{-i,j}^o V_j^o \leq (1 - a_j)M \qquad \forall\ j \in T_{-i} \tag{7}$$

$$s_j = v - \sum_o q_{i,j}^o V_j^o \qquad \forall\ j \in T_i \tag{8}$$

$$s_j = v - \sum_o q_{-i,j}^o V_j^o \qquad \forall\ j \in T_{-i} \tag{9}$$

$$a_j + M s_j \geq 1 \qquad \forall\ j \in T \tag{10}$$

$$u = f(q,a), \tag{11}$$

where $M$ is a very large number, $a_j = 1$ if the attacker chooses to attack target $j$, and

$$f(q,a) = \frac{1}{\sum_{j \in T} a_j} \left( \sum_{j \in T_i} a_j \sum_{o \in O} q_{i,j}^o U_{i,j}^o + \sum_{j \in T_{-i}} a_j \sum_{o \in O} q_{-i,j}^o U_{i,j}^o \right).$$

While constraint 11 is non-linear, we can linearize it using McCormick inequalities. Constraints 2 and 3 ensure that the defender's strategy is a valid probability distribution. Constraint 5 ensures that at least one target is chosen by the attacker, since ASE allows for the attacker to choose uniformly at random from a set of optimal targets to attack. Constraints 6 and 7 compute the optimal attacker utility $v$; alone, they ensure that this utility corresponds to *some* attack target. For example, if $a_j = 1$ for some target $j \in T_i$, this forces the difference between the optimal attacker value $v$ and the attacker value at the target $j$ to be 0 (this is an optimal target to attack). Constraints 8 and 9 compute an auxiliary variable $s_j$, which

is 0 if and only if attacking a target $j$ yields an optimal utility to the attacker. These variables, together with constraints 10 and 6-7 ensure that the binary variable $a_j = 1$ if and only if the attacker (weakly) prefers to attack target $j$; that is, these jointly compute the set of optimal attack targets. Finally, constraint 11 computes the expected utility to the defender if the attacker chooses one of his most preferred targets uniformly at random.

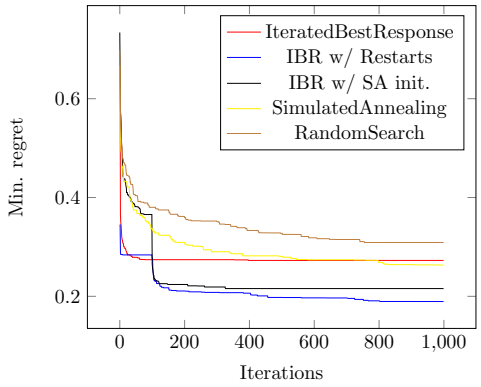


Fig. 2. Comparison of NE approximation algorithms ($|N| = 5$ and $|T| = 20$, $c = 0.2$.)

### B. Approximating ASE

Scenarios with a large set of targets and/or interdependencies, which arise frequently in realistic complex systems, may be intractable to compute the ASE exactly. As a result, formal equilibrium analysis for general scenarios with interdependencies is not feasible. Previously, [16] presented a convergent equilibrium approximation algorithm based on *simulated annealing (SA)* that would be applicable in our setting. They additionally showed in simulation that SA is outperformed by a heuristic based on *iterated best response (IBR)* dynamics. We interpret IBR as a local search heuristic, with the property that if the starting point is a Nash equilibrium, IBR will never deviate from it (i.e., Nash equilibrium is a fixed point). Clearly, then, the choice of a starting point can be significant for the performance of IBR, making it natural to consider coupling it with random restarts. Our contribution in this section is to present evidence that IBR with random restarts is an effective equilibrium approximation approach in our setting (and outperforms several alternatives). We use this algorithm for our analyses below.

We compare the following Nash equilibrium approximation algorithms executed for 1000 iterations: random search (RS), which generates 1000 strategy profiles randomly, computes the game theoretic regret of each, and chooses a profile with the smallest regret; simulated annealing (SA), with the temperature exponentially increasing with iterations; and iterated best response (IBR). We also include two additional variations of IBR: the first uses SA for the first 100 iterations, and then switches to IBR for the remainder (starting with the best approximation produced by SA); the second is IBR with random restarts (RIBR). RIBR includes initial corner cases that may be hard to converge to in a limited amount of time (i.e., all defenders not defending, all defenders defending completely). We execute our comparison on games with 2 players and 10 targets and games with 5 players and 20 independent, randomly valued targets. We found that RIBR outperforms other alternatives in both settings. Figure 2 shows the comparison visually for the setting with 5 players and 20 targets.

## V. Analysis of Interdependent Multidefender Security Games

In many practical applications, a defender may have an indirect valuation of targets not under their control as a result of functional dependence, or the possibility of a failure cascading from another defender's target(s) to a target of explicit value. Interdependence can also be present within a defender's controlled targets, which may reshape the original valuations in nontrivial ways. As a result, we analyze interdependent multidefender games empirically to gain insight into the strategic behavior of defenders in realistic, complex systems.

To model interdependencies, we construct a graph $(T, E)$ with $T$ the set of targets (nodes) as above, and $E$ the set of edges $(j, j')$, where an edge from $j$ to $j'$ means that a successful attack on $j$ may have impact on $j'$. Each target $j$ has associated with it a value, $v_{ij}$, for the defender $i$, which is the loss to $i$ if $j$ is affected (e.g., compromised, broken). The security configuration determines the probability $z_j^o(j)$ that target $j$ is affected if the attacker attacks it *directly* and the defense configuration is $o$. We model the interdependencies between the nodes as independent cascade contagion [14], [15]. The contagion proceeds starting at an attacked node $j$, affecting its network neighbors $j'$ each with probability $r_{j,j'}$, the contagion then spreads from the newly affected nodes $j'$ to their neighbors, and so on. The contagion can only occur one time along any network edge, and once a node is affected it stays affected through the diffusion process. Each player's valuation for each target is then updated based on the probability of a failure cascading to one of the player's owned targets. Figure 3 illustrates a sample scenario with interdependencies, with defenders' utilities before and after the interdependent cascade contagion process. Initially each defender only values their own target; however, after observing that an attack's consequences can be transmitted from another previously unvalued target to their target with probability $r$, the defender's now also have a non-zero value for the target outside of their control. Note that after the contagion process, even though Defender 1 has completely covered their node, there is still negative utility associated with Defender 2 being attacked as a result of the interdependency.

In this model, the defender $i$'s utility $U_{i,j}^o = \mathbb{E}[z_j^o(j) \sum_k (-v_{ik}) \rho_k(j)]$, where $\rho_k(j)$ is the probability that a target $k$ is affected by the contagion process if the attacked target $j$ is successfully compromised. In the experiments below, we restrict $O$ to be binary (corresponding to coverage decisions), and let $z_j^o(j) = 0$ when the target is covered when attacked, and $z_j^o(j) = 1$ if it is not covered, for all targets $j$. We also let $V_j^o = -\sum_i U_{ij}^o$.

We now (approximately) compute ASE for several synthetic classes of interdependency networks popular in network sci-

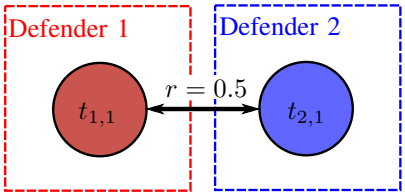


| | | |
|---|---|---|
| **Target Values:** | $v_{1,t_{1,1}} = 1$ | $v_{2,t_{1,1}} = 0$ |
| | $v_{1,t_{2,1}} = 0$ | $v_{2,t_{2,1}} = 1$ |
| **Utilities before contagion:** | $U_{1,t_{1,1}} = 0$ | $U_{2,t_{1,1}} = 0$ |
| | $U_{1,t_{2,1}} = 0$ | $U_{2,t_{2,1}} = -1$ |
| **Utilities after contagion:** | $U_{1,t_{1,1}} = 0$ | $U_{2,t_{1,1}} = 0$ |
| | $U_{1,t_{2,1}} = -0.5$ | $U_{2,t_{2,1}} = -1$ |

Fig. 3. Simple interdependency graph with 2 defenders, 2 targets in $T = \{t_{1,1}, t_{2,1}\}$, and interdependency edges $E = \{(t_{1,1}, t_{2,1}), (t_{2,1}, t_{1,2})\}$. The probability of a failure spread between the targets is $r = r_{t_{1,1},t_{2,1}} = r_{t_{2,1},t_{1,1}} = 0.5$. Defenders' values ($v_{ik}$) are shown for each target $k$. Target $t_{1,1}$ is fully defended, while target $t_{2,1}$ is not defended, and is attacked. Utilities of both players are shown before contagion (if contagion does not occur) and after contagion (if one does occur).

ence and graph theory literature, and for networks derived from real power grid systems.

### A. Analysis of Multi-Defender Games on Synthetic Networks

For our first set of experiments, we use RIBR on 3 artificially generated networks, with 40 samples for each parameter variation. First, we will illustrate and compare the results of our interdependent multidefender game on artificial networks. We use 3 commonly analyzed network structures: a grid, Erdős-Rényi networks, and preferential attachment networks. In all of the generated networks, there are 64 nodes or targets. For the latter two, we use the Metis graph partitioning software to partition the nodes (targets) among defenders. This software partitions nodes to minimize connectivity among the targets belonging to different defenders, a property that we expect to commonly hold in real networks due to efficiency considerations.

We begin by considering average strategies and social welfare for the three different synthetic networks (grid, Erdős-Rényi, and preferential attachment), as a function of the number of players (degree of decentralization) and the cascade probability (interdependent risk). These results are shown in Figure 4. The first rather stark observation is that network structure makes little difference when each node is controlled by a single player, but it makes a significant qualitative difference both for social welfare and actual strategies utilized by the players in all other cases.

We first discuss social welfare in greater detail (shown in Figure 4, top). When interdependent risk is low ($0.1 \leq r \leq 0.3$), social welfare follows a relatively simple pattern: increasing decentralization makes initially almost no difference, until sufficiently many players are involved, at which point social welfare falls dramatically; this pattern is roughly monotonic with increasing decentralization, with worst outcomes emerging when each player controls a single node, and mirrors previous findings [13]. Both Erdős-Rényi and preferential attachment networks are less susceptible to the negative effects of decentralization in this case than the grid network, where the dropoff occurs with fewer players (less decentralization). This may be largely a consequence of the fact that network partitioning tools we use attempt to minimize interdependence among players—something that is likely to mirror reality—and far more opportunities for doing so exist in Erdős-Rényi and preferential attachment models.

When $r$ is higher (greater interdependencies), the results exhibit an entirely new phenomenology. Across all three network models, for sufficiently large $r$, the impact of decentralization is non-monotonic: an intermediate level of decentralization has the most detrimental impact on security, while a highly decentralized system becomes near-optimal!

Investigating actual (average) strategic decisions by the players yields deeper insights into the findings above. When interdependencies are weak, optimal decision is to invest relatively little in security, in any generative model. Increasing decentralization, therefore, gives rise primarily to over-investment, mirroring our analytical results for the limiting case when targets were independent. The tendency to over-invest, however, is quite weak until the network is extremely decentralized, except in the grid network. When $r$ is high, positive externalities prevail, and the predominant phenomenon is underinvestment. What is surprising is, again, non-monotonicity in the level of decentralization: when decentralization is moderate, underinvestment can be quite dramatic. On the other hand, a high level of decentralization often appears to dull this effect, and the level of investment in security becomes much closer to optimal.

### B. Results on power grid networks

Although Erdős-Rényi and preferential attachment models were developed in part to resemble real networks, the approximate equilibrium results applied to three snippets of actual power networks most resemble the phenomena observed for the grid, (shown in Figure 5). In particular, just as in the grid above, over-investment in security appears to dominate, even at relatively high levels of interdependence, but only when decentralization is significant. Most other levels of decentralization remain relatively near-optimal (these are, in fact, more robust to decentralization than the grid network above).

To dig somewhat deeper into the rather complex phenomenology we have observed, Figure 6 shows several examples of actual strategy realizations. First, consider the top series of plots for the grid with cascade probability $r = 0.1$. As previously described, we can clearly see that the optimal security configuration involves no security investment (leftmost grid), whereas an increasing level of decentralization gives rise to increased security investment, culminating, ultimately, with full protection in the extreme level of decentralization. The contrast between the two extremes offers some guidance: even though optimal global configuration involves no security, when each player controls (and values) only a single node, the best response of an attacked node is to defend it just enough to force the attacker to attack another; for example, slightly more than the next weakest node. Iterating on this

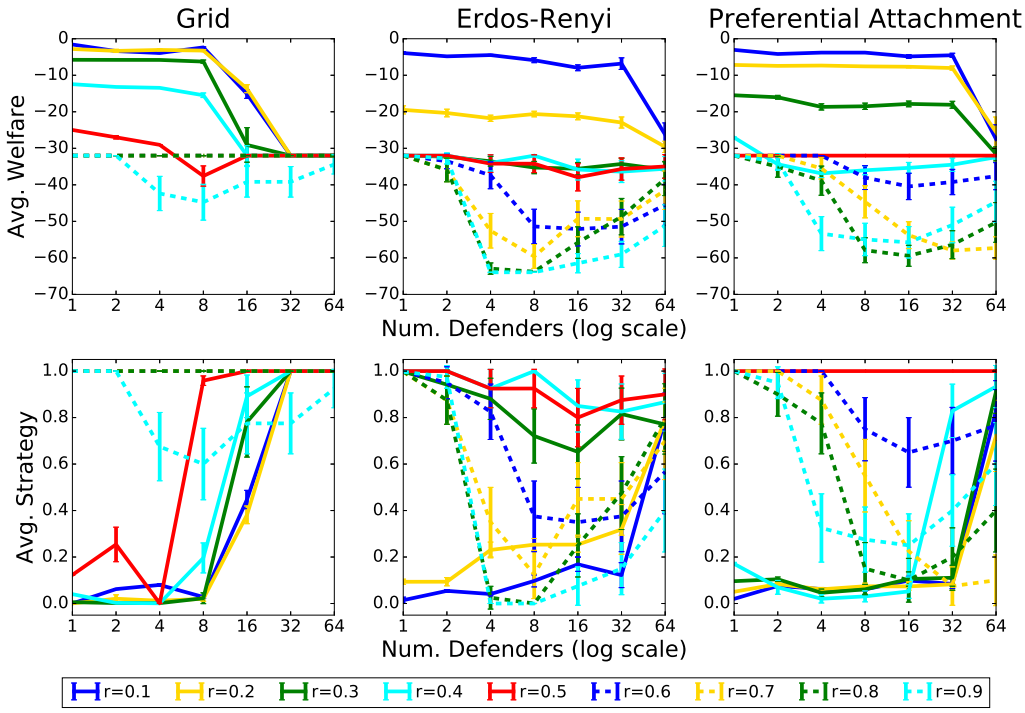


Fig. 4. Approximate equilibrium security outcomes for varying cascade probability $r$, as a function of (log of) the number of players (level of decentralization) for the three synthetic networks. Top: Social welfare. Bottom: Average strategy (higher strategy corresponds to higher average probability of defense).

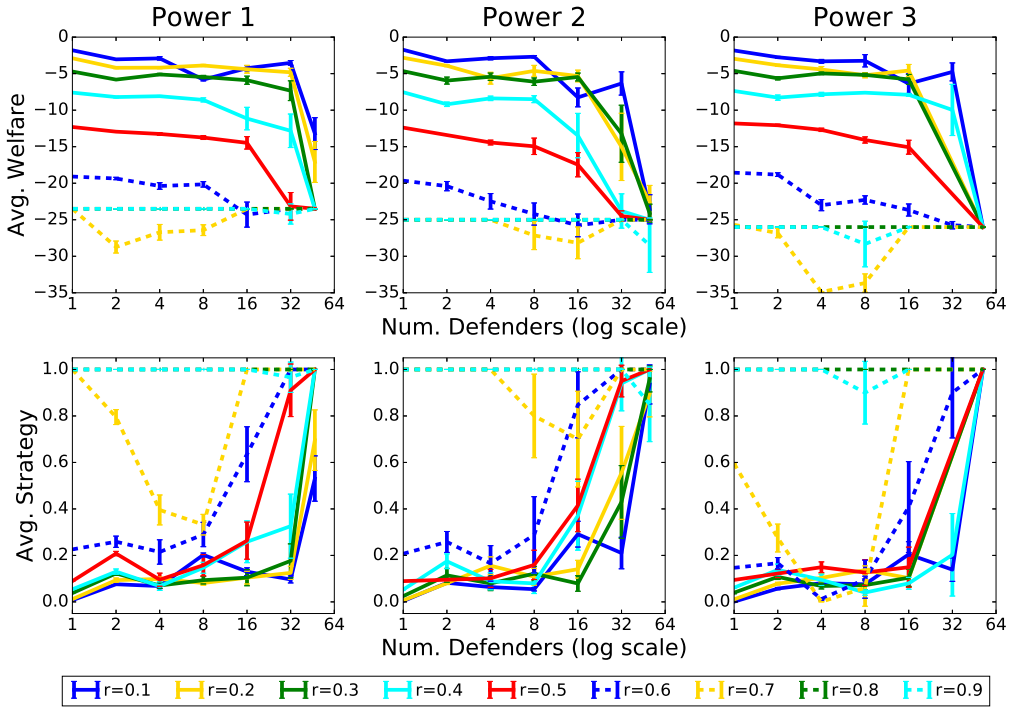


Fig. 5. Approximate equilibrium security outcomes for varying cascade probability $r$, as a function of (log of) the number of players (level of decentralization) for the three real power networks. Top: Social welfare. Bottom: Average strategy (higher strategy corresponds to higher average probability of defense).

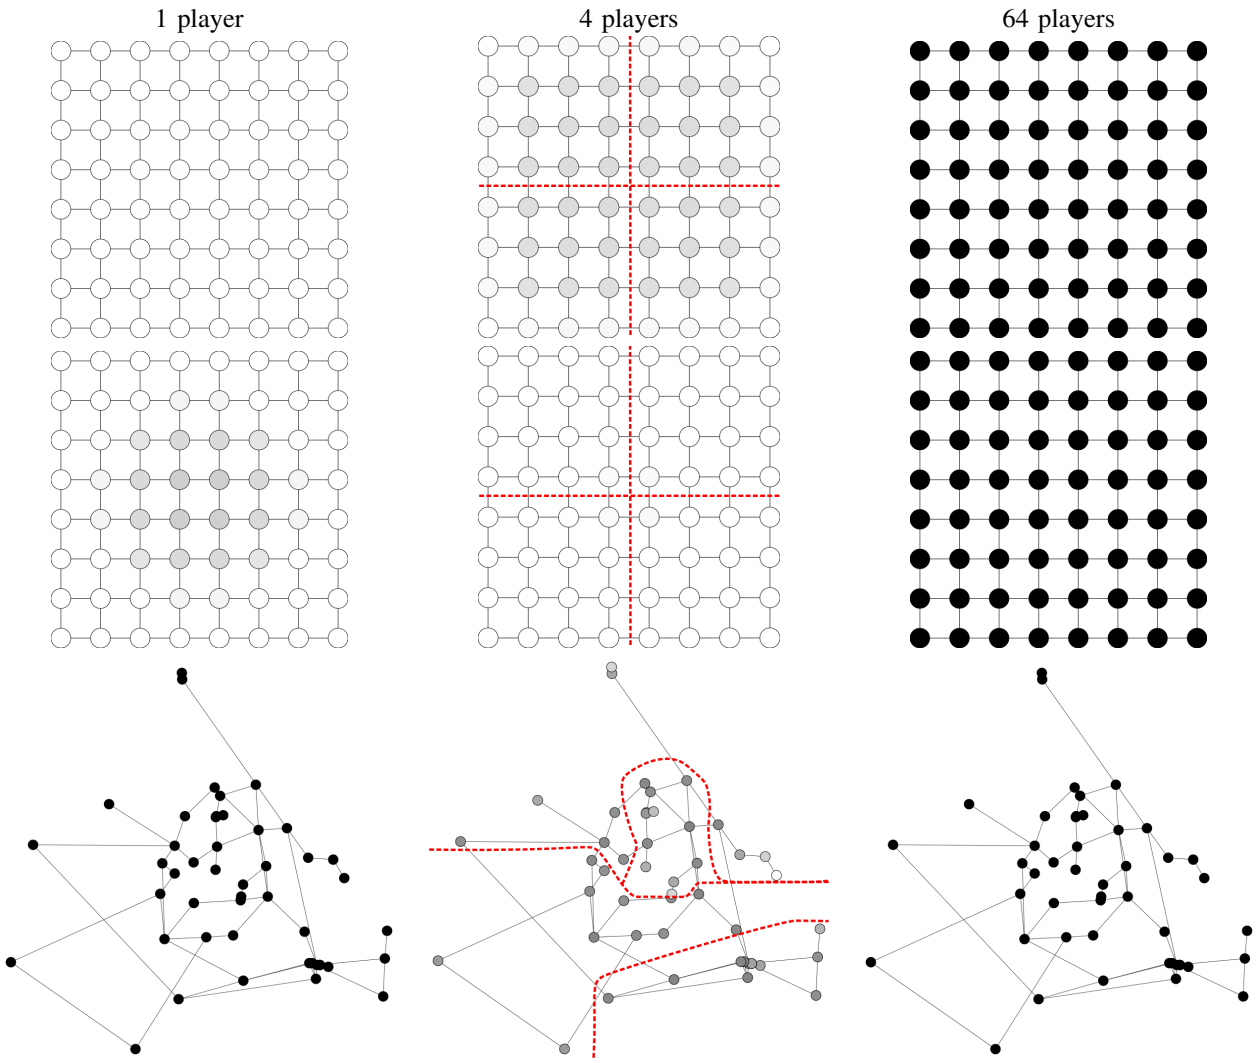


Fig. 6. Strategic realizations for representative games on grid and distribution network topologies. Top: Grid, $r = 0.1$, middle: Grid, $r = 0.4$, bottom: Power Network, $r = 0.7$. Darker node colors indicate higher probability of defense (coverage). Dotted red lines indicate the partition of nodes among players (except in the rightmost case, when each player controls a single node).

idea, strategies "cascade" to full defense. When the player controls more than one node, however, there is suddenly strategic tension: higher security on one node may well push the attacker to attack another node under this player's control. Positive externalities become more significant as well: pushing the attacker to attack another node "nearby" is likely to gain little when cascade probabilities are high and multiple nodes owned by the defender could be affected. For sufficiently high cascade probabilities, and sufficiently low number of players, such positive network effects can actually sway players to under-invest in security, as we can see both in the middle and last rows of Figure 6 (the 4-player case). Here, strategic complementarities make security investment not worthwhile in equilibrium: the nodes that need to be defended are relatively central, and cut across different players (i.e., the critical central nodes create a kind of "buffer" between defenders). This behavior diminishes as decentralization increases.

## VI. Conclusion

In this work, we have extended the current state of Stackelberg security games to include multiple defenders in non-cooperative scenarios with independent or interdependent targets. For the independent case with homogeneous targets, we provided complete characterizations of Nash and approximate equilibria, socially optimal solutions, and price of anarchy (PoA). Our analysis showed that defenders generally overprotect the targets. For the interdependent case, we developed a novel computational framework to overcome the difficulties of providing a concise formal analysis of such a complex model. One of our most stark findings is the non-monotonicity of welfare and strategic choices as a function of the number of players: in a number of cases, higher levels of decentralization become near-optimal, even while intermediate decentralization leads to very poor outcomes. Our findings enable a deeper understanding of practical security considerations, highlighting the importance of *both*, over- and under-investment in security, and the dependence of each on network structure, the magni-

tude of network externalities, and the level of decentralization. Finally, we have shown how security behavior in our model on real-world power networks relates to those in synthetic networks, highlighting similar behaviors with grid networks.

## REFERENCES


[1] V. Conitzer and T. Sandholm, "Computing the optimal strategy to commit to," in *Proceedings of the 7th ACM conference on Electronic commerce*, ser. EC '06. New York, NY, USA: ACM, 2006, pp. 82–90.

[2] P. Paruchuri, J. P. Pearce, J. Marecki, M. Tambe, F. Ordonez, and S. Kraus, "Playing games with security: an efficient exact algorithm for bayesian stackelberg games," in *Proceedings of the Seventh International Conference on Autonomous Agents and Multiagent Systems*, 2008, pp. 895–902.

[3] C. Kiekintveld, M. Jain, J. Tsai, J. Pita, F. Ordonez, and M. Tambe, "Computing optimal randomized resource allocations for massive security games," in *Proceedings of the Eighth International Conference on Autonomous Agents and Multiagent Systems*, 2009.

[4] E. Shieh, 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 *Proceedings of the Eleventh International Conference on Autonomous Agents and Multiagent Systems*, 2012, pp. 13–20.

[5] J. Lazar, "Electricity regulation in the US: A guide," Regulatory Assistance Project, Tech. Rep., 2011.

[6] J. Lou and Y. Vorobeychik, "Equilibrium analysis of multi-defender security games," in *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, IJCAI 2015, Buenos Aires, Argentina, July 25-31, 2015*, 2015, pp. 596–602.

[7] H. Kunreuther and G. Heal, "Interdependent security," *Journal of Risk and Uncertainty*, vol. 26, no. 2-3, pp. 231–249, 2003.

[8] E. Shieh, A. X. Jiang, A. Yadav, P. Varakantham, and M. Tambe, "Unleasing dec-mdps in security games: Enabling effective defender teamwork," *European Conference on Artificial Intelligence*, 2014.

[9] H. Chan, M. Ceyko, and L. E. Ortiz, "Interdependent defense games: Modeling interdependent security under deliberate attack," in *Twenty-Eighth Conference on Uncertainty in Artificial Intelligence*, 2012, pp. 152–162.

[10] D. L. Alderson, G. G. Brown, W. M. Carlyle, and R. K. Wood, "Solving defender-attacker-defender models for infrastructure defense," *INFORMS Computing Society Conference*, 2011.

[11] Y. Bachrach, M. Draief, and S. Goyal, "Contagion and observability in security domains," in *Allerton Conference*, 2013.

[12] D. Acemoglu, A. Malekian, and A. Ozdaglar, "Network security and contagion," 2013, working paper.

[13] Y. Vorobeychik, J. Mayo, R. Armstrong, and J. Ruthruff, "Noncooperatively optimized tolerance: Decentralized strategic optimization in complex systems," *Physical Review Letters*, vol. 107, no. 10, p. 108702, 2011.

[14] J. Letchford and Y. Vorobeychik, "Computing optimal security strategies for interdependent assets," in *Conference on Uncertainty in Artificial Intelligence*, 2012, pp. 459–468.

[15] D. Kempe, J. M. Kleinberg, and Éva Tardos, "Maximizing the spread of influence in a social network," in *Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2003, pp. 137–146.

[16] Y. Vorobeychik and M. P. Wellman, "Stochastic search methods for nash equilibrium approximation in simulation-based games," in *Seventh International Conference on Autonomous Agents and Multiagent Systems*, 2008, pp. 1055–1062.

## APPENDIX

**Theorem 1.** *In the* Independent Multidefender setting*, Nash equilibrium among defenders (ASE) exists* if and only if $U^c - U^u \geq kc - \frac{(n-1)(\Omega-U^c)}{n}$. *In this equilibrium all targets are protected with probability 1.*

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

When all targets have the same coverage probability $q$ to be protected, the utility of each defender is

$$u = \frac{(U^c - U^u - nkc)q + U^u + (nk-1)\Omega}{n}.$$

If $q < 1$, then some defender $i$ could increase $q$ to $q+\delta$ for all of her targets to ensure none of them are attacked, and obtain utility of $u' = k\Omega - k(q+\delta)c$, so that

$$u' - u = \frac{(U^c - U^u)(1-q) + (\Omega - U^c) - nkc\delta}{n}.$$

As $U^c \geq U^u$, $\Omega \geq U^c$, and $\delta$ can be arbitrarily small, $u' - u > 0$ when $q < 1$, which means that this cannot be a Nash equilibrium. Thus, the only possible equilibrium can be $q_{ij} = 1$ for all targets $t_{ij}$.

When all targets have the same coverage probability $q = 1$, each defender's utility is

$$u = \frac{U^c - nkc + (nk-1)\Omega}{n}.$$

We claim that if a defender $i$ has an incentive to deviate, it is optimal for this defender to use the same coverage probability for all her targets. Otherwise, for some target $t_{ik}$ which has probability 0 of being attacked, she could decrease $q'_{ik}$ to obtain higher utility. If probabilities of targets protected by defender $i$ are all $q'$ $(0 \leq q' < 1)$, then her expected utility is $u' = (U^c - U^u - c)q' + U^u + (k-1)(\Omega - q'c)$, and

$$u' - u = (U^c - U^u - kc)(q' - 1) + \frac{(n-1)(U^c - \Omega)}{n}.$$

We therefore have two cases:

1) If $U^c - U^u \geq kc$, then $u' - u \leq 0$, and $q = 1$ for all targets is a Nash equilibrium.
2) If $U^c - U^u < kc$, the maximal value of $u' - u$ corresponds to $q' = 0$:

$$\max_{0 \leq q' < 1} u' - u = -(U^c - U^u - kc) - \frac{(n-1)(\Omega - U^c)}{n}.$$

If $kc - \frac{(n-1)(\Omega-U^c)}{n} \leq U^c - U^u < kc$, $u' - u \leq 0$, it is a Nash equilibrium; otherwise, it is not.

To sum up, a Nash equilibrium exists *if and only if* $U^c - U^u \geq kc - \frac{(n-1)(\Omega-U^c)}{n}$, and the equilibrium corresponds to all targets having probability 1 of being protected. □

**Theorem 2.** *In* Independent Multidefender setting*, in the optimal $\epsilon$-equilibrium ($\epsilon$-ASE) all targets are protected with probability $\frac{\Omega-U^u}{kc}$. The corresponding $\epsilon$ is $\frac{(\Omega-U^u)(kc-U^c+U^u)}{cnk}$.*

*Proof.* When all targets have the same coverage probability $q$, the expected utility of each defender is

$$u = \frac{(U^c - U^u - nkc)q + U^u + (nk-1)\Omega}{n}.$$

Suppose $0 \leq q < 1$. If some defender $i$ increases $q$ to $q + \delta_{ij}$ for each of her target $t_{ij}$, then she would obtain utility $u' = \sum_{j=1}^{k} \Omega - (q + \delta_{ij})c$, and

$$\begin{aligned} u' - u &= \frac{\Omega - (U^c - U^u)q - U^u}{n} - \sum_{j=1}^{k} \delta_{ij} c \\ &\leq \frac{\Omega - (U^c - U^u)q - U^u}{n}. \end{aligned} \tag{12}$$

Now we consider scenarios in which a defender $i$ could obtain higher utility by decreasing protection probability. We claim that if a defender $i$ has an incentive to deviate, it is optimal for this defender to use the same coverage probability for all her targets. Otherwise, for some target $t_{ik}$ which has probability 0 of being attacked, she could decrease $q'_{ik}$ to obtain higher utility. Thus, we need only consider cases in which a defender deviates by decreasing coverage probabilities for all her targets to $q - \delta$. Her utility will become $u'' = (U^c - U^u - kc)(q - \delta) + U^u + (k-1)\Omega$. Since $U^c - U^u < kc$, $\delta = q$ (the maximal value of $\delta$) maximizes $u'' - u$:

$$\max_{0 < \delta \leq q} u'' - u = \frac{\Omega - (U^c - U^u)q - U^u}{nk} + kcq + U^u - \Omega. \tag{13}$$

By comparing the value of equation (12) and equation (13), we get different values of $\epsilon$ for $\epsilon$-equilibrium:

$$\epsilon = \begin{cases} \frac{\Omega-(U^c-U^u)q-U^u}{n}, & \text{if } 0 \leq q \leq \frac{\Omega-U^u}{kc}; \\ \frac{\Omega-(U^c-U^u)q-U^u}{n} + kcq + U^u - \Omega, & \text{if } \frac{\Omega-U^u}{kc} < q \leq 1. \end{cases}$$

When $q = \frac{\Omega-U^u}{kc}$, we get the minimal $\epsilon = \frac{(\Omega-U^u)(kc-U^c+U^u)}{cnk}$.

We claim that the $\frac{(\Omega-U^u)(kc-U^c+U^u)}{cnk}$-equilibrium can appear *only if* all targets have the same coverage probability $q$. 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); and 2) Some defender has different coverage probabilities for her targets. In case 1), there exist $\beta$ defenders $(1 \leq \beta < n)$ who have the same minimal coverage probability $q'$. The expected utility for each defender among these $\beta$ is

$$u = \frac{(U^c - U^u - k\beta c)q' + U^u + (k\beta - 1)\Omega}{\beta}.$$

When $\frac{\Omega-U^u}{kc} < q' \leq 1$, some defender $i$ among these $\beta$ could decrease the coverage probability of all her targets to 0 and obtain the utility of $u_1 = U^u + (k-1)\Omega$, so that

$$\begin{aligned} u_1 - u &= \frac{\Omega - (U^c - U^u)q' - U^u}{\beta} + kcq' + U^u - \Omega \\ &> \frac{\Omega - (U^c - U^u)q' - U^u}{n} + kcq' + U^u - \Omega. \end{aligned}$$

When $0 \leq q' \leq \frac{\Omega - U^u}{kc}$, some defender $i$ among these $\beta$ can increase coverage probabilities of all her targets to $q' + \delta_3$ to obtain utility of $u_2 = k\Omega - k(q' + \delta_3)c$, with

$$\begin{aligned} u_2 - u &= \frac{\Omega - (U^c - U^u)q' - U^u - k\beta c\delta_3}{\beta} \\ &> \frac{\Omega - (U^c - U^u)q' - U^u}{n}, \end{aligned}$$

where the inequality holds because $\delta_3$ can be arbitrarily small. Thus, no profile in case $1)$ can be a $\frac{(\Omega - U^u)(kc - U^c + U^u)}{cnk}$-equilibrium. In case $2)$, any defender who has different coverage probabilities among her targets can always increase her payoff by decreasing the coverage probabilities of the targets with higher coverage to yield identical coverage for all targets. Consequently, no profile in case $2)$ can be a $\frac{(\Omega - U^u)(kc - U^c + U^u)}{cnk}$-equilibrium. □