Stackelberg Security Games with Multiple Uncoordinated Defenders

Jiarui Gan
University of Oxford
Oxford, United Kingdom
jiarui.gan@cs.ox.ac.uk

Edith Elkind
University of Oxford
Oxford, United Kingdom
elkind@cs.ox.ac.uk

Michael Wooldridge
University of Oxford
Oxford, United Kingdom
mjw@cs.ox.ac.uk

ABSTRACT

Stackelberg security games have received much attention in recent years. While most existing work focuses on single-defender settings, there are many real-world scenarios that involve multiple defenders (e.g., multi-national anti-crime actions in international waters, different security agencies patrolling the same area). In this paper, we consider security games with uncoordinated defenders who jointly protect a set of targets, but may have different valuations for these targets; each defender schedules their own resources and selfishly optimizes their own utility. We generalize the standard (single-defender) model of Stackelberg security games to this setting and formulate an equilibrium concept that captures the nature of strategic interaction among the players. We argue that an exact equilibrium may fail to exist, and, in fact, deciding whether it exists is NP-hard. However, under mild assumptions, every multi-defender security game admits an ϵ-equilibrium for every $\epsilon > 0$, and the limit points corresponding to $\epsilon \to 0$ can be efficiently approximated.

KEYWORDS

Security games; Stackelberg games; multiple defenders

ACM Reference Format:
Jiarui Gan, Edith Elkind, and Michael Wooldridge. 2018. Stackelberg Security Games with Multiple Uncoordinated Defenders. In *Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018), Stockholm, Sweden, July 10–15, 2018,* IFAAMAS, 9 pages.

1 INTRODUCTION

Stackelberg security games (SSGs) offer a framework to optimize allocation of defense resources against strategic adversaries. There is a large body of research on this topic with many successful applications [1, 25]. Much of the existing work focuses on single-defender games, where a defender acts first, allocating resources to protect a set of targets, and an attacker, after having observed the allocation, responds optimally by attacking a most profitable target. However, some real-world defense scenarios may involve multiple defenders. For example, multiple countries may carry out anti-crime actions in international waters at the same time: It was reported that there had been actions against oil-siphoning in the waters between Singapore, Indonesia and Malaysia [26], and against illegal fishing in the Palk Strait between Sri Lanka and India [10]. As another example, different security agencies may have overlapping areas of responsibility and protect targets therein simultaneously: The US Coast Guard and local police departments are known to patrol independently at some major ports in the USA [8]. To model such scenarios, we need to consider *security games with multiple defenders*.

Sometimes the defenders can be assumed to have the same valuation over the targets, and there are a few papers in the literature that study this scenario [2, 8] (we survey the related work in Section 1.1). Another special case considered in prior work is one where each defender is associated with a disjoint subset of targets, and can only protect targets in this subset [15, 16, 23]. In contrast, in our work we consider the more general setting where defenders may value the targets differently, yet each defender can allocate security resources to any of the targets. We assume that all defenders choose their strategies simultaneously, and the attacker responds optimally, choosing a single target to attack. Thus, our primary focus is the game *among the defenders* and the equilibria of this game; we refer to our solution concept as the *Nash-Stackelberg equilibrium (NSE).*

Now, an important feature of SSGs, both with a single defender and with multiple defenders, is that in an equilibrium, an attacker usually finds several targets equally attractive: if, on the contrary, there was a unique target t that he preferred to attack, the defender(s) could then shift some resources towards t from other targets, to make t less vulnerable to the attack while still ensuring that it is the attacker's preferred choice. As ties are ubiquitous, the defenders' beliefs about the attacker's behavior in case of a tie play a crucial role in the analysis. We argue that the standard *optimistic* assumption that the attacker would break ties in favor of the defender(s) is inappropriate in the case of a game with heterogeneous defenders. Instead, we assume that defenders are *pessimistic* regarding the attacker's choices. Under this assumption, we are able to obtain several results concerning the existence and computation complexity of NSE.

Specifically, we establish that an exact equilibrium may not exist due to discontinuity of the defenders' utility functions, and it is NP-hard even just to decide its existence. Nevertheless, the key finding of our work is that every multi-defender game admits an ϵ-equilibrium for every $\epsilon > 0$, and the respective limit points can be efficiently approximated. This suggests that NSE is a promising solution concept for multi-defender security games, which deserves further attention.

1.1 Related Work

There are several papers that consider security games with multiple defenders; however, we believe that our model is richer and more realistic than those considered in prior work. Chan et al. [3] consider games in which multiple players decide the amount of security

Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018), M. Dastani, G. Sukthankar, E. André, S. Koenig (eds.), July 10–15, 2018, Stockholm, Sweden. © 2018 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

investment, but unlike in SSGs the attacker is assumed to act simultaneously with the defenders. Jiang et al. [8] and Basilico et al. [2] study multi-defender games in the leader-follower framework, but limit their consideration to defenders with identical interests. A series of papers [12, 15–17, 23] then consider a multi-defender model where, like in our model, defenders may have different utility functions. However, in this line of work it is assumed that each defender only protects their own set of targets and these sets are pairwise disjoint, which simplifies the problem considerably. In addition, these papers further depart from the standard model of SSGs by assuming that the attacker breaks ties uniformly at random. There are also some research efforts toward multi-leader games beyond the security domain, such as normal-form multi-leader Stackelberg games [11, 14] and oligopoly models in which leaders choose the level of investment or production to maximize profits [6, 21, 22]. However, these papers either employ a very different modeling approach or fall short of offering any algorithmic analysis. We refer the reader to a comprehensive survey by Lou et al. [15].

It is worth mentioning that, in contrast to multi-defender games, games involving multiple attackers (or attacker types) have been studies quite extensively in various settings. These include Bayesian SSGs [5, 18–20, 27], coalitional SSGs with attacker networks [7, 28], and, more recently, generic multi-follower Stackelberg games [4]. However, games with multiple leaders, which are the focus of our work, appear to require fundamentally different techniques than games with multiple followers.

2 A MULTI-DEFENDER MODEL

A multi-defender security game is played between n defenders and an attacker. The defenders protect a set of m targets which the attacker wants to attack. As in an SSG, the defenders act first, simultaneously and independently; the attacker acts after having observed the defenders' strategies.

Specifically, each defender i has k_i defense resources, which she allocates to protect the targets. A target is said to be *protected*, or *covered*, if at least one resource is allocated to it; and *unprotected*, or *uncovered*, otherwise. In the pure strategy setting, an attack on a protected target j is unsuccessful and results in the attacker receiving a penalty p^{a}_j, and each defender i receiving a reward r^{d}_{ij}. An attack on an unprotected target is successful: the attacker receives a reward r^{a}_j, and each defender receives a penalty p^{d}_{ij}.[1] We assume that $r^{\mathrm{d}}_{ij} > p^{\mathrm{d}}_{ij}$ and $r^{\mathrm{a}}_j > p^{\mathrm{a}}_j$, i.e., each defender prefers an attack to be unsuccessful, and the attacker prefers the opposite. Thus, although defenders are heterogeneous, they all want all targets to be safe.

In the mixed strategy setting, the strategy of a defender i is a distribution over all feasible ways of allocating resources. Such a strategy can be compactly represented by a vector $\mathbf{x}_i = \langle x_{ij} \rangle_{j \in [m]} \in \mathcal{X}_i$,[2] where x_{ij} is the probability that target j is protected by defender i, and $\mathcal{X}_i$ denotes the feasible strategy space which, under the resource budget constraint, is $\mathcal{X}_i = \{\mathbf{x}_i : 0 \le x_{ij} \le 1 \text{ for all } j \in [m], \text{ and } \sum_j x_{ij} \le k_i\}$. We do not consider other scheduling constraints in this paper. The joint strategy profile of all defenders is written as $X = \langle x_{ij} \rangle_{i \in [n], j \in [m]}$. The players' utilities depend on the overall probability that the targets are protected, which is referred to as the *coverage*. We denoted by c_j the coverage of target j, and by $\mathbf{c} = \langle c_j \rangle_{j \in [m]}$ the coverage (vector) of all targets. Since each defender acts independently, $\mathbf{c}$ is given by

$$c_j = \mathrm{cov}_j(X) := 1 - \textstyle\prod_{i \in [n]} (1 - x_{ij});$$

we write $\mathbf{c} = \mathrm{cov}(X)$. When the attacker chooses to attack target j, the expected utilities of each defender i and the attacker are given by, respectively:

$$U^{\mathrm{d}}_i(\mathbf{c}, j) = c_j \cdot r^{\mathrm{d}}_{ij} + (1 - c_j) \cdot p^{\mathrm{d}}_{ij};$$
$$U^{\mathrm{a}}(\mathbf{c}, j) = (1 - c_j) \cdot r^{\mathrm{a}}_j + c_j \cdot p^{\mathrm{a}}_j.$$

We do not consider mixed strategy responses of the attacker for reasons that will be clear after we introduce our solution concept.

2.1 Nash–Stackelberg Equilibrium (NSE)

Our solution concept, NSE, captures both the leader-follower structure of the game and the simultaneous moves of the leaders. In an NSE, no defender has the incentive to deviate, assuming that, if she deviates, other defenders will stick to their strategies and the attacker will respond optimally to the new profile. As argued in Section 1, a tie-breaking rule is needed to specify which best response the attacker will choose when more than one is available. In single-defender models, it is normally assumed that the attacker always chooses the target that is optimal for the defender; the respective equilibrium concept is called the *strong Stackelberg equilibrium* (SSE) [13]. Though counter-intuitive, the SSE is justified by the fact that the defender can make the attacker strictly prefer a specific target among the tied ones, by reducing protection of this target by an infinitesimal amount. The resulting strategies, as well as the players' utilities, are arbitrarily close to those under the SSE irrespective of the actual tie-breaking rule. This argument extends to multi-defender games where defenders have *identical* utility functions as in the work of Jiang et al. [8]. A natural counterpart to the SSE is the *weak Stackelberg equilibrium* (WSE) in which the defender pessimistically assumes that the attacker will always choose the worst target. An SSE always exists in a single-defender game, but a WSE may not [24].

Unfortunately, optimistic tie-breaking is inappropriate for our setting because it may lead the defenders to hold inconsistent beliefs as to which target the attacker will choose. Indeed, if multiple defenders simultaneously shift resources away from their least preferred targets, the resulting profile may be far from an equilibrium, in the sense that some defenders can increase their payoff considerably by deviating. That is, unlike in the case of a single defender, if we compute the defenders' strategies under the assumption that they are optimistic, we cannot ensure that, no matter how the attacker breaks ties, there is an ϵ-equilibrium in the neighborhood of the strategy profile that we have computed. On the other hand, a closer look reveals that, in the single-defender setting, while a WSE may fail to exist, the infinitesimal deviation scheme actually generates an approximate WSE, in which the incentive to deviate is infinitesimal. In this sense, the pessimistic tie-breaking is *not* inconsistent with the standard model. We therefore adopt this form of tie-breaking rule in our work. As we will see later, in the multi-defender setting, this approach results in a uniform belief among the defenders about the attacker's response. In addition, it

[1] Both penalties and rewards are values to be *added* to a player's utility.

[2] We write $[z] = \{1, \ldots, z\}$ for any integer $z > 0$.

guarantees a lower bound on the defenders' utilities irrespective of the tie-breaking rule.

We denote by $\mathsf{BR}(\mathbf{c}) := \arg\max_j U^{\mathrm{a}}(\mathbf{c}, j)$ the set of attacker's best (pure strategy) responses to a coverage $\mathbf{c}$, and by $\mathsf{br}_i(\mathbf{c})$ defender i's belief about the attacker's response. In general, $\mathsf{br}_i(\mathbf{c})$ is a mixed strategy supported on pure strategies in $\mathsf{BR}(\mathbf{c})$. However, a pessimistic defender may assume that $\mathsf{br}_i(\mathbf{c})$ is a pure strategy, because the worst-case defender utility can be achieved under a pure strategy response. Thus, we have

$$\mathsf{br}_i(\mathbf{c}) \in \arg\min_{j \in \mathsf{BR}(\mathbf{c})} U_i^{\mathrm{d}}(\mathbf{c}, j).$$

We define NSE in Definition 2.1. Note that, since the attacker's strategy is fully determined by the best response function, we can treat our game as one among the defenders. Thus, we use the term *strategy profile* to refer to the profile of the *defenders'* strategies. We also overload the notations and let $\mathsf{BR}(X) := \mathsf{BR}(\mathsf{cov}(X))$, $\mathsf{br}_i(X) := \mathsf{br}_i(\mathsf{cov}(X))$, and $U_i^{\mathrm{d}}(X, j) := U_i^{\mathrm{d}}(\mathsf{cov}(X), j)$.

Definition 2.1 (NSE). *A (defender) strategy profile X is an ϵ-NSE if for all $i \in [n]$ and all $X' \in \left\{\langle \mathbf{x}_i', X_{-i}\rangle : \mathbf{x}_i' \in \mathcal{X}_i\right\}$ it holds that*

$$U_i^{\mathrm{d}}\left(X, \mathsf{br}_i(X)\right) \geq U_i^{\mathrm{d}}\left(X', \mathsf{br}_i\left(X'\right)\right) - \epsilon, \tag{1}$$

where X_{-i} denotes the strategy profile of defenders $[n] \setminus \{i\}$, and $\langle \mathbf{x}', X_{-i}\rangle$ denotes X with $\mathbf{x}_i$ replaced by $\mathbf{x}'$. An (exact) NSE is a 0-NSE.

We conclude this section with two negative results on NSE. First, the fact that WSE may not exist in single-defender games implies that in multi-defender games the existence of exact NSE cannot be guaranteed. Second, the problem of deciding the existence of NSE is NP-hard; see Theorem 2.2. These results motivate us to focus on approximate equilibria in the remainder of the paper.

Theorem 2.2. *Deciding whether there exists an NSE in a multi-defender game is* NP*-hard.*

Proof. We provide a reduction from the classic NP-complete problem *exact cover by 3-sets* (X3C). Given a set $S = \{e_1, \ldots, e_{3k}\}$ and a collection of subsets $\mathcal{S} = \{S_i\}_{i \in [\ell]}$ of S, each of size 3, $\ell \geq k$, X3C asks if there is a cover of S consisting of k subsets.

For an instance of X3C, we construct a multi-defender game with ℓ defenders, with one resource each, and a set of targets $T \cup T'$ with $T = [3k]$ and $T' = \{3k+1, \ldots, 2k+\ell\}$; note that $|T'| = \ell - k$. Let $T_i = \{j \in T : e_j \in S_i\}$ and $T_{-i} = \{j \in T : e_j \notin S_i\}$; the payoff parameters are set as follows (the value in each entry applies to all targets in the corresponding subset).

	T_i	T_{-i}	T'
r_{ij}^{d}	6	8	2
p_{ij}^{d}	0	7	1

	T	T'
r_j^{a}	3	3
p_j^{a}	0	2

We will now argue that S admits an exact cover if and only if the above game admits an NSE.

(i) *Exact cover ⇒ NSE.* Let $I \subseteq [\ell]$ be an exact cover of S of size k, i.e., $\cup_{i \in I} S_i = S$. One can verify that the following strategy profile X forms an NSE. In X, each of the k defenders $i \in I$ protects the three targets in T_i, each with probability 1/3 (so that $\mathsf{cov}_j(X) = \frac{1}{3}$ for all $j \in T$, as I is an exact cover), and each of the remaining $\ell - k$ defenders protects a unique target in T' with probability 1 (so that $\mathsf{cov}_j(X) = 1$ for all $j \in T'$).

(ii) *∄ exact cover ⇒ ∄ NSE.* Consider an NSE X, First, in X, the attacker finds all targets equally appealing to attack in terms of the expected utilities: if some target is overly protected, defenders can increase their utility by shifting some coverage from this target to the ones that are more attractive to the attacker. Second, for each defender i, $x_{ij} = 0$ for all $j \in T_{-i}$: otherwise i would be tempted to reduce x_{ij} in order to attract the attacker to attack j (even the penalty on targets in T_{-i} is higher than rewards on other targets). Now suppose that S does not admit an exact cover. Then we need more than k subsets to cover S, and correspondingly, more than k defenders to protect T given our second observation. As a result, defenders cannot fully cover all targets in T', and, for targets in T and T' to be equally appealing for the attacker, we have to have $\mathsf{cov}_j(X) < 1/3$ for all $j \in T$, which gives $U_i^{\mathrm{d}}(X, j) < U_i^{\mathrm{d}}(X, j') < U_i^{\mathrm{d}}(X, j'')$ for all $j \in T_i$, $j' \in T'$ and $j'' \in T_{-i}$. By the pessimistic tie-breaking, $\mathsf{br}_i(X) \in T_i$. Thus, a defender i who protects T' would be better off reducing the protection to attract the attacker to attack T', which contradicts the assumption that X is an NSE. □

3 APPROXIMATE EQUILIBRIUM

Clearly, every game admits an ϵ-NSE if ϵ is sufficiently large, so we are interested in the smallest ϵ for which an ϵ-NSE exists. We say that a multi-defender game is *consistent* if $r_j^{\mathrm{a}} > p_j^{\mathrm{a}}$ and $r_{ij}^{\mathrm{d}} > p_{ij}^{\mathrm{d}}$ for all $i \in [n], j \in [m]$. Our key result is that for consistent games ϵ-NSE exists for arbitrarily small values of ϵ.[3] In the remainder of this paper, we only consider consistent games. Theorem 3.2 (below) presents a stronger result, which also establishes the existence of a limit point of ϵ-NSE as $\epsilon \to 0$; we term such a limit point a 0^+-NSE (Definition 3.1). In this section, we first present sufficient conditions for a strategy profile X to be a 0^+-NSE. We use these conditions to prove Theorem 3.2 for a special class of consistent games, which we call *basic games* (Section 4); we extend the proof to all consistent multi-defender games in Section 5.

Definition 3.1 (0^+-NSE). *A strategy profile X forms a 0^+-NSE if there exists a sequence of strategy profiles $\langle X^{(\ell)}\rangle_{\ell=1}^{\infty}$ and a sequence of real numbers $\langle \epsilon^{(\ell)}\rangle_{\ell=1}^{\infty}$, such that every $X^{(\ell)}$ is an $\epsilon^{(\ell)}$-NSE, $\lim_{\ell\to\infty} X^{(\ell)} = X$, and $\lim_{\ell\to\infty} \epsilon^{(\ell)} = 0$.*

Theorem 3.2. *Every consistent multi-defender security game admits a 0^+-NSE.*

3.1 Sufficient Conditions for a 0^+-NSE

We start by introducing a useful notion: the *level coverage*. For the ease of description, we adopt hereinafter the shorthand $\check{u} = \max_j p_j^{\mathrm{a}}$ and $\hat{u} = \max_j r_j^{\mathrm{a}}$.

Definition 3.3 (Height of a coverage). *The* height *of a coverage $\mathbf{c}$ is the optimal attacker utility it induces. We write* $\mathsf{height}(\mathbf{c}) := \max_j U^{\mathrm{a}}(\mathbf{c}, j)$ *and* $\mathsf{height}(X) := \mathsf{height}(\mathsf{cov}(X))$.

Definition 3.4 (Level coverage / strategy profile). *A coverage vector $\mathbf{c}$ is said to be* level*, or to be a* level coverage *if $c_j = 0$ for all $j \notin \mathsf{BR}(\mathbf{c})$. A strategy profile X is* level *if $\mathsf{cov}(X)$ is level.*

[3] However, in the presence of malicious defenders, who prefer some target to be attacked successfully, i.e., $r_{ij}^{\mathrm{d}} < p_{ij}^{\mathrm{d}}$, a ϵ-NSE may not exist for every $\epsilon > 0$. We provide an example in the full version of the paper.

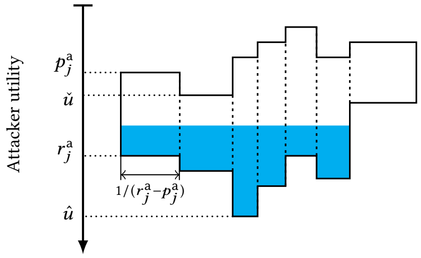

Figure 1: Visualizing a level coverage.

Proposition 3.5. *Let $\overline{\mathrm{cov}}(u)$ be the coverage vector whose j-th element is $\overline{\mathrm{cov}}_j(u) := \max\left\{0, (r_j^{\mathrm{a}} - u)/(r_j^{\mathrm{a}} - p_j^{\mathrm{a}})\right\}$. Then $\overline{\mathrm{cov}}(u)$ is level and it is the only level coverage whose height is u.*

We omit the formal proof of Proposition 3.5 (all proofs can be found in the full version of the paper). To gain some intuition about a level coverage, consider the tank-filling model illustrated in Figure 1. We think of the defense resources (as probability mass under mixed strategies, which is fractional) as water, and each target as a water tank to be filled; the width of tank j is $1/(r_j^{\mathrm{a}} - p_j^{\mathrm{a}})$, so that if water is added to the tank, the *height* of the water surface, as indicated on the axis on the left, is exactly the attacker's utility of attacking j when the coverage c_j is equal to the volume of the water in the tank; see Definition 3.3 (note that the height decreases when the water surface in the visualization rises). The attacker's best responses then correspond to tanks with the lowest surface. If we connect the tanks, gravity will balance the water surfaces so that a grand *level* is maintained across the tanks; the respective lowest surface is maximized and the corresponding coverage is exactly a level coverage, by which no target is overly or insufficiently protected. This also establishes a one-to-one correspondence between a level coverage and its height (Proposition 3.5). Adding all resources into the connected tanks naturally leads to an SSE in the single-defender setting.[4]

This intuition extends to multi-defender games: a 0^+-NSE must be a level strategy, as otherwise defenders would be better off shifting resources from overly protected targets to targets that are less protected. More precisely, Theorem 3.6 offers a set of conditions sufficient for a strategy profile X to be a 0^+-NSE.

Theorem 3.6. *A strategy profile X forms a 0^+-NSE if it satisfies all of the following conditions:*

C1. *X is level.*
C2. *$\sum_j x_{ij} = k_i$ for all i, or* height$(X) = \check{u}$.
C3. *There exist i^*, j^* such that $x_{i^*j^*} > 0$, and for every deviation $\mathbf{x}'_i \in \mathcal{X}_i$ and $X' = \langle \mathbf{x}'_i, X_{-i}\rangle$ it holds that $U_i^{\mathrm{d}}(X, \mathrm{br}_i(X')) \le U_i^{\mathrm{d}}(X, j^*)$ for all $i \in [n]$.*

Intuitively, for any X that satisfies **C1-C3**, if $x_{i^*j^*}$ is reduced by an infinitesimal amount δ, then the attacker will *strictly* prefer attacking j^* since X is level by **C1**. By **C2**, no defender can redirect the attacker to other targets simply by adding resources, so any meaningful deviation will result in coverage decrease on some target, in particular, targets in the attacker's best response set (formally, Lemma 3.7). However, by **C3**, no attacker best response provides a higher defender utility than j^* does even when their coverage stays unchanged. Thus, no deviation offers a utility improvement greater than δ up to a constant multiplier, and hence X is a 0^+-NSE. We present a formal proof after Lemma 3.7.

Lemma 3.7. *Suppose a strategy profile X satisfies **C1** and **C2**. Then, for all i and all $X' \in \{\langle \mathbf{x}'_i, X_{-i}\rangle : \mathbf{x}'_i \in \mathcal{X}_i\}$ it holds that*

$$U^{\mathrm{a}}(X', \mathrm{br}_i(X')) \ge U^{\mathrm{a}}(X, \mathrm{br}_i(X')) = U^{\mathrm{a}}(X, \mathrm{br}_i(X)). \qquad (2)$$

Proof. Since $U^{\mathrm{a}}(X, \mathrm{br}_i(X')) \le U^{\mathrm{a}}(X, \mathrm{br}_i(X))$ holds automatically as $\mathrm{br}_i(X)$ is the best response, we only need to show that $U^{\mathrm{a}}(X', \mathrm{br}_i(X')) \ge U^{\mathrm{a}}(X, \mathrm{br}_i(X')) \ge U^{\mathrm{a}}(X, \mathrm{br}_i(X))$. Suppose for a contradiction that at least one of the following inequalities holds:

(a) $U^{\mathrm{a}}(X', \mathrm{br}_i(X')) < U^{\mathrm{a}}(X, \mathrm{br}_i(X'))$;
(b) $U^{\mathrm{a}}(X, \mathrm{br}_i(X')) < U^{\mathrm{a}}(X, \mathrm{br}_i(X))$.

We claim that either will lead to the following for all $j \in \mathrm{BR}(X)$:

$$\begin{aligned} U^{\mathrm{a}}(X', j) &\le^{\dagger} U^{\mathrm{a}}(X', \mathrm{br}_i(X')) \\ &\le^{\ddagger} U^{\mathrm{a}}(X, \mathrm{br}_i(X')) \le^{\S} U^{\mathrm{a}}(X, j)\,. \end{aligned}$$

Specifically, $\le^{\dagger}$ holds because $\mathrm{br}_i(X')$ is the best response to X'; $\le^{\ddagger}$ holds directly with assumption (a), while with assumption (b), it holds because (b) implies $\mathrm{br}_i(X') \notin \mathrm{BR}(X)$, so $\mathrm{cov}_{\mathrm{br}_i(X')}(X) = 0$ as X is level; finally, $\le^{\S}$ holds because $j \in \mathrm{BR}(X)$ is a best response to X. In particular, $\le^{\ddagger}$ and $\le^{\S}$ are strictly satisfied by (a) and (b), respectively, so the above inequalities eventually give

$$U^{\mathrm{a}}(X', j) < U^{\mathrm{a}}(X, j) \quad \text{for all } j \in \mathrm{BR}(X). \qquad (3)$$

We show that this contradicts both conditions in **C2**.

(i) Suppose $\sum_{j\in[m]} x_{ij} = k_i$ for all i. Since $U^{\mathrm{a}}(X, j)$ is monotone decreasing w.r.t. x_{ij}, Eq. (3) implies that $x'_{ij} > x_{ij}$ for all $j \in \mathrm{BR}(X)$. Furthermore, by the definition of level coverage, for those $j \notin \mathrm{BR}(X)$, we have $\mathrm{cov}_j(X) = 0$ and naturally $x_{ij} = 0 \le x'_{ij}$. It follows that $\sum_j x'_{ij} > \sum_j x_{ij} = k_i$ for all i, which violates the constraint on the number of available resources.

(ii) Suppose $\mathrm{cov}(X) = \overline{\mathrm{cov}}(\check{u})$. We have $U^{\mathrm{a}}(X, \mathrm{br}_i(X)) = \check{u}$, so by Eq. (3) $U^{\mathrm{a}}(X', j) < U^{\mathrm{a}}(X, j) = U^{\mathrm{a}}(\mathrm{br}_i(X), X) = \check{u}$ for all $j \in \mathrm{BR}(X)$. In particular, if we take $j = \mathrm{br}_i(X)$, this becomes $U^{\mathrm{a}}(X', \mathrm{br}_i(X)) < \check{u} = \max_j p_j^{\mathrm{a}}$, which is a contradiction. □

Proof of Theorem 3.6. We construct an ϵ-NSE $\tilde{X}$ by letting defender i^* reduce $x_{i^*j^*}$ by a small amount $\delta > 0$, and show that $\tilde{X}$ forms a $\lambda\cdot\delta$-NSE for some constant λ, so when $\delta \to 0$, we have $\lambda\cdot\delta \to 0$ and $\tilde{X} \to X$, which implies that X is a 0^+-NSE.

Consider an arbitrary deviation $\mathbf{x}'_i \in \mathcal{X}_i$. Let $\tilde{X}' = \langle \mathbf{x}'_i, \tilde{X}\rangle$ and $X' = \langle \mathbf{x}'_i, X_{-i}\rangle$. $\tilde{X}'$ can be viewed as X' with $x'_{i^*j^*}$ reduced by at most δ (exactly δ when $i \ne i^*$, and 0 otherwise). Thus, we have

$$\mathrm{cov}_{j^*}(X') - \delta \le \mathrm{cov}_{j^*}(\tilde{X}') \le \mathrm{cov}_{j^*}(X'); \qquad (4)$$

and for all other $j \ne j^*$,

$$\mathrm{cov}_j(X') = \mathrm{cov}_j(\tilde{X}'). \qquad (5)$$

[4] The visualization is inspired by the classic algorithm ORIGAMI [9], though, as the reader will see later, the algorithm we develop based on these ideas is very different from ORIGAMI, even in the degenerate single-defender case.

Since $x_{i^*j^*} > 0$ by **C3**, we have $\mathrm{cov}_{j^*}(X) > 0$ and, thus, $j^* \in \mathrm{BR}(X)$ by the definition of level coverage. Let $j' = \mathrm{br}_i(X')$. By Lemma 3.7,

$$U^{\mathrm{a}}(X', j') \geq U^{\mathrm{a}}(X, j') = U^{\mathrm{a}}(X, j^*), \tag{6}$$

from which we also know that $\mathrm{cov}_{j'}(X') \leq \mathrm{cov}_{j'}(X)$ as $U^{\mathrm{a}}(\mathbf{c}, j)$ decreases with c_j, so as $U_i^{\mathrm{d}}(\mathbf{c}, j)$ increases with c_j,

$$U_i^{\mathrm{d}}(X', j') \leq U_i^{\mathrm{d}}(X, j'). \tag{7}$$

Consider the following two cases:

(i) If $\mathrm{br}_i(\tilde{X}') = \mathrm{br}_i(X')$, we have

$$\begin{aligned} U_i^{\mathrm{d}}(\tilde{X}', \mathrm{br}_i(\tilde{X}')) = U_i^{\mathrm{d}}(\tilde{X}', j') &\leq^{\dagger} U_i^{\mathrm{d}}(X', j') \leq^{\ddagger} U_i^{\mathrm{d}}(X, j') \\ \leq^{\S} U_i^{\mathrm{d}}(X, j^*) &= U_i^{\mathrm{d}}(\tilde{X}, j^*) + \delta \cdot (r_{ij^*}^{\mathrm{d}} - p_{ij^*}^{\mathrm{d}}) \\ &=^{\P} U_i^{\mathrm{d}}(\tilde{X}, \mathrm{br}_i(\tilde{X})) + \delta \cdot (r_{ij^*}^{\mathrm{d}} - p_{ij^*}^{\mathrm{d}}), \end{aligned} \tag{8}$$

where $\leq^{\dagger}$ holds by Eq. (4); $\leq^{\ddagger}$ holds by Eq. (7); $\leq^{\S}$ holds by **C3**; and $=^{\P}$ holds because $\mathrm{br}_i(\tilde{X}) = j^*$ as j^* is strictly preferred by the attacker after the deviation.

(ii) If $\mathrm{br}_i(\tilde{X}') \neq \mathrm{br}_i(X')$, we must have $\mathrm{br}_i(\tilde{X}') = j^* \neq \mathrm{br}_i(X')$ by Eqs. (4) and (5). It follows that

$$\begin{aligned} U^{\mathrm{a}}(\tilde{X}', j^*) &\geq^{\dagger} U^{\mathrm{a}}(\tilde{X}', j') =^{\ddagger} U^{\mathrm{a}}(X', j') \\ &\geq^{\S} U^{\mathrm{a}}(X, j^*) \geq^{\P} U^{\mathrm{a}}(\tilde{X}, j^*) - \delta \cdot (r_{j^*}^{\mathrm{a}} - p_{j^*}^{\mathrm{a}}), \end{aligned}$$

where $\geq^{\dagger}$ holds because $j^* = \mathrm{br}_i(\tilde{X}')$; $=^{\ddagger}$ holds by Eq. (5); $\geq^{\S}$ holds by Eq. (6); and $\geq^{\P}$ holds as, by definition, $\tilde{X}$ is obtained from X by reducing $x_{i^*j^*}$ by δ. This implies $\mathrm{cov}_{j^*}(\tilde{X}') \leq \mathrm{cov}_{j^*}(\tilde{X}) + \delta$, and finally $U_i^{\mathrm{d}}(\tilde{X}', j^*) \leq U_i^{\mathrm{d}}(\tilde{X}, j^*) + \delta \cdot (r_{ij^*}^{\mathrm{d}} - p_{ij^*}^{\mathrm{d}})$, or

$$U_i^{\mathrm{d}}(\tilde{X}', \mathrm{br}_i(\tilde{X}')) \leq U_i^{\mathrm{d}}(\tilde{X}, \mathrm{br}_i(\tilde{X})) + \delta \cdot (r_{ij^*}^{\mathrm{d}} - p_{ij^*}^{\mathrm{d}}). \tag{9}$$

As Eqs. (8) and (9) imply, in both cases $\tilde{X}$ forms a λ·δ-NSE with $\lambda = \max_i (r_{ij^*}^{\mathrm{d}} - p_{ij^*}^{\mathrm{d}})$ being a constant. This concludes the proof. □

4 0^+-NSE IN BASIC GAMES

In this section, we prove Theorem 3.2 for *basic games*. These are games where the defenders' payoff parameters are such that their preferences concerning the attacker's response are independent of the strategy profile being played.

Definition 4.1 (Basic game). *A consistent multi-defender game is called a* basic game *if for every defender i* $\in [n]$ *and every pair of distinct targets* $j, j' \in [m]$, *one of the following conditions holds:*

(i) $r_{ij}^{\mathrm{d}} < p_{ij'}^{\mathrm{d}}$, *in which case we write* $j \succ_i j'$ *(in the sense that* $U_i^{\mathrm{d}}(X, j) \leq U_i^{\mathrm{d}}(X, j')$ *holds for any X, i.e., an attack on j is always more detrimental to i than an attack on j′, and protecting j is a higher priority).*

(ii) $r_{ij'}^{\mathrm{d}} < p_{ij}^{\mathrm{d}}$, *in which case we write* $j \prec_i j'$.

We also write $j \succeq_i j'$ *if* $j \succ_i j'$ *or* $j = j'$; *and* $j \preceq_i j'$ *if* $j \prec_i j'$ *or* $j = j'$.

In a basic game, the defenders' utility functions can be represented concisely by a *preference profile*, i.e., a list of the defenders' preference orders over the targets. We denote by $j_{i\ell}$ the ℓ-th most important target for defender i (so that $j_{i1} \succ_i j_{i2} \succ_i \ldots j_{im}$), and by $J = \langle j_{i\ell} \rangle_{i \in [n], \ell \in [m]}$ the preference profile. A basic game is then a tuple $\langle J, \mathbf{r}^{\mathrm{a}}, \mathbf{p}^{\mathrm{a}}, \mathbf{k} \rangle$. For basic games, condition **C3** in Theorem 3.6 can be restated as follows.

C3′. There exist i^*, j^* such that $x_{i^*j^*} > 0$, and for every deviation $\mathbf{x}_i' \in \mathcal{X}_i$ and $X' = \langle \mathbf{x}_i', X_{-i} \rangle$ it holds that $\mathrm{br}_i(X') \succeq_i j^*$.

4.1 A Proof by Construction

We construct an X that satisfies **C1**–**C3** (**C3′**). Since $\mathrm{cov}(X)$ must be level by **C1** and there is a one-to-one correspondence between a level coverage and its height, our approach is to scan through the range $[\check{u}, \hat{u}]$ to find a u such that $\overline{\mathrm{cov}}(u)$ can be implemented by some X satisfying **C2** and **C3**. Implementability is tested by Procedure 1. Though the approach adopted by Procedure 1 is *not* the only way to implement $\overline{\mathrm{cov}}(u)$, we simply discard candidates that do not pass the test; this suffices for the purpose of our proof.

Procedure 1: Check the implementability of $\overline{\mathrm{cov}}(u)$

```
for all $i$, $j$ do $x_{ij} \leftarrow 0$;
for $i = 1, \ldots, n$ do
    for $j = j_{i1}, j_{i2}, \ldots, j_{im}$ do
        Increase $x_{ij}$ until $\mathrm{cov}_j(X)$ reaches $\overline{\mathrm{cov}}_j(u)$ or $\sum_{j'} x_{ij'}$ reaches $k_i$;
        if $x_{ij} > 0$ then $i^* \leftarrow i$, and $j^* \leftarrow j$;
Check if $\mathrm{cov}(X) = \overline{\mathrm{cov}}(u)$ and $\sum_j x_{ij} = k_i$ for all $i$.
```

Procedure 1 allocates resources according to the importance of the targets; those with higher importance receive resources first. Every time i allocates a positive amount of resource, j^* is updated to the current target j, so that it holds that $x_{ij} = 0$ for all $j \prec_i j^*$ throughout the procedure. Now, suppose that in the end $\mathrm{cov}(X) = \overline{\mathrm{cov}}(u)$ and $\sum_j x_{ij} = k_i$ for all i. Then X satisfies **C1** and **C2**. It also satisfies **C3**, for the following reason: now that $\sum_j x_{ij} = k_i$, the defenders have no resources left, so they have to reduce the coverage of some target if they deviate, but this will only redirect the attacker to a target $j \succeq_i j^*$ as $x_{ij} = 0$ for all $j \prec_i j^*$. Thus, X is a 0^+-NSE by Theorem 3.6. Otherwise, there can be two cases:

(i) Deficit: $\sum_j x_{ij} = k_i$ for all i, but $\mathrm{cov}_j(X) < \overline{\mathrm{cov}}_j(u)$ for some j, in which case resources are used up but $\overline{\mathrm{cov}}(u)$ is not yet reached, so there is a *deficit* as more resources are needed.

(ii) Surplus: $\mathrm{cov}(X) = \overline{\mathrm{cov}}(u)$, but $\sum_j x_{ij} < k_i$ for some i, in which case $\overline{\mathrm{cov}}(u)$ is reached, but some defenders have spare resources, so there is a resource *surplus*.

In case of a deficit, we increase u (i.e., lower the goal line in the tank-filling model) to reduce the demand, and in case of a surplus, we do the opposite. The aim is to find a zero point, i.e., a point with neither a deficit nor a surplus. Observe that there is always a surplus at $u = \hat{u}$, as no resource is needed when the goal line is set at the bottom. Further, the demand changes continuously with the goal line u (Lemma 4.2). Therefore, if there is a deficit at $u = \check{u}$, there must be a zero point in $[\check{u}, \hat{u}]$. However, if there is a surplus at $u = \check{u}$, we cannot further reduce u while maintaining the level. To deal with this special case, we introduce a stronger procedure called Alloc (Procedure 2).

A Two-phase Allocation Alloc has two phases. In the first phase each defender i only allocates resources to targets $j \succ_i$

Procedure 2: ALLOC

```
input  : a basic game $\langle J, \mathbf{r}^{\mathrm{a}}, \mathbf{p}^{\mathrm{a}}, \mathbf{k}\rangle$; $u \in [\check{u}, \hat{u}]$.
output : (1) $X = \langle x_{ij}\rangle$; (2) $i^*, j^*$; (3) surp.
1 Initialize $x_{ij} \leftarrow 0$, $\Delta_j \leftarrow \overline{\mathrm{cov}}_j(u)$, and $c_j \leftarrow 0$, for all $i, j$;
  /* --------------------- Phase 1 --------------------- */
2 for $i = 1, \dots, n$ do
3     for $j = j_{i1}, \dots, j_{im}$ such that $j \succ_i \mathrm{br}_i(\mathbf{1}^m)$ do
4         $x_{ij} \leftarrow \min\left\{\Delta_j,\ k_i - \sum_{j' \succ_i j} x_{ij'}\right\}$;
5         $c_j \leftarrow 1 - (1 - c_j)(1 - x_{ij})$;
6         $\Delta_j \leftarrow \begin{cases} 0, & \text{if } c_j = 1 \\ \frac{\overline{\mathrm{cov}}_j(u) - c_j}{1 - c_j}, & \text{otherwise} \end{cases}$;
7         if $x_{ij} > 0$ then $i^* \leftarrow i$ and $j^* \leftarrow j$;
  /* --------------------- Phase 2 --------------------- */
8 Repeat Lines 2–7 with $j = \mathrm{br}_i(\mathbf{1}^m), \dots, j_{im}$ in Line 3;
9 surp $\leftarrow \sum_i \left(k_i - \sum_j x_{ij}\right) - \sum_j \Delta_j$.
```

$\mathrm{br}_i(\mathbf{1}^m)$,[5] and in the second phase each defender i allocates the remaining resources to targets $j \preceq_i \mathrm{br}_i(\mathbf{1}^m)$. Here $\mathrm{br}_i(\mathbf{1}^m)$ is the attacker's best response to coverage $\mathbf{1}^m$, which is also the best response to any coverage $\mathbf{c} > \overline{\mathrm{cov}}(\check{u})$, coordinate-wise (if one keeps adding water after the level has reached $\check{u}$, the tanks with the lowest surface will stay the same). Lines 4–6 detail Line 4 of Procedure 1, where c_j is the current coverage of target j, and Δ_j is the amount of resources still needed in order for the coverage to reach $\overline{\mathrm{cov}}_j(u)$. In the end, as part of the output, surp records the amount of resource surplus (which represents a deficit if negative).

LEMMA 4.2. *Fix all other input parameters and let* $\mathrm{surp}(u)$ *be the output* surp *of* ALLOC *on input* u. $\mathrm{surp}(u)$ *is continuous and strictly monotone increasing in* $[\check{u}, \hat{u}]$.

We are ready to prove Theorem 3.2 for basic games, restated as Theorem 4.3 below.

THEOREM 4.3 (THEOREM 3.2 RESTRICTED TO BASIC GAMES). *Every basic game admits a* 0^+*-NSE.*

PROOF. $\mathrm{surp}(u)$ is continuous by Lemma 4.2, and $\mathrm{surp}(\hat{u}) = \sum_i k_i > 0$, as no resource is demanded when the goal line is set to $\hat{u}$. Thus, if $\mathrm{surp}(\check{u}) \le 0$, there must be a zero point u^* at which $\mathrm{surp}(u^*) = 0$; the output X satisfies **C1**–**C3** and hence forms a 0^+-NSE.

On the other hand, if $\mathrm{surp}(\check{u}) > 0$, the output X with $u = \check{u}$ satisfies **C1** and **C2**. To see that it also satisfies **C3** (**C3′**), consider the following cases:

(i) If a defender i has no spare resource, she needs to reduce the coverage of a target in order to redirect the attacker to that target. This, however, can only be done on targets that are worse for i since $x_{ij} = 0$ for all $j \prec_i j^*$. Therefore, $\mathrm{br}_i(X') \succeq_i j^*$.

(ii) If i has spare resources, coverage can be increased. We make two observations. First, the coverage of all $j \succ_i \mathrm{br}_i(\mathbf{1}^m)$ must have reached the goal line in Phase 1. Second, those in $\mathrm{BR}(\mathbf{1}^m)$ are all skipped in Phase 1, so they must be filled in Phase 2; as a result, j^* will be updated to $\mathrm{br}_i(\mathbf{1}^m)$ or even beyond that in i's preference order, so $j^* \preceq_i \mathrm{br}_i(\mathbf{1}^m)$. Note that defender i cannot remove $\mathrm{br}_i(\mathbf{1}^m)$ from the best response set by *adding* more resources on top of X as $\mathrm{br}_i(\mathbf{1}^m)$ is already fully covered. Therefore, by the pessimistic tie-breaking rule i cannot redirect the attacker to any $j \prec_i \mathrm{br}_i(\mathbf{1}^m)$, nor to any $j \prec_i j^*$, since $j^* \preceq_i \mathrm{br}_i(\mathbf{1}^m)$ by our second observation. Therefore, $\mathrm{br}_i(X') \succeq_i j^*$.

C3′ is satisfied in both cases, so by Theorem 3.6 X is a 0^+-NSE. This completes the proof. □

5 0^+-NSE IN GENERAL CONSISTENT GAMES

Unlike in basic games, in the general case the defenders' preferences over targets depend on the actual coverage vector. Thus, if we take a preference profile J and attempt to construct a 0^+-NSE X using our previous approach (assuming that defenders' preferences are specified by J as in a basic game), the defenders' preferences at X may end up being different from J and hence X may fail to be a 0^+-NSE of the original game. Therefore, we want to find a "stationary point" where the preference profile J used to construct X coincides with the profile induced by X. We define *induced preference profiles* formally in Definition 5.1. In contrast with basic games, we allow players to be indifferent among different targets in an induced preference profile.

DEFINITION 5.1 (INDUCED PREFERENCE PROFILE). *A preference profile* J induced by a coverage $\mathbf{c}$ *is a list of binary relations* $\langle \succeq_i^J \rangle_{i \in [n]}$ *such that* $j_1 \succeq_i^J j_2$ *iff* $U_i^{\mathrm{d}}(\mathbf{c}, j_1) \le U_i^{\mathrm{d}}(\mathbf{c}, j_2)$. *A preference profile* induced by a strategy profile X *is the preference profile induced by* $\mathrm{cov}(X)$.

With the above definition in hand, **C3** can be restated as follows:

C3″. There exist i^*, j^* such that $x_{i^*j^*} > 0$ and for every deviation $\mathbf{x}'_i \in \mathcal{X}_i$ and $X' = \langle \mathbf{x}'_i, X_{-i}\rangle$ it holds that $\mathrm{br}_i(X') \succeq_i^J j^*$, *where* J *is induced by* X.

5.1 Existence of a Stationary Point

If a strategy profile X is a 0^+-NSE, by **C1** it must be level. Thus, we limit our attention to level strategy profiles, and hence to preferences induced by level strategy profiles. Consider moving u from $\check{u}$ to $\hat{u}$. The corresponding level coverage changes from $\overline{\mathrm{cov}}(\check{u})$ to $\overline{\mathrm{cov}}(\hat{u})$, and in the meantime the preference profile induced by $\overline{\mathrm{cov}}(u)$ changes when some defender's utility functions on different targets intersect, i.e., when $U_i^{\mathrm{d}}(\overline{\mathrm{cov}}(u), j) = U_i^{\mathrm{d}}(\overline{\mathrm{cov}}(u), j')$ for some $j \ne j'$. Observe that $U_i^{\mathrm{d}}(\overline{\mathrm{cov}}(\cdot), j)$ is continuous in $[\check{u}, \hat{u}]$ for all i and j. Let

$$\mathcal{U} = \left\{u \in [\check{u}, \hat{u}] : U_i^{\mathrm{d}}(\overline{\mathrm{cov}}(u), j) = U_i^{\mathrm{d}}(\overline{\mathrm{cov}}(u), j'), i \in [n], j \ne j'\right\} \quad (10)$$

be the values of u at the intersections.[6] Let $u^{(\ell)}$ be the ℓ-th smallest element in $\mathcal{U} \cup \{\check{u}, \hat{u}\}$. For each ℓ, all $\overline{\mathrm{cov}}(u)$ with $u \in (u^{(\ell)}, u^{(\ell+1)})$ thus induce the same preference profile; denote it by $J^{(\ell)}$. Further, we map each $J^{(\ell)}$ to a 0^+-NSE $X^{(\ell)}$ of a basic game with preferences $J^{(\ell)}$ (let indifference between targets be resolved by target indices):

[5] $\mathbf{1}^m$ denotes an all-1 vector of length m.

[6] If two utility functions are the same in an interval, there can be infinitely many intersection points. In this case, we only include the endpoints of that interval in $\mathcal{U}$. This suffices, because our goal is to identify distinct preference profiles induced by the utility functions, and the defender's preferences do not change in the interval.

fix the preferences to $J^{(\ell)}$, find the zero point (or $\check{u}$ if there is no zero point) of surp in the procedure Alloc, and take the strategy profile at this zero point. Let $h^{(\ell)} = \mathsf{height}(X^{(\ell)})$.

Now each interval $(u^{(\ell)}, u^{(\ell+1)})$ is associated with a value $h^{(\ell)}$. Intuitively, if some $h^{(\ell)}$ happens to fall back in $(u^{(\ell)}, u^{(\ell+1)})$, then $h^{(\ell)}$ is a stationary point, as desired. The following lemma establishes the existence of two types of stationary points. We use it to prove Theorem 3.2 below by showing the existence of an X satisfying **C1–C3 (C3′′)**.

Lemma 5.2. *Let $u^{(1)}, \dots, u^{(L)}$ and $h^{(1)}, \dots, h^{(L-1)}$ be sequences of real numbers, such that $u^{(1)} < u^{(2)} < \cdots < u^{(L)}$ and $u^{(1)} \le h^{(\ell)} \le u^{(L)}$ for all ℓ. Then there exists either $h^{(\ell)} \in [u^{(\ell)}, u^{(\ell+1)}]$ or $u^{(\ell)} \in (\!|h^{(\ell-1)}, h^{(\ell)}|\!)$, where $(\!|a, b|\!) := \{x : a < x < b \vee b < x < a\}$ denotes the open interval between points $a, b \in \mathbb{R}$.*

Proof. Suppose for the sake of contradiction that neither $h^{(\ell)}$ nor $u^{(\ell)}$ exists. Then for all ℓ it holds that (i) $h^{(\ell)} \notin [u^{(\ell)}, u^{(\ell+1)}]$ and (ii) $u^{(\ell)} \notin (\!|h^{(\ell-1)}, h^{(\ell)}|\!)$. In particular, $h^{(1)} \notin [u^{(1)}, u^{(2)}]$; since $u^{(1)} \le h^{(\ell)} \le u^{(L)}$ for all ℓ, this implies $h^{(1)} > u^{(2)}$. Now, suppose that $h^{(t)} > u^{(t+1)}$ for all $t < \ell$. If $h^{(\ell+1)} \le u^{(\ell+2)}$ then either $u^{(\ell+1)} \le h^{(\ell+1)} \le u^{(\ell+2)}$, which contradicts (i), or $h^{(\ell+1)} < u^{(\ell+1)}$, which implies $h^{(\ell+1)} < u^{(\ell+1)} < h^{(\ell)}$ and contradicts (ii). Thus, we have $h^{(\ell+1)} > u^{(\ell+2)}$, and, by induction, $h^{(\ell)} > u^{(\ell+1)}$ for all $\ell \le L$. However, this contradicts the assumption that $h^{(L-1)} \le u^{(L)}$. □

Theorem 3.2 (restated). *Every consistent multi-defender security game admits a 0^+-NSE.*

Proof. By Lemma 5.2, there exists either (i) $h^{(\ell)} \in [u^{(\ell)}, u^{(\ell+1)}]$, or (ii) $u^{(\ell)} \in (\!|h^{(\ell-1)}, h^{(\ell)}|\!)$. W.l.o.g. we assume $h^{(\ell-1)} < u^{(\ell)} < h^{(\ell)}$ for Case (ii). In Case (i), the argument in the proof of Theorem 4.3 shows that $X^{(\ell)}$ satisfies **C1–C3 (C3′′)**. In Case (ii), $X^{(\ell-1)}$ and $X^{(\ell)}$ satisfy **C1** and **C2**, but neither of them satisfies **C3**. We show that we can "merge" them into a profile X^* that satisfies all three conditions **C1–C3 (C3′′)**. Procedure Alloc# (Procedure 3), developed on the basis of Alloc, accomplishes the task.

Alloc# is again a two-phase procedure that calls Alloc twice on two preference profiles $J^{(\ell-1)}$ and $J^{(\ell)}$: Phase 1 is exactly the same as Alloc, with the preferences set to $J^{(\ell-1)}$ and amount of resources set to $\theta{\cdot}k_i$ for each i; Phase 2 changes the preferences to $J^{(\ell)}$ and allocates the remaining amount of resources, i.e., $(1-\theta){\cdot}k_i$, on top of the allocation of Phase 1. Slightly differently from Alloc, in Line 5, the amount of resources demanded is set to $\Delta_j{\cdot}(1 - x_{ij})$ to account for resources already allocated (to target j by defender i) in Phase 1: one can verify that, if i adds $\Delta_j{\cdot}(1 - x_{ij})$ resources to the current allocation $X = \langle x_{ij} \rangle$, resulting in X', then $\mathsf{cov}_j(X') = \overline{\mathsf{cov}}_j(u)$.

Having all the other parameters fixed, we view the outputs as functions of θ. We find that, first, $\mathsf{surp}^{\#}(\theta)$ is continuous w.r.t. θ, by showing inductively that all variables can be expressed as continuous functions in the closed form. Second, when $\theta = 1$, Phase 2 is actually skipped and Alloc# degenerates to Alloc with preferences $J^{(\ell-1)}$; whereas when $\theta = 0$, Phase 1 is skipped and Alloc# degenerates to Alloc with preferences $J^{(\ell)}$. Formally,

$$\mathsf{surp}^{\#}(1) = \mathsf{surp}(J^{(\ell-1)}, u^{(\ell)}) > \mathsf{surp}(J^{(\ell-1)}, h^{(\ell-1)}) = 0, \text{ and}$$

$$\mathsf{surp}^{\#}(0) = \mathsf{surp}(\ J^{(\ell)}\ ,\ u^{(\ell)}) < \mathsf{surp}(\ J^{(\ell)}\ ,\ \ h^{(\ell)}\ ) = 0,$$

Procedure 3: Alloc#

```
   input  : $J^{(\ell-1)}, J^{(\ell)}; \mathbf{r}^{\mathrm{a}}, \mathbf{p}^{\mathrm{a}}, \mathbf{k}; u^{(\ell)}; \theta \in [0, 1]$.
   output : (1) $X = \langle x_{ij} \rangle$; (2) $i^*, j^*$; (3) $\mathsf{surp}^{\#}$.
1  Initialize $x_{ij} \leftarrow 0$, $\Delta_j \leftarrow \overline{\mathsf{cov}}_j(u^{(\ell)})$, and $c_j \leftarrow 0$, for all $i, j$;
   /* ---------- Phase 1: allocation by $J^{(\ell-1)}$ ---------- */
2  Set $J \leftarrow J^{(\ell-1)}$, and $\tilde{k}_i \leftarrow \theta \cdot k_i$ for all $i$;
3  for $i = 1, \dots, n$ do
4      for $j = j_{i1}, \dots, j_{im}$ such that $j \succ_i^J \mathsf{br}_i^J(\mathbf{1}^m)$ do
5          $x_{ij} \leftarrow x_{ij} + \min\left\{\Delta_j{\cdot}(1 - x_{ij}),\ \tilde{k}_i - \sum_{j' \succ_i^J j} x_{ij'}\right\}$;
6          $c_j \leftarrow \mathsf{cov}(X)$;
7          $\Delta_j \leftarrow \frac{\overline{\mathsf{cov}}_j(u^*) - c_j}{1 - c_j}$;[7]
8          if $x_{ij} > 0$ then $j^* \leftarrow j$, $i^* \leftarrow i$;
9  Repeat Lines 3–8 with $j = \mathsf{br}_i^J(\mathbf{1}^m), \dots, j_{im}$ in Line 4;
   /* ------------ Phase 2: allocation by $J^{(\ell)}$ ---------- */
10 Reset $J \leftarrow J^{(\ell)}$, and $\tilde{k}_i \leftarrow (1 - \theta) \cdot k_i$ for all $i$;
11 Repeat Lines 3–9;
12 $\mathsf{surp}^{\#} \leftarrow \sum_i \left(k_i - \sum_j x_{ij}\right) - \sum_j \Delta_j$.
```

where $\mathsf{surp}(J, u)$ denotes the output surp of Alloc, as a function of u and the preference profile J. Thus, there exists a zero point $\theta^* \in (0, 1)$ such that $\mathsf{surp}^{\#}(\theta^*) = 0$. Let $X^* = X(\theta^*)$. From $\mathsf{surp}^{\#}(\theta^*) = 0$ we know that $\overline{\mathsf{cov}}_j(u^{(\ell)}) - \mathsf{cov}_j(X^*) = 0$ for all j and $k_i - \sum_j x^*_{ij} = 0$ for all i, so **C1** and **C2** are satisfied by X^*.

To see that **C3′′** is satisfied, we show that $\mathsf{br}_i(X') \succeq_i^J j^*$ for the preference profile J induced by X^* (equivalently, J is induced by $\overline{\mathsf{cov}}(u^{(\ell)})$). Observe that $j \succeq_i^{J^{(\ell-1)}} j^*$ for each target j that receives a positive amount of resources from some defender i in Phase 1, and $j \succeq_i^{J^{(\ell)}} j^*$ for each target j that receives a positive amount of resources from some defender i in Phase 2. These are all the targets defender i can redirect the attacker to through strategy deviation. Thus, for any X' as the result of a deviation of defender i, either $\mathsf{br}_i(X') \succeq_i^{J^{(\ell-1)}} j^*$ or $\mathsf{br}_i(X') \succeq_i^{J^{(\ell)}} j^*$. Suppose that $\mathsf{br}_i(X') \succeq_i^{J^{(\ell-1)}} j^*$. (The same argument applies in the other case.) Since $\overline{\mathsf{cov}}(u)$ induces $J^{(\ell-1)}$ for all $u \in (u^{(\ell-1)}, u^{(\ell)})$, by Definition 5.1, the inequality $\mathsf{br}_i(X') \succeq_i^{J^{(\ell-1)}} j^*$ then implies, for all $u \in (u^{(\ell-1)}, u^{(\ell)})$,

$$U_i^{\mathrm{d}}(\overline{\mathsf{cov}}(u), \mathsf{br}_i(X')) - U_i^{\mathrm{d}}(\overline{\mathsf{cov}}(u), j^*) \le 0,$$

which also holds for $u = u^{(\ell)}$ because $U_i^{\mathrm{d}}(\overline{\mathsf{cov}}(\cdot), j)$ is continuous. Thus, $U_i^{\mathrm{d}}(\overline{\mathsf{cov}}(u^{(\ell)}), \mathsf{br}_i(X')) - U_i^{\mathrm{d}}(\overline{\mathsf{cov}}(u^{(\ell)}), j^*) \le 0$, and again by Definition 5.1, this implies $\mathsf{br}_i(X') \succeq_i^J j^*$ for J induced by $u^{(\ell)}$.

To summarize, in Case (i), $X^{(\ell)}$ satisfies **C1–C3 (C3′′)** and forms a 0^+-NSE; in Case (ii), we merge $X^{(\ell-1)}$ and $X^{(\ell)}$ into X^*, which achieves the same. Therefor, a 0^+-NSE always exists. □

[7] Unlike in Alloc, it always holds that $c_j < 1$ because $u^{(\ell)} > h^{(\ell-1)} \ge \check{u}$.

6 COMPUTING A 0^+-NSE

Now that we have shown the existence of a 0^+-NSE, the proof also yields a method to compute a 0^+-NSE. Since a 0^+-NSE may involve irrational numbers,[8] we are interested in approximate solutions with given precision. Specifically, for a $\delta > 0$, we say that a 0^+-NSE X is *approximated to precision* δ if we obtain an X' such that $|x_{ij} - x'_{ij}| \le \delta$ for all i, j. Below, Theorem 6.3 shows that such a solution can be computed in polynomial time. The proof relies on Lemmas 6.1 and 6.2. Without loss of generality, all results in this section assume that all payoff parameters are positive integers encoded in binary.

Lemma 6.1. *Let u^* be the zero point of* $\mathrm{surp}(u)$ *as in* Alloc *(or $u^* = \check{u}$ if no zero point exists). For any given $\delta > 0$, $X(u^*)$ can be approximated to precision δ in time polynomial in the size of the input parameters and* $\log \frac{1}{\delta}$.

Proof Sketch. Since Alloc runs in polynomial time, we can run binary search on the input u to approximate $\tilde{u}^*$, and then take $X(\tilde{u}^*)$ as an approximation of $X^{(\ell)}$. For this approach to be valid, we need to show that the outputs of Alloc do not change too quickly with the input, so that it only takes polynomial time for the binary search to find a $\tilde{u}^*$ that is sufficiently close to u^* for $X(\tilde{u}^*)$ to be within a distance of δ to $X(u^*)$. This is not a trivial task, as Δ_j changes very quickly with c_j when c_j is close to 1 (i.e., $\frac{\partial \Delta_j}{\partial c_j} \to \infty$ when $c_j \to 1$). We observe that, since the input parameters are integers, the problem we deal with is essentially discrete, and the gap between c_j and 1 can be lower-bounded by an exponentially small term. Given this, the outputs can be proven to change at an exponential rate only, so binary search is able to find a solution in polynomial time. □

Lemma 6.2. *Let θ^* be such that* $\mathrm{surp}^\#(\theta^*) = 0$*, and $X^* = X(\theta^*)$. For any given $\delta > 0$, X^* can be approximated to precision δ in time polynomial in the size of the input of* Alloc$^\#$ *and* $\log \frac{1}{\delta}$.

Proof Sketch. The proof is similar to the proof of Lemma 6.1: in polynomial time, binary search can find a $\tilde{\theta}^*$ sufficiently close to θ^* such that $X(\tilde{\theta}^*)$ approximates X^* to precision δ. □

Theorem 6.3. *For any $\delta > 0$, a 0^+-NSE can be approximated to precision δ in time polynomial in the size of the parameters of the game and* $\log \frac{1}{\delta}$.

Proof. The following algorithm finds a 0^+-NSE with precision δ and runs in polynomial time.

Step 1. Compute all intersection points $\mathcal{U}$ as defined in Eq. (10). This can be done in time $\mathrm{poly}(n, m, \log M)$ by enumerating all the triples (i, j, j'), where M is the bound of the payoff parameters. We assume an arithmetic operation on two rational numbers to take time polynomial in the size of their binary representation.

Step 2. Approximate $X^{(\ell)}$ and $h^{(\ell)}$. First, compute the preference profiles $J^{(\ell)}$. Since all points in the same interval $(u^{(\ell)}, u^{(\ell+1)})$ induce the same preference, we take an arbitrary $v \in (u^{(\ell)}, u^{(\ell+1)})$ and compute $J^{(\ell)}$ by sorting $U_i^{\mathrm{d}}(\overline{\mathrm{cov}}(v), 1), \dots, U_i^{\mathrm{d}}(\overline{\mathrm{cov}}(v), m)$. Then fix the preferences to $J^{(\ell)}$ and compute the zero point of surp, with which $X^{(\ell)}$ and $h^{(\ell)}$ can be obtained using Alloc. In particular, we use binary search to approximate the zero point, and hence obtain approximations of $X^{(\ell)}$ and $h^{(\ell)}$. By Lemma 6.1, this can be done in time polynomial in the size of the input parameters and $\log \frac{1}{\delta}$ for any desired precision δ.

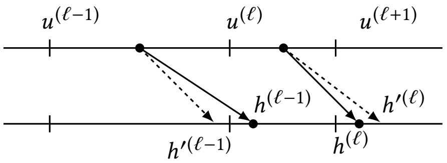

Figure 2: The solid arrows point to the exact values and the dashed arrows to the approximated ones. $u^{(\ell)}$ is falsely detected as a fixed point as $u^{(\ell)} \notin [h^{(\ell-1)}, h^{(\ell)}]$ but $u^{(\ell)} \in [h'^{(\ell-1)}, h'^{(\ell)}]$.

Step 3. Find out the stationary point. We enumerate all ℓ to find either an $h^{(\ell)} \in [u^{(\ell)}, u^{(\ell+1)}]$ or a $u^{(\ell)} \in (\!|h^{(\ell-1)}, h^{(\ell)}|\!)$. If we find an $h^{(\ell)} \in [u^{(\ell)}, u^{(\ell+1)}]$, we are done, as $X^{(\ell)}$ is within a distance of δ from a 0^+-NSE. Otherwise, we need an additional step to find a value $\theta^* \in [0, 1]$ such that $\mathrm{surp}^\#(\theta^*) = 0$ in order to generate a strategy profile with height $u^{(\ell)}$. Again, we use binary search to approximate θ, which, for any desired precision $\delta > 0$, takes time polynomial in the input size and $\log \frac{1}{\delta}$ according to Lemma 6.2.

Note that there might be an issue of false detection due to the approximation error, if the actual value of $h^{(\ell)}$ and the approximated one fall into different intervals (Figure 2). To avoid this, for each $J^{(\ell)}$, we can first fix the interval ℓ such that $\mathrm{surp}(u^{(\ell)}) \ge 0$ and $\mathrm{surp}(u^{(\ell+1)}) \le 0$, and search within this interval only. This ensures that the approximated value is always in the correct interval. □

7 CONCLUSION

This work can be seen as a starting point of an ambitious research agenda aimed at understanding security games with heterogeneous defenders. We list a few research directions below, which in our opinion should be pursued in the immediate future; however, essentially all research questions that have been considered in the context of single-defender security games have a natural counterpart in the context of multi-defender games.

Perhaps the most immediate question is to analyze the 'price of anarchy' in multi-defender games, i.e., to understand the loss in protection caused by the heterogeneity of defenders' interests. Further, it would be interesting to consider settings where groups of defenders can cooperate by coordinating their actions, and study the benefits of such cooperation and ways to promote it. Another possibility is to allow for other forms of utility definitions, e.g., where defenders divide rewards or penalties. We would also like to extend our analysis to schedule-based games, and to more realistic attacker behavior models (e.g., quantal response).

ACKNOWLEDGMENTS

Jiarui Gan would like to acknowledge the funding received towards his study from the EPSRC International Doctoral Scholars Grant EP/N509711/1 and his department.

[8] We provide an example in the full version of the paper.

REFERENCES

[1] Bo An. 2017. Game theoretic analysis of security and sustainability. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI'17)*. 5111–5115.

[2] Nicola Basilico, Andrea Celli, Giuseppe De Nittis, and Nicola Gatti. 2017. Computing the team-maxmin equilibrium in single-team single-adversary team games. *Intelligenza Artificiale* 11, 1 (2017), 67–79.

[3] Hau Chan, Michael Ceyko, and Luis E. Ortiz. 2012. Interdependent defense games: modeling interdependent security under deliberate attacks. In *Proceedings of the 28th Conference on Uncertainty in Artificial Intelligence (UAI'12)*. 152–162.

[4] Stefano Coniglio, Nicola Gatti, and Alberto Marchesi. 2017. Pessimistic leader-follower equilibria with multiple followers. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI'17)*. 171–177.

[5] Vincent Conitzer and Tuomas Sandholm. 2006. Computing the optimal strategy to commit to. In *Proceedings of the 7th ACM Conference on Electronic Commerce (ACM EC'06)*. 82–90.

[6] Victor DeMiguel and Huifu Xu. 2009. A stochastic multiple-leader Stackelberg model: analysis, computation, and application. *Operations Research* 57, 5 (2009), 1220–1235.

[7] Qingyu Guo, Bo An, Yevgeniy Vorobeychik, Long Tran-Thanh, Jiarui Gan, and Chunyan Miao. 2016. Coalitional security games. In *Proceedings of the 2016 International Conference on Autonomous Agents and Multiagent Systems (AAMAS'16)*. 159–167.

[8] Albert Xin Jiang, Ariel D Procaccia, Yundi Qian, Nisarg Shah, and Milind Tambe. 2013. Defender (mis) coordination in security games.. In *Proceedings of the 23rd International Joint Conference on Artificial Intelligence (IJCAI'13)*. 220–226.

[9] Christopher Kiekintveld, Manish Jain, Jason Tsai, James Pita, Fernando Ordóñez, and Milind Tambe. 2009. Computing optimal randomized resource allocations for massive security games. In *Proceedings of The 8th International Conference on Autonomous Agents and Multiagent Systems (AAMAS'09)*. 689–696.

[10] Natalie Klein. 2017. Can International Litigation Solve the India-Sri Lanka Fishing Dispute? *The Diplomat* (July 2017). Available at: http://thediplomat.com/2017/07/can-international-litigation-solve-the-india-sri-lanka-fishing-dispute/.

[11] Ankur A Kulkarni and Uday V Shanbhag. 2014. A shared-constraint approach to multi-leader multi-follower games. *Set-Valued and Variational Analysis* 22, 4 (2014), 691–720.

[12] Aron Laszka, Jian Lou, and Yevgeniy Vorobeychik. 2016. Multi-defender strategic filtering against spear-phishing attacks.. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI'16)*. 537–543.

[13] George Leitmann. 1978. On generalized Stackelberg strategies. *Journal of Optimization Theory and Applications* 26, 4 (1978), 637–643.

[14] Sven Leyffer and Todd Munson. 2010. Solving multi-leader–common-follower games. *Optimisation Methods & Software* 25, 4 (2010), 601–623.

[15] Jian Lou, Andrew M Smith, and Yevgeniy Vorobeychik. 2017. Multidefender Security Games. *IEEE Intelligent Systems* 32, 1 (2017), 50–60.

[16] Jian Lou and Yevgeniy Vorobeychik. 2015. Equilibrium analysis of multi-defender security games. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI'15)*. 596–602.

[17] Jian Lou and Yevgeniy Vorobeychik. 2016. Decentralization and security in dynamic traffic light control. In *Proceedings of the Symposium and Bootcamp on the Science of Security*. ACM, 90–92.

[18] Praveen Paruchuri, Jonathan P. Pearce, Janusz Marecki, Milind Tambe, Fernando Ordóñez, and Sarit Kraus. 2008. Efficient algorithms to solve Bayesian Stackelberg games for security applications. In *Proceedings of the 23rd AAAI Conference on Artificial Intelligence (AAAI'08)*. 1559–1562.

[19] Praveen Paruchuri, Jonathan P Pearce, Janusz Marecki, Milind Tambe, Fernando Ordonez, and Sarit Kraus. 2008. Playing games for security: An efficient exact algorithm for solving Bayesian Stackelberg games. In *Proceedings of the 7th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS'08)*. 895–902.

[20] Praveen Paruchuri, Jonathan P Pearce, Milind Tambe, Fernando Ordonez, and Sarit Kraus. 2007. An efficient heuristic approach for security against multiple adversaries. In *Proceedings of the 6th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS'07)*. 311–318.

[21] Hanif D. Sherali. 1984. A multiple leader stackelberg model and analysis. *Operation Research* 32, 2 (1984), 390–404.

[22] Ankur Sinha, Pekka Malo, Anton Frantsev, and Kalyanmoy Deb. 2014. Finding optimal strategies in a multi-period multi-leader-follower Stackelberg game using an evolutionary algorithm. *Computers & Operation Research* 41 (2014), 374–385.

[23] Andrew Smith, Yevgeniy Vorobeychik, and Joshua Letchford. 2014. MultiDefender security games on networks. *ACM SIGMETRICS Performance Evaluation Review* 41, 4 (2014), 4–7.

[24] Bernhard von Stengel and Shmuel Zamir. 2004. Leadership with commitment to mixed strategies. (2004).

[25] Milind Tambe. 2011. *Security and game theory: algorithms, deployed systems, lessons learned*. Cambridge University Press.

[26] The Economist. 2015. Malacca buccaneers. *The Economist (Jun 27th 2015)* (June 2015). Available at: https://www.economist.com/news/asia/21656237-step-aside-somalia-south-east-asia-new-piracy-capital-world-malacca-buccaneers.

[27] Jason Tsai, Yundi Qian, Yevgeniy Vorobeychik, Christopher Kiekintveld, and Milind Tambe. 2013. Bayesian security games for controlling contagion. In *Proceedings of the ASE/IEEE International Conference on Social Computing(SocialCom)*. IEEE, 33–38.

[28] Zhen Wang, Yue Yin, and Bo An. 2016. Computing optimal monitoring strategy for detecting terrorist plots. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI'16)*. 637–643.