

# Robust Solutions for Multi-Defender Stackelberg Security Games\*

Dolev Mutzari Yonatan Aumann Sarit Kraus

Department of Computer Science, Bar Ilan University, Ramat Gan,  
Israel  
dolevmu@gmail.com, aumann@cs.biu.ac.il, sarit@cs.biu.ac.il

## Abstract

Multi-defender Stackelberg Security Games (MSSG) have recently gained increasing attention in the literature. However, the solutions offered to date are highly sensitive, wherein even small perturbations in the attacker’s utility or slight uncertainties thereof can dramatically change the defenders’ resulting payoffs and alter the equilibrium. In this paper, we introduce a robust model for MSSGs, which admits solutions that are resistant to small perturbations or uncertainties in the game’s parameters. First, we formally define the notion of robustness, as well as the robust MSSG model. Then, for the non-cooperative setting, we prove the existence of a robust approximate equilibrium in any such game, and provide an efficient construction thereof. For the cooperative setting, we show that any such game admits a robust approximate  $\alpha$ -core, provide an efficient construction thereof, and prove that stronger types of the core may be empty. Interestingly, the robust solutions can substantially increase the defenders’ utilities over those of the non-robust ones.

## 1 Introduction

Stackelberg Security Games (SSGs) have attracted much attention in multi-agent community [Paruchuri *et al.*, 2008b, Tambe, 2011, Nguyen *et al.*, 2013, An and Tambe, 2017]. They were applied to a variety of security problems, including the allocation of security resources at the Los Angeles International Airport [Pita *et al.*, 2008], and protecting biodiversity in conservation areas [Basak *et al.*, 2016] (see [Sinha *et al.*, 2018] for an overview). The original SSG model postulates a single defender facing a single attacker. In many applications, however, different targets may be valuable — to varying extents — to multiple disparate parties, each of which is interested in defending its targets of interest. Accordingly, the model of Multi-defender Stackelberg Security Games

---

\*This research has been partly supported by the Israel Science Foundation under grant 1958/20 and the EU Project TAILOR under grant 952215.

(MSSGs), wherein multiple defenders (leaders) protect a set of targets against a strategic attacker (follower) who is their common enemy, has recently gained increasing attention. Solution concepts for this game were developed over time, starting from  $\zeta$ -Nash Equilibrium (NE) [Gan *et al.*, 2018], coordination mechanism [Gan *et al.*, 2020], coalition formation [Mutzari *et al.*, 2021] and correlated equilibrium (CE) via negotiation [Castiglioni *et al.*, 2021].

However, for reasons we will detail shortly, to date, the solutions offered for MSSGs are highly sensitive, wherein small perturbations in the attacker’s utility, or slight uncertainties thereof, can dramatically alter the defenders’ resulting payoffs and alter the equilibrium.

**Tie-Breaking and Discontinuity.** The key, and somewhat subtle, reason for the high sensitivity of the previous solutions has to do with the issue of tie-breaking. Typically, in SSGs, at equilibrium, the attacker has several equally attractive targets. Indeed, otherwise the defender(s) could shift resources from the overly protected targets to increase the protection of the attacker’s chosen target. So, as ties are ubiquitous, the attacker’s tie-breaking behaviour plays a crucial role in the analysis. Different works have considered different tie-breaking policies. In the case of a single defender, it is commonly assumed that the attacker breaks ties in favor of the defender, the reason being that the defender can essentially enforce his choice by shifting an arbitrarily small amount of resources away from his desirable target, absorbing only an arbitrarily small utility loss.

In the case of multiple heterogeneous defenders, however, optimistic tie-breaking is not well-defined (and not justified), as the best target for each defender may be different. Hence, [Gan *et al.*, 2018] use *pessimistic tie-breaking*, wherein each defender acts as if (among its choices) the attacker will attack the target that is *worst* for the defender. So, different defenders act as if the attacker, deterministically, attacks a different target. This tie-breaking rule, however, is not only overly pessimistic and unrealistic (as the attacker cannot simultaneously attack multiple targets), but also results in a sharp non-continuity in the defenders’ payoffs. This is because the pessimistic tie-breaking occurs only at the point of *exactly identical* utilities (of the attacker) for the disparate targets. If the attacker’s utility is slightly perturbed (or if there is even a slight uncertainty with regards to the attacker’s true utilities), then the pessimistic tie-breaking is no longer relevant and the defenders’ payoff may change dramatically. This discontinuity in payoff, in turn, shapes the structure of the solution concepts, which also tend to exhibit a discontinuity at the exact equilibrium point. Indeed, this sensitivity, and discontinuity, is the core reason why, in the existing MSSG models, exact-NE need not exist (see Example D.3).

In practice, however, such sharp discontinuity is often hard to justify. Indeed, while it is standard in game theory to assume that players’ utilities are common knowledge, it is unrealistic to assume that these utilities are known accurately to any level of precision. Evidently, the defenders must somehow *infer* the attacker’s utilities using some combination of data and reasoning, but this data (and reasoning) may be incomplete, imprecise, or noisy. Additionally, the attacker itself may have bounded computational resources, or be subject to other forms of noise, and thus not operate exactly as dictated

by infinite precision calculations. Finally, the attacker’s strategy is also determined by its knowledge/belief of the defender’s coverage strategies, which the attacker needs to somehow infer from possibly incomplete and noisy data.

**Robustness in MSSGs.** Accordingly, as detailed in the related work, in the *single* defender setting, a whole line of research is devoted to addressing uncertainties (see [Nguyen *et al.*, 2014]). Most of these papers tackle uncertainty by applying robust optimization techniques ([Pita *et al.*, 2009, Jiang *et al.*, 2013, Qian *et al.*, 2015, Nguyen *et al.*, 2014]). Some works consider a Bayesian approach [Kiekintveld *et al.*, 2011, Yin and Tambe, 2012, Yang *et al.*, 2012, Amin *et al.*, 2016]. We believe, however, that this avenue is ill-suited for the multi-defender setting, with its tie-breaking subtlety. Robust optimization takes a *worst-case* approach: if some parameters are only known to lay within some set, then one assumes — pessimistically — that these parameters obtain values that minimize the objective function. In the multi-defender setting, however, there is no *one* worst case; the worst-case for one defender may well be good for another. So, worst-case robustness is not well defined in this setting. We note that one could possibly propose a multi-player worst case analysis wherein each defender acts as if the parameters are worst for itself. But, coupled with pessimistic tie-breaking, this would result in an unreasonably, doubly pessimistic and unrealistic perspective, wherein different defenders postulate attackers with different parameters, attacking different targets; as if they inhabit parallel universes. The multi-defender setting therefore calls for a different approach.

**Our Contribution.** In this paper, we offer a robust model for the analysis of MSSGs. First, we formally define the notion of a *robust solution*, which formalizes the notion that the game solutions (e.g. Nash equilibrium, core) remain valid even after small perturbations or uncertainties in the game’s parameters. We then introduce a formal model wherein small perturbations in the attacker’s utility result in only small changes in the attacker’s expected behavior, and hence also in that of the defenders. Essentially, we model the attacker’s behavior as being probabilistic, with a continuous distribution concentrated around the behavior dictated by the presumed (and possibly inaccurate) utility. Importantly, we do not suppose any specific form for this distribution, only that it is concentrated as stated. Thus, this one model captures and unifies many possible scenarios and sources of noise and uncertainty.

Once we have formally defined the notions of robustness and the robust MSSG model, we show that the robust model indeed offers robust solutions. For the non-cooperative setting, we provide an efficient algorithm for constructing a robust approximate NE. For the cooperative setting, we consider the core of the game, for different variants of the core  $(\alpha, \gamma)$ . For the  $\alpha$  version, we prove that the robust approximate core is always non-empty, and give an efficient algorithm for finding solutions therein. Importantly, since the utility of the defenders is no longer determined by their pessimistic beliefs, the resulting core solutions allow greater utility for the defenders than in previous models. Finally, we show that the  $\gamma$ -core may be empty.

We note that an additional benefit of the model, besides robustness, is that it renders moot the entire issue of tie-breaking — optimistic or pessimistic. With a continuous

probability distribution, there is zero probability for ties.

**Selected Related Work.** Robust analysis of the single-defender case has attracted considerable attention. Many of these employ robust optimization techniques to address uncertainties, while others model uncertainties using Bayesian models. Robust optimization is employed by: [Kiekintveld *et al.*, 2013] in a setting where the attacker’s utilities are known to lay in some interval, but the exact value is not known; and by [Pita *et al.*, 2009] in a settings wherein the attacker may exhibit bounded rationality; by [Jiang *et al.*, 2013] in a setting wherein the attacker’s type is only known to be monotone. [Qian *et al.*, 2015] consider a worst-case approach for studying risk-averse attackers, for which the level of risk-adverseness is unknown. Finally, [Nguyen *et al.*, 2014] offer a unified framework and methods for simultaneously addressing multiple types of uncertainties in single defender SSGs using robust optimization. As detailed in the introduction, we argue that the worst-case approach of robust optimization is ill-suited for the multi-defender setting, wherein there is no one worst-case for all defenders.

Other works, still in the single defender setting, employ a Bayesian approach to address uncertainties (as we do for the multi-defender setting). A general analysis of a single defender Bayesian SSG was introduced in [Paruchuri *et al.*, 2008a]. [Kiekintveld *et al.*, 2011] study Bayesian SSGs with a continuous payoff distribution for the attacker. They demonstrate scenarios where there are too many possible attacker types to consider, and known solutions for a finite number of attacker types don’t scale well. [Yin and Tambe, 2012] provide a unified Bayesian approach to handling both discrete and continuous uncertainty in attacker types. [Yang *et al.*, 2012] model human adversaries with bounded rationality as quantal response adversaries, which with some probability do not choose the best response. They specifically assume that the attacker’s strategy is determined by a scaled *soft-max* function of the expected utilities of the attacker on each target, and utilize the specific structure of the *soft-max* function to obtain a solution for the setting (see also [Amin *et al.*, 2016], who consider general SGs).

Our work differs from all of the above in that it considers the multi-defender case. The Bayesian modeling we introduce is also different. On the one hand, unlike the unrestricted, general Bayesian models (e.g. [Paruchuri *et al.*, 2008a, Kiekintveld *et al.*, 2011, Yin and Tambe, 2012]), [Paruchuri *et al.*, 2008a, Yin and Tambe, 2012]), we assume a concentrated distribution (Definition 3.2). This, we believe, adequately models most frequent sources of uncertainty, including noisy and imprecise information, bounded computation power, and most incarnations of bounded rationality. On the other hand, we do not suppose any specific form for this concentrated distribution (e.g. *soft-max*), as such an assumption would fail to model many real-world uncertainties.

## 2 Defining Robustness

The main objective of this work is to develop a robust MSSG framework. We now formally define the notion of robustness. We deliberately do so in the most general terms, not confining the definition to the specific game, or solution concepts that we consider. We later show how the model we offer indeed exhibits robustness as defined here.

**Solution Concepts.** Consider the Nash equilibrium solution concept. Technically, given a game  $G$ , there is a set  $\text{NE}(G)$  of strategy profiles for which the NE property holds. Similarly, the core of  $G$  is a set of coalition structures  $\text{core}(G)$ , for which the core properties hold. Thus, in the most general sense, a *solution concept* is a function  $\mathcal{X}$  from the set of games (of some class) to some space  $S$ , which maps each game  $G$  (of the class) to the corresponding structures for which the specific property of the solution concept holds.

**Nearness.** Conceptually, robustness of a solution  $\mathbf{x}$  states that  $\mathbf{x}$  remains a solution even under “small” perturbations in game parameters. This requires a notion of *distance* amongst games. Given a distance function  $d(\cdot, \cdot)$  over game pairs, say that games  $G, \hat{G}$  are  $\eta$ -near if  $d(G, \hat{G}) \leq \eta$  (later,  $d$  will be instantiated based on the MSSG’s specific parameters).

**Definition 2.1 (Robust Solution).** Let  $\mathcal{X}$  be a solution concept,  $G$  a game, and  $x \in \mathcal{X}(G)$ . We say that  $x$  is an  $\eta$ -robust  $\mathcal{X}$  of  $G$ , if  $x \in \mathcal{X}(\hat{G})$  for any  $\hat{G}$  that is  $\eta$ -near to  $G$ .

Thus, for example, a strategy profile  $\mathbf{x}$  is said to be an  $\eta$ -robust  $\epsilon$ -NE of  $G$  if it is an  $\epsilon$ -NE of any  $\hat{G}$  that is  $\eta$ -near to  $G$  ( $G$  itself included).

### 3 The Robust MSSG Model

**The Standard model.** In an MSSG, there is a set  $\mathcal{N} = \{1, \dots, n\}$  of *defenders*, a set  $\mathcal{T} = \{t_1, \dots, t_m\}$  of *targets* that the defenders wish to protect, and an attacker, who seeks to attack the targets. Each defender  $i \in \mathcal{N}$  has  $k_i \in \mathbb{N}$  *security resources*, each of which can be allocated to *protect* a target. The attacker chooses one target to attack. The attack is *successful* if the target is unprotected by any security resource. A successful (res. unsuccessful) attack at  $t$  yields utility  $r^a(t)$  (res.  $p^a(t)$ ) to the attacker and  $p_i^d(t)$  (res.  $r_i^d(t)$ ) to defender  $i$ , for all  $i$  ( $r_i^d(t) > p_i^d(t)$ ,  $r^a(t) > p^a(t)$ ).

An MSSG is thus a 5-tuple  $G = (\mathcal{N}, \mathcal{T}, \mathcal{K}, \mathcal{R}, \mathcal{P})$ , where  $\mathcal{K} = (k_1, \dots, k_n)$  are the numbers of security resources of the defenders,  $\mathcal{R}$  is the sequence of rewards,  $\mathcal{R} = (r^a(t_1), \dots, r^a(t_m), r_1^d(t_1), \dots, r_n^d(t_m))$ , and similarly  $\mathcal{P}$  is the sequence of penalties.

The defenders’ allocation of security resources to the targets may be randomized. Specifically, each defender  $i$  chooses a *coverage vector*  $\mathbf{x}_i \in \mathcal{C}_{k_i}$  (where  $\mathcal{C}_k = \{x \in [0, 1]^m \mid \sum x_t \leq k\}$ ). Here,  $x_{i,t}$  is the probability that target  $t$  is protected by one of defender  $i$ ’s security resources.<sup>1</sup> We denote by  $\mathbf{X} = (\mathbf{x}_i)_{i \in \mathcal{N}}$ .

If the defenders are uncoordinated, the probability that target  $t \in \mathcal{T}$  is protected is

$$c_t = \text{cov}_t(X) = 1 - \prod_{i \in \mathcal{N}} (1 - x_{i,t}).$$

The *overall coverage vector*  $\mathbf{c} = (c_t)_{t \in \mathcal{T}}$  is assumed to be known to the attacker, which chooses its action after the defenders have committed to their distribution. Thus, the

<sup>1</sup>Provably, any such coverage vector can be implemented by a distribution over deterministic allocation strategies, each employing at most  $k$  resources.

attacker and defenders' utilities upon an attack at  $t$  are:

$$U^a(\mathbf{c}, t) = U^a(c_t, t) = (1 - c_t) \cdot r^a(t) + c_t \cdot p^a(t) \quad (1)$$

$$U_i^d(\mathbf{c}, t) = U_i^d(c_t, t) = c_t \cdot r_i^d(t) + (1 - c_t) \cdot p_i^d(t) \quad (2)$$

In the classic (non-robust) model, all players are assumed to be rational. As such, the attacker's best response is to attack a target in the set  $\text{BR}(\mathbf{c}) := \arg \max_{t \in \mathcal{T}} U^a(\mathbf{c}, t)$ , which maximizes its expected utility. However, as discussed in the introduction,  $\text{BR}(\mathbf{c})$  typically consists of multiple targets, which brings about the issues of tie-breaking and discontinuity.

**The Robust Model.** The non-robustness of the standard MSSG model arises from the assumption of exact deterministic behavior of the attacker, whereby it always plays the *exact* optimal play, even if the difference between the optimal play and the next in line is minuscule. Therefore, small changes in the attacker's utility function may lead to abrupt changes in the attacker's strategy, in turn causing abrupt changes in the defender's utilities. Accordingly, to obtain a robust model, we model the attacker's behavior as being *probabilistic*, with a continuous distribution concentrated around the deterministic behavior dictated by the presumed (possibly inaccurate) utility. Importantly, we seek a general model which can capture the multitude of possible reasons for the mentioned perturbations. The formal details follow.

Given a coverage vector  $\mathbf{c}$ , and the resultant attacker's utility vector  $\mathbf{u} = \mathbf{u}(\mathbf{c}) = (U^a(\mathbf{c}, t))_{t \in \mathcal{T}}$  (over the different targets), the attacker's actual behavior is assumed to be determined by a probability distribution  $\omega(\mathbf{u}) \in [0, 1]^m$ , specifying the probability that the attacker will actually attack each target. The function  $\omega$  is termed the *Attacker's Behavior Function*, and is assumed to have the following properties:

**Definition 3.1 (Attacker's Behaviour Function (ABF)).** A continuously differentiable function  $\omega : \mathbb{R}_+^m \rightarrow \mathbb{R}_+^m$  is an attacker's behaviour function if the following axioms hold:

1.  $\omega(\mathbf{u})$  is a probability distribution.
2.  $\omega$  is monotone increasing at each coordinate.
3. For each  $\emptyset \neq S \subseteq \mathcal{T}$  and  $t \in S : \frac{\omega_t(\mathbf{u})}{\sum_{t' \in S} \omega_{t'}(\mathbf{u})}$  is independent of any  $u_{t'}, t' \notin S$ .

Axiom 3 states that, given that the attack is within the set  $S$ , the conditional probability of an attack on any specific target  $t \in S$  is only determined by the inter-relationships between the utilities of targets within  $S$ .

Thus, given a coverage vector  $\mathbf{c}$ , inducing attacker utilities  $\mathbf{u}(\mathbf{c})$ , and ABF  $\omega$ , defender  $i$ 's expected utility is given by:

$$U_i^d(\mathbf{c}) = \sum_{t \in \mathcal{T}} U_i^d(\mathbf{c}, t) \omega_t(\mathbf{u}(\mathbf{c})) \quad (3)$$

The following definition provides that the attacker's behavior is centered around the optimal deterministic one.

**Definition 3.2** ( $(\delta, \epsilon)$ -ABF). *Let  $G$  be an (M)SSG, and  $\delta, \epsilon > 0$ . An ABF  $\omega$  is a  $(\delta, \epsilon)$ -ABF if:*

$$u_{t'} < u_t - \delta \Rightarrow \omega_{t'}(\mathbf{u}) < \epsilon \quad (4)$$

Thus, with a  $(\delta, \epsilon)$ -ABF, any target offering a utility  $\delta$  less than the optimal one will be attacked with probability  $< \epsilon$ . In Appendix D.1, we show that one example of an ABF is the softmax function  $\sigma(x)$ , and that for any  $\delta, \epsilon$ , there exists a constant  $d$ , such that  $\sigma(d \cdot)$  is a  $(\delta, \epsilon)$ -ABF.

We believe this  $(\delta, \epsilon)$ -ABF formulation captures, under one unifying definition, the many possible sources for noise and uncertainty in the attacker’s utility, as discussed in the introduction. With these definitions, a *robust MSSG* is a pair  $\mathcal{G} = \langle G, \omega \rangle$ , where  $G$  is a MSSG and  $\omega$  is a  $(\delta, \epsilon)$ -ABF.<sup>2</sup>

## 4 Solution Concepts

For completeness, we review the MSSG solution concepts, as defined in previous works. Throughout, we consider the approximate versions.

**$\zeta$ -Nash Equilibrium.** Given the players’ utility functions, the definitions of approximate NE is standard:

**Definition 4.1** ( $\zeta$ -Nash Equilibrium). *A strategy profile  $\mathbf{X} = (\mathbf{x}_i)_{i \in \mathcal{N}}$  is an  $\zeta$ -NE if for any defender  $i \in \mathcal{N}$  and any strategy  $\mathbf{x}'_i$  of  $i$ :*

$$U_i^d(\text{cov}(X)) \geq U_i^d(\text{cov}(\langle \mathbf{x}'_i, X_{-i} \rangle)) - \zeta. \quad (5)$$

**Coalitions.** When the defenders are uncoordinated, independent choices may result in inefficient resource use. Therefore, [Mutzari *et al.*, 2021] consider a model for coalition formation in MSSG, which we adopt. Any subset  $P \subseteq \mathcal{N}$  of defenders may form a *coalition*, in which case they act as a single defender with  $k_P = \sum_{i \in P} k_i$  resources. The coalition consisting of all of the defenders is called the *grand coalition*.

Coalitions partition the set of defenders:  $\mathcal{N} = \{P_1, \dots, P_\ell\}$ . Each coalition  $P_i$  chooses a coverage vector  $\mathbf{x}_j \in \mathcal{C}_{k_{P_j}}$ . Denoting  $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_\ell)$ , the probability that target  $t$  is protected is:

$$c_t = \text{cov}_t(\mathbf{X}) := 1 - \prod_{i=1}^{\ell} (1 - x_{i,t}). \quad (6)$$

Given a strategy profile  $\mathbf{X}$  and a target  $t$ , the defenders’ utilities are defined in the same way as in (3). A *coalition structure* is any such pair  $\mathcal{CS}(\mathcal{P}, \mathbf{X})$ .

Since  $\mathcal{CS}$  includes a coverage vector  $\mathbf{X}$  that (by (6)) induces the overall coverage vector  $\text{cov}(X)$ , for notational simplicity, the players’ utilities can be viewed as a function of the coalition structure, which we simply denote  $U_i^d(\mathcal{CS})$ .

<sup>2</sup>We note that inevitably the proper values for  $\delta, \epsilon$  are dependent on  $G$ . The reason is that  $\delta$  must be small compared to the attacker’s utilities, thus must be scaled in accordance, and  $\epsilon$  must depend on the number of targets as it has to be small relative to  $1/m$ .

**Deviations.** Given a coalition structure  $\mathcal{CS}$ , a subset of defenders may choose to deviate and form a new coalition. In such a case, one must define the strategy played by all defenders. Following [Moulin and Peleg, 1982, Chalkiadakis *et al.*, 2011, Shapley, 1973], the assumption is that in order to protect the status quo, the remaining defenders take revenge against the deviators, employing a strategy that makes at least one not gain by the deviation. It is assumed that this revenge strategy is chosen *after* the deviators choose their strategy. The deviation is *successful* if no such revenge is possible.

**Definition 4.2 ( $\zeta$ -successful deviation).** A deviation  $(D, \mathbf{x}_D)$  from a coalition structure  $\mathcal{CS}$  is  $\zeta$ -successful if for any deviator  $i \in D$  and for any revenge strategy  $\mathbf{x}_R$ ,  $U_i^d(\text{cov}(\langle \mathbf{x}_D, \mathbf{x}_R \rangle)) > U_i^d(\text{cov}(\mathcal{CS})) + \zeta$ .

**Definition 4.3 ( $\zeta$ -approximate core).** A coalition structure  $\mathcal{CS}$  is in the  $\zeta$ -approximate  $\alpha$ -core if it admits no  $\zeta$ -approximate successful deviation.

So, in the  $\zeta$ -approximate core, a deviation cannot add more than  $\zeta$  to the utility of at least one of the deviation's members.

## 5 The Robustness Theorem

We now provide the central robustness theorem, which essentially states that in the robust MSSG model, any approximate-NE or approximate-core automatically translates to their respective *robust* counterparts. We note that, technically, the theorem essentially follows directly from the continuity of  $\omega$ , but the result is exactly what is necessary.

First, recall that the definition of a robust solution (Definition 2.1) requires a distance function between games. In MSSG, for  $G = (\mathcal{N}, \mathcal{T}, \mathcal{K}, \mathcal{R}, \mathcal{P})$ ,  $\hat{G} = (\mathcal{N}, \mathcal{T}, \mathcal{K}, \hat{\mathcal{R}}, \hat{\mathcal{P}})$ , and  $\mathcal{G} = (G, \omega)$ ,  $\hat{\mathcal{G}} = (\hat{G}, \hat{\omega})$ , we define:

$$d(\mathcal{G}, \hat{\mathcal{G}}) = \max\{\|\mathcal{R} - \hat{\mathcal{R}}\|_\infty, \|\mathcal{P} - \hat{\mathcal{P}}\|_\infty, \|\omega - \hat{\omega}\|_\infty\}$$

The following lemma bounds the change in utility arising from small changes in the game parameters.

**Lemma 5.1.** Let  $\mathcal{G}$  be a robust MSSG and let  $\eta > 0$ . Then there exists a constant  $b$  such that, for any coverage vector  $\mathbf{c}$  and any game  $\hat{\mathcal{G}}$   $\eta$ -near to  $\mathcal{G}$ , the following holds:

$$|\hat{U}_i^d(\mathbf{c}) - U_i^d(\mathbf{c})| \leq \frac{b}{2}\eta,$$

where  $\hat{U}_i^d(\mathbf{c})$ ,  $U_i^d(\mathbf{c})$  are the utilities of defender  $i$  in games  $\hat{\mathcal{G}}$ ,  $\mathcal{G}$  respectively.

The proof, which essentially follows from the continuity of the ABF is provided in the Appendix A.

**Theorem 5.2.** Let  $\mathcal{G}$  be a MSSG. With  $b$  of Lemma 5.1, for any  $\eta$ :

1. any  $\zeta$ -NE of  $\mathcal{G}$  is also an  $\eta$ -robust  $(\zeta + b\eta)$ -NE of  $\mathcal{G}$ .



2. any  $\zeta$  core of  $\mathcal{G}$  is also an  $\eta$ -robust  $(\zeta + b\eta)$ -core of  $\mathcal{G}$  (for all three variants of the core  $\alpha, \beta, \gamma$ ).

*Proof.* We prove (1), and the proof of (2) is analogous. Let  $\mathbf{X}$  be a  $\zeta$ -NE of  $\mathcal{G}$ , and  $\hat{\mathcal{G}}$   $\eta$ -near to  $\mathcal{G}$ . Then, for any possible  $\mathbf{x}'_i$  of player  $i$

$$\hat{U}_i^d(\text{cov}(\langle \mathbf{x}'_i, \mathbf{X}_{-i} \rangle)) \geq U_i^d(\text{cov}(\langle \mathbf{x}'_i, \mathbf{X}_{-i} \rangle)) - b\eta/2 \quad (7)$$

$$\geq U_i^d(\text{cov}(\langle \mathbf{x}_i, \mathbf{X}_{-i} \rangle)) - \zeta - b\eta/2 \quad (8)$$

$$\geq \hat{U}_i^d(\text{cov}(\langle \mathbf{x}_i, \mathbf{X}_{-i} \rangle)) - \zeta - b\eta \quad (9)$$

where we use Lemma 5.1 for (7) and (9), and (8) follows by the definition of  $\zeta$ -NE.  $\square$

Accordingly, in order to find robust solutions in our model it suffices to find regular solutions in the model, and robustness then follows automatically. The remainder of the paper is thus dedicated to finding such solutions.

## 6 Nash Equilibrium in Robust MSSGs

We now show how to compute an approximate NE in robust MSSGs. For simplicity, we restrict the discussion to *non-saturated* games. A robust MSSG  $\mathcal{G}$  is  $\alpha$ -saturated if there exists a coverage  $\mathbf{c} \in \mathcal{C}_{k_{\mathcal{N}}}$  and  $t \in \mathcal{T}$  s.t.  $c_t \geq 1 - \alpha$  and  $\omega_t(\mathbf{c}) \geq \epsilon$ . In a saturated game, there are sufficient resources to induce an attack that is almost surely caught. The results can be extended to saturated games using a similar approach to the one used by [Gan *et al.*, 2018].

**Theorem 6.1.** *There exists a polynomial algorithm, such that for any  $G$  there exist  $A, B, C, \epsilon_0, \delta_0$  s.t. for any  $(\delta, \epsilon)$ -ABF  $\omega$  with  $\epsilon < \epsilon_0, \delta < \delta_0$ , on input  $(G, \omega)$  the algorithm outputs a strategy profile  $\mathbf{X}$  that is a  $\zeta$ -NE, for  $\zeta = B\delta + C\epsilon$  (provided  $(G, \omega)$  is not  $A\delta$ -saturated).*

The exact constants, together with a detailed proof, appear in the Appendix B. The proof builds upon the techniques of [Gan *et al.*, 2018], but requires significant adaptations for our setting. For simplicity, in the main body of the paper we consider the case where all targets carry identical penalties and identical rewards for the attacker (but not the defenders). This case allows us to explain the core elements of the construction and the proof, while avoiding the technicalities. Note that in the explanation we do not seek to obtain the best constants. Better bounds are offered in the complete proof.

For the case we are considering, we can further normalize the attacker's utilities so that  $p^a(t) = 0$  and  $r^a(t) = 1$  for all  $t$ . With this assumption,  $U^a(c, t) = 1 - c_t$ , for all  $c, t$ .

The core procedure underlying the algorithm is ALLOC (Algorithm 1). ALLOC gets as input a parameter  $c$ , and aims to have  $c_t = c$ , for all targets. To this end, the defenders — one by one in order — iterate through the targets, allocating resources until either (i) the  $c$  level is reached, or (ii) their resources are fully depleted. This allocation is performed in Line 4, where  $x_{i,t} = 1 - \frac{1-c}{1-c_t}$  brings the coverage to exactly  $c$ , and  $x_{i,t} = k_i$  depletes the defender's resources. Importantly, each defender considers

---

**Algorithm 1** ALLOC

---

**Input:** value  $c, \bar{c} \in [0, 1]$ ,  $G = (\mathcal{N}, \mathcal{T}, \mathcal{K}, \mathcal{P}, \mathcal{R})$

**Output:** a strategy profile  $\mathbf{X} = (x_{i,t})$

- 1:  $c_t \leftarrow 0$ , for all  $t$
  - 2: **for**  $i \in \mathcal{N}$  **do**
  - 3:   **for**  $t \in \mathcal{T}$  by reverse  $\preceq_i^{\bar{c}}$  precedence order **do**
  - 4:      $x_{i,t} \leftarrow \min\{1 - \frac{1-c}{1-c_t}, k_i\}$
  - 5:      $k_i \leftarrow k_i - x_{i,t}$
  - 6:      $c_t \leftarrow 1 - (1 - c_t)(1 - x_{i,t})$ ;
- 

the targets in *reverse preference order* of its expected utility from the target, assuming some identical coverage level  $\bar{c}$  (with  $\bar{c}$  possibly different from  $c$ ). Specifically, for targets  $t, t'$ , we denote  $t \preceq_i^{\bar{c}} t'$  if  $U_i^d(\bar{c}, t) \leq U_i^d(\bar{c}, t')$ . In ALLOC, defenders iterate through the targets from the least to the highest in the  $\preceq_i^{\bar{c}}$  order.

Clearly, the actual coverage level obtained by ALLOC is determined by  $c$ : if  $c$  is too small then ALLOC may complete without depleting the defenders' resources, and if  $c$  is too big then some targets may have  $c_t < c$ . However, there exists a  $\bar{c}$  for which the algorithm completes with  $c_t = \bar{c}$  for all targets (assuming the game is non-saturated). This  $\bar{c}$  can be approximated to within any level of accuracy by binary search, with multiple runs of ALLOC. Set  $\beta = (1 - \bar{c})/2$ .

Set  $\check{c} = \bar{c} + m\delta/\beta$ . For  $\delta$  sufficiently small,  $\check{c} < 1 - \beta$ . Consider the outcome of running ALLOC with  $c = \check{c}$  and  $\bar{c} = \bar{c}$ . Let  $\hat{\mathbf{X}}$  be the resultant allocation, and for each  $t$ , let  $\hat{c}_t$  be the resultant coverage of target  $t$ . Since  $\check{c} > \bar{c}$ , there will necessarily be targets  $t$  for which  $\hat{c}_t < \bar{c}$ , but for  $\delta$  sufficiently small there will be only one such target. Denote this target  $t^*$ . For all other  $t$ 's,  $\hat{c}_t = \bar{c}$ .

We now argue that  $\hat{\mathbf{X}}$  is a  $\zeta$ -NE for  $\zeta = B\delta + C\epsilon$ , for some constants  $B, C$  dependant only on  $G$ .

In the following, we focus on the attacker's behavior occurring with probability  $\geq \epsilon$ . For ease of exposition, we say that an event is *likely* if it happens with probability  $\geq \epsilon$ .

If all defenders play by  $\hat{\mathbf{X}}$  then an attack is only likely at  $t^*$ . This is since  $\hat{c}_t - \hat{c}_{t^*} \geq \delta$ , and hence  $\hat{u}_{t^*} - \hat{u}_t \geq \delta$ , for all  $t$ , and  $\omega$  is a  $(\delta, \epsilon)$ -ABF. Consider a deviation of defender  $i$ , and let  $\mathbf{c}'$  be the coverage vector resulted from this deviation. We now explain why  $\mathbf{c}'$  cannot "substantially" increase the utility of  $i$ . Let:  $L$  be the targets - aside from  $t^*$  to which  $i$  allocated resources,  $L^+ = L \cup \{t^*\}$ , and  $H = \mathcal{T} \setminus L^+$ . We consider what  $\mathbf{c}'$  can do to attacks on  $H, L$  and  $t^*$ .

**Attacks on  $H$ .** We argue that under  $\mathbf{c}'$  the attacker is unlikely to attack any target of  $H$ . The reason is that since  $i$  did not allocate resources to  $H$ , the only way that it can induce an attack on  $H$  is by increasing the coverage of all targets in  $L^+$  to  $\check{c} + \delta$ . But, since  $c_{t^*} < \bar{c} = \check{c} - m\delta/\beta$ , bringing  $t^*$ 's coverage to  $\check{c} + \delta$  can only be accomplished by taking coverage from the members of  $L$ . In doing so, at least one will result with coverage less than  $\check{c} - \delta$  (the value of  $\check{c}$  was so chosen).

**Attacks on  $L$ .** A deviation *can* induce attacks on  $L$ . However, it cannot provide  $i$  substantially more than it originally obtained, where *substantially* means adding more than  $\delta B$  utility. The reason is that since  $t^*$  is not  $\check{c}$  covered, it must be that  $t^*$  was considered by  $i$  after all elements of  $L$  - or else  $i$  would either bring  $t^*$ 's coverage to  $\check{c}$  or fully deplete its resources. So,  $t \preceq_{\check{c}}^i t^*$ , for all  $t \in L$ . This means that - at coverage level  $\bar{c}$  - attacks on targets of  $L$  offer  $i$  no more utility than attacks on  $t^*$ . So, if  $\bar{c}$  where the coverage, then inducing an attack on  $L$  would offer no gain to  $i$ . In practice, the elements of  $\mathbf{c}'$  and  $\hat{\mathbf{c}}$  are not exactly  $\bar{c}$ , but they are  $O(\delta)$  away. So, since  $U_i^d(c, t)$  is linear in  $c$ , the differences in utility can also be only  $O(\delta)$ . So attacks on  $L$  cannot offer *substantially* more utility than what  $t^*$  initially offered.

**Attack on  $t^*$ .** A deviation can also possibly increase the coverage of  $t^*$ , thus offering more utility if and when attacked. However, one can only add  $O(\delta)$  coverage (while keeping an attack on  $t^*$  likely), so that this addition cannot add more than  $O(\delta)$  utility.

**Putting It All Together.** We obtain that with probability  $1 - \epsilon$  the added utility due to the deviation is at most  $\delta B$ , for some constant  $B$ . With probability  $\epsilon$  the increase can be larger, but clearly bounded by  $C = \max_{j,t} r_j^d(t)$ . So, the utility increase is bounded by  $\zeta = \delta B + \epsilon C$ . This completes the construction of  $\zeta$ -NE. By Theorem 5.2, this strategy is also an  $\eta$ -robust  $(\zeta + b\eta)$ -NE, for any  $\eta$ .

## 7 The Robust Core

In this section we explain the algorithm to construct a robust  $\zeta$ -approximate  $\alpha$ -core. Importantly, we only outline the construction and proof ideas. Full details appear in Appendix C.

**Resistance to Subset Deviations.** With minor adaptations, the ALLOC procedure can be applied to the cooperative setting (wherein probabilities are additive rather than multiplicative). Also, the procedure can easily be configured to accept a target utility level  $u$  as input (rather than target coverage  $c$ ) - see Algorithm 4. As before, repeated calls to ALLOC allow us to find a  $\bar{u}$  and strategy  $\bar{\mathbf{X}}$  such that all targets offer utility  $\bar{u}$  to the attacker (except for those with  $r^a(t) < \bar{u}$ ), using all resources. Next, we re-run ALLOC with  $\tilde{u} = \bar{u} - mO(\delta)$ . Let  $\hat{\mathbf{X}}$  be the resulted strategy profile,  $\hat{\mathbf{c}} = \text{cov}(\hat{\mathbf{X}})$ , and  $t^*$  the target not covered to  $\tilde{u}$ . Then,  $U^a(\hat{\mathbf{c}}, t^*) > \bar{u} + \delta$  and  $U^a(\hat{\mathbf{c}}, t) = \tilde{u}$  for all other  $t$ . So, the attacker is only likely to attack  $t^*$ .

We argue that  $\hat{\mathbf{X}}$  is resilient to any deviation  $\mathbf{x}_D$  of any proper subset  $D \subset \mathcal{N}$  of defenders. Set  $\mathbf{c}' = \text{cov}(\hat{\mathbf{X}}_{-D}, \mathbf{x}_D)$ . Let  $L$  be the targets, aside from  $t^*$ , to which  $D$  allocated resources in  $\hat{\mathbf{X}}$ , and  $L^+ = L \cup \{t^*\}$ . Then, analogously to the NE case, no deviation of  $D$  can induce a likely attack outside  $L \cup \{t^*\}$ . They may, however, be able to alter the attack probabilities within  $L^+$ . Let  $\mathcal{A}$  be the targets that are likely to be attacked under  $\mathbf{x}_D$ . So,  $|U^a(\mathbf{c}', t_i) - U^a(\mathbf{c}', t_j)| < \delta$ , for any  $t_i, t_j \in \mathcal{A}$ . Now, recall that given  $D$  and  $\mathbf{x}_D$ , the remaining defenders -  $\bar{D}$  - can change their strategy. So, using only  $O(\delta)$  resources per target,  $\bar{D}$  can raise the coverage of all but one target

of  $t_0 \in \mathcal{A}$ , so that  $t_0$  offers at least  $\delta$  more utility than *all* other targets. Since  $t_0 \in L^+$ , there exists at least one  $i \in D$  for which  $t_0 \preceq_i^{\text{cov}(\mathbf{X}^*)} t^*$ . So, for this  $i$ , the deviation does not offer any (substantial) gain.

**Resistance to Grand Coalition Deviations.** We now show how transform the above strategy to one resistant to grand coalition deviations. Let  $\hat{u}_i^d$  be the utility of defender  $i$  under  $\hat{\mathbf{X}}$ . Consider the following linear program (with variables  $p_t$ ):

$$\begin{aligned} \max \quad & \sum_{t \in \mathcal{T}} \sum_{i \in \mathcal{N}} p_t \cdot U_i^d(\hat{\mathbf{c}}, t) \\ \text{s.t.} \quad & \sum_{t \in \mathcal{T}} p_t \cdot U_i^d(\hat{\mathbf{c}}, t) \geq \hat{u}_i^d, & i = 1, \dots, n \\ & \sum_{t \in \mathcal{T}} p_t = 1; p_t \geq 0, & t = 1, \dots, m \end{aligned}$$

Let  $\mathbf{p}^* = (p_t^*)_{t=1}^m$  the solution provided by the programs (there necessarily exists a solution, since  $\mathbf{p} = \omega(\mathbf{u}(\hat{\mathbf{c}}))$  is a feasible one). The following proposition (proved in Appendix C.3) provides that with small changes to  $\bar{\mathbf{u}} = (\bar{u}_1, \dots, \bar{u}_n)$  the distribution  $\mathbf{p}^*$  can (essentially) be induced.

**Proposition 7.1.** *For any probability distribution  $\mathbf{p}$  and any  $u > 0$ , there exists a utility vector  $\mathbf{u}$  s.t.  $\|\omega(\mathbf{u}) - \mathbf{p}\|_\infty < m^2 \epsilon$ , and  $u \leq u_t < u + 2m\delta$  for all  $t \in \mathcal{T}$ .*

Let  $\bar{\mathbf{u}}^*$  be the utility vector provided by invoking Proposition 7.1 with  $u = \bar{u}$  and  $\mathbf{p} = \mathbf{p}^*$ . Since  $\bar{\mathbf{u}}$  is implementable (by  $\bar{\mathbf{X}}$ ) and  $\bar{u}_t^* \geq \bar{u}$  for all  $t$ , then  $\bar{\mathbf{u}}^*$  is also implementable, by some strategy  $\bar{\mathbf{X}}^*$ . We now argue that  $\bar{\mathbf{X}}^*$  is an the  $\zeta$ -approximate  $\alpha$ -core, for  $\zeta = A\epsilon + B\delta$  (for some constants  $A, B$  fully determined in the full proof). In the following arguments, interpret all statements “up to at most  $\zeta$  gains”.

To see that  $\bar{\mathbf{X}}^*$  resists grand coalition deviations, first note that we may assume that in any such deviation  $\mathbf{X}'$ , all elements of  $\mathbf{u}' = \mathbf{u}(\text{cov}(\mathbf{X}'))$  are close (within  $O(\delta)$ ) to  $\bar{\mathbf{u}}$ . Otherwise, there must be targets that are more than  $\delta$  apart in their utility, in which case it is possible to shift resources from the those with lesser  $u'_t$  to the higher ones, without substantially changing the attack distribution (by Proposition 7.1). Now, any target where  $u'_t > \bar{u}$  is not substantially better for any deviator, and any target where  $u'_t < \bar{u} - \delta$  is not likely to be attacked. By construction,  $\mathbf{p}^*$  maximizes the sum of defender utilities, subject to each getting at least as in  $\hat{\mathbf{X}}$ . So, it is impossible that *all* defenders simultaneously get even more.

For subset deviations, consider  $D \subset \mathcal{N}$  and suppose they have a deviation  $\mathbf{x}_D$  that guarantees all members of  $D$  more than in  $\bar{\mathbf{X}}^*$  (regardless of how the others play). But  $\bar{\mathbf{X}}^*$  offers essentially as much as the output of the linear program, which, by its constraints, offers each defender as much as  $\hat{\mathbf{X}}$ . So,  $\mathbf{x}_D$  would also constitute a successful deviation from  $\hat{\mathbf{X}}$ , which we proved cannot be. We thus obtain:

**Theorem 7.2.** *There exists a polynomial algorithm, such that for any  $G$  there exist  $A, B, C, \epsilon_0, \delta_0$  s.t. for any  $(\delta, \epsilon)$ -ABF  $\omega$  with  $\epsilon < \epsilon_0, \delta < \delta_0$ , on input  $(G, \omega)$  the algorithm outputs a strategy profile  $\mathbf{X}$  that is a  $\zeta$ -approximate  $\alpha$ -core, for  $\zeta = B\delta + C\epsilon$  (provided that  $(G, \omega)$  is not  $A\delta$ -saturated).*

By Theorem 5.2, for any  $\eta > 0$ , any such solution is also an  $\eta$ -robust  $(\zeta + b\eta)$ -approximate  $\alpha$ -core.

**$\gamma$ -Core.** Unlike the  $\alpha$ -core, the approximate  $\gamma$ -core (see [Chander, 2010], Definition D.6) may be empty.

Player	$t_1$	$t_2$	$t_3$	$t_4$	$t_5$	$t_6$	$k$
Defender 1,2	1, 0	1, 0	1, 0	6, 5	6, 5	6, 5	1
Defender 3,4	6, 5	6, 5	6, 5	1, 0	1, 0	1, 0	1
Attacker	1, 0	1, 0	1, 0	1, 0	1, 0	1, 0	–

Table 1: Example of a MSSG where the approximate  $\gamma$ -core is empty

Basically, defenders 1, 2 can always  $\gamma$ -deviate and get a utility  $\geq 4$  each, as can defenders 3, 4, but no coalition structure can yield all defenders  $\geq 4$  utility. See a full explanation in Appendix D.7.

## References

- [Amin *et al.*, 2016] K. Amin, S. Singh, and M. P. Wellman. Gradient methods for stackelberg security games. In *Proc. of UAI*, pages 2–11, 2016.
- [An and Tambe, 2017] B. An and M. Tambe. *Stackelberg Security Games (SSG) Basics and Application Overview*, pages 485–507. Cambridge U. Press, 2017.
- [Basak *et al.*, 2016] A. Basak, F. Fang, T. H. Nguyen, and C. Kiekintveld. Abstraction methods for solving graph-based security games. In *AAMAS*, pages 13–33, 2016.
- [Castiglioni *et al.*, 2021] M. Castiglioni, A. Marchesi, and N. Gatti. Committing to correlated strategies with multiple leaders. *Artificial Intelligence*, page 103549, 2021.
- [Chalkiadakis *et al.*, 2011] G. Chalkiadakis, E. Elkind, and M. Wooldridge. Computational aspects of cooperative game theory. *Synthesis Lectures on AI and ML*, 5(6), 2011.
- [Chander and Tulkens, 2006] Parkash Chander and Henry Tulkens. Cooperation, stability and self-enforcement in international environmental agreements: a conceptual discussion. Technical report, AgEcon Search, 2006.
- [Chander, 2007] Parkash Chander. The gamma-core and coalition formation. *International Journal of Game Theory*, 35(4):539–556, 2007.
- [Chander, 2010] Parkash Chander. Cores of games with positive externalities. *CORE DP*, 4:2010, 2010.
- [Gan *et al.*, 2018] J. Gan, E. Elkind, and M. Wooldridge. Stackelberg security games with multiple uncoordinated defenders. In *AAMAS*. ACM Press, 2018.

- [Gan *et al.*, 2020] J. Gan, E. Elkind, S. Kraus, and Mi. Wooldridge. Mechanism design for defense coordination in security games. In *AAMAS*, pages 402–410, 2020.
- [Jiang *et al.*, 2013] A. X. Jiang, Thanh H. Nguyen, M. Tambe, and A. D Procaccia. Monotonic maximin: A robust stackelberg solution against boundedly rational followers. In Sajal K. Das, Cristina Nita-Rotaru, and Murat Kantarcioglu, editors, *Decision and Game Theory for Security*, pages 119–139, 2013.
- [Kiekintveld *et al.*, 2011] C. Kiekintveld, J. Marecki, and M. Tambe. Approximation methods for infinite bayesian stackelberg games: Modeling distributional payoff uncertainty. In *AAMAS*, pages 1005–1012, 2011.
- [Kiekintveld *et al.*, 2013] C. Kiekintveld, T. Islam, and Vl. Kreinovich. Security games with interval uncertainty. In *AAMAS*, page 231–238, 2013.
- [Moulin and Peleg, 1982] H Moulin and B Peleg. Cores of effectivity functions and implementation theory. *Journal of Mathematical Economics*, 10(1):115 – 145, 1982.
- [Mutzari *et al.*, 2021] D. Mutzari, J. Gan, and S. Kraus. Coalition formation in multi-defender security games. In *AAAI*, volume 35, pages 5603–5610, 2021.
- [Nguyen *et al.*, 2013] Thanh Hong Nguyen, Rong Yang, Amos Azaria, Sarit Kraus, and Milind Tambe. Analyzing the effectiveness of adversary modeling in security games. In *AAAI*, 2013.
- [Nguyen *et al.*, 2014] Thanh Hong Nguyen, Albert Xin Jiang, and Milind Tambe. Stop the compartmentalization: unified robust algorithms for handling uncertainties in security games. In *AAMAS*, pages 317–324, 2014.
- [Paruchuri *et al.*, 2008a] Praveen Paruchuri, Sarit Kraus, Jonathan P Pearce, Janusz Marecki, Milind Tambe, and Fernando Ordonez. Playing games for security: An efficient exact algorithm for solving bayesian stackelberg games. *Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems*, 2008.
- [Paruchuri *et al.*, 2008b] Praveen Paruchuri, Jonathan P Pearce, Janusz Marecki, Milind Tambe, Fernando Ordonez, and Sarit Kraus. Efficient algorithms to solve bayesian stackelberg games for security applications. In *AAAI*, pages 1559–1562, 2008.
- [Pita *et al.*, 2008] James Pita, Manish Jain, Janusz Marecki, Fernando Ordóñez, Christopher Portway, Milind Tambe, Craig Western, Praveen Paruchuri, and Sarit Kraus. Deployed armor protection: the application of a game theoretic model for security at the los angeles international airport. In *Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems: industrial track*, pages 125–132, 2008.
- [Pita *et al.*, 2009] James Pita, Manish Jain, Fernando Ordóñez, Milind Tambe, Sarit Kraus, and Reuma Magori-Cohen. Effective solutions for real-world stackelberg games: When agents must deal with human uncertainties. In *Proceedings of The 8th*

*International Conference on Autonomous Agents and Multiagent Systems - Volume 1*, AAMAS '09, page 369–376, 2009.

- [Qian *et al.*, 2015] Yundi Qian, B William, and Milind Tambe. Robust strategy against unknown risk-averse attackers in security games. In *AAMAS*, pages 1341–1349, 2015.
- [Shapley, 1973] Shubik Shapley. *Game Theory in Economics - Chapter 6: Characteristic Function, Core and Stable Set*. RAND Corporation, 1973.
- [Sinha *et al.*, 2018] Arunesh Sinha, Fei Fang, Bo An, Christopher Kiekintveld, and Milind Tambe. Stackelberg security games: Looking beyond a decade of success. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence*. IJCAI, 2018.
- [Tambe, 2011] Milind Tambe. *Security and game theory: algorithms, deployed systems, lessons learned*. Cambridge University Press, 2011.
- [Yang *et al.*, 2012] Rong Yang, Fernando Ordonez, and Milind Tambe. Computing optimal strategy against quantal response in security games. In *AAMAS*, pages 847–854, 2012.
- [Yin and Tambe, 2012] Zhengyu Yin and Milind Tambe. A unified method for handling discrete and continuous uncertainty in bayesian stackelberg games. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems*, AAMAS '12, page 855–862, 2012.

# Appendix

## A Proof of Lemma 5.1

*Proof.* The change of utility can be bounded, using the triangular inequality, by three additive factors in the following manner:

$$\begin{aligned}
 & |\hat{U}_i^d(\mathbf{c}) - \hat{U}_i^d(\mathbf{c})| = \\
 & \left| \sum_{t \in \mathcal{T}} \hat{U}_i^d(c_t, t) \cdot \hat{\omega}_t(\hat{\mathbf{u}}(\mathbf{c})) - \sum_{t \in \mathcal{T}} U_i^d(c_t, t) \cdot \omega_t(\mathbf{u}(\mathbf{c})) \right| \leq \\
 & \left| \sum_{t \in \mathcal{T}} \hat{U}_i^d(c_t, t) \hat{\omega}_t(\hat{\mathbf{u}}(\mathbf{c})) - \sum_{t \in \mathcal{T}} U_i^d(c_t, t) \hat{\omega}_t(\hat{\mathbf{u}}(\mathbf{c})) \right| \tag{10} \\
 & + \left| \sum_{t \in \mathcal{T}} U_i^d(c_t, t) \hat{\omega}_t(\hat{\mathbf{u}}(\mathbf{c})) - \sum_{t \in \mathcal{T}} U_i^d(c_t, t) \omega_t(\hat{\mathbf{u}}(\mathbf{c})) \right| \tag{11} \\
 & + \left| \sum_{t \in \mathcal{T}} U_i^d(c_t, t) \omega_t(\hat{\mathbf{u}}(\mathbf{c})) - \sum_{t \in \mathcal{T}} U_i^d(c_t, t) \omega_t(\mathbf{u}(\mathbf{c})) \right| \tag{12}
 \end{aligned}$$

Intuitively, each term corresponds to the contribution of the change defender  $i$ 's parameters, the attacker's behaviour function, and the attacker's reward and penalty parameters, to the overall change of utility.

We now bound each of (10)-(12). For (10):

$$\sum_{t \in \mathcal{T}} |\hat{U}_i^d(c_t, t) - U_i^d(c_t, t)| \cdot \hat{\omega}_t(\hat{\mathbf{u}}(\mathbf{c})) \leq \sum_{t \in \mathcal{T}} \eta \hat{\omega}_t(\hat{\mathbf{u}}(\mathbf{c})) = \eta$$

For (11):

$$\begin{aligned}
 & \sum_{t \in \mathcal{T}} U_i^d(c_t, t) |\hat{\omega}_t(\hat{\mathbf{u}}(\mathbf{c})) - \omega_t(\hat{\mathbf{u}}(\mathbf{c}))| \\
 & \leq \eta \sum_{t \in \mathcal{T}} U_i^d(c_t, t) \leq \eta \sum_{t \in \mathcal{T}} r_i^d(t)
 \end{aligned}$$

For (12) we must first understand how a change in the attacker's utility vector alters the attack distribution of the targets. Consider  $\omega$  restricted to the domain is  $\prod_{t \in \mathcal{T}} [\max(0, p^a(t) - \eta), r^a(t) + \eta]$  which is closed and bounded and therefore compact. Therefore, since  $\omega$  is continuously differentiable,  $|\nabla \omega|$  is continuous and therefore bounded, say by some constant  $K$  (in other words,  $\omega$  is  $K$ -Lipschitz continuous).<sup>3</sup>

Thus, by the mean value theorem,

$$\begin{aligned}
 \|\omega(\hat{\mathbf{u}}(\mathbf{c})) - \omega(\mathbf{u}(\mathbf{c}))\|_\infty & \leq \|\omega(\hat{\mathbf{u}}(\mathbf{c})) - \omega(\mathbf{u}(\mathbf{c}))\|_2 \\
 & \leq K \|\hat{\mathbf{u}}(\mathbf{c}) - \mathbf{u}(\mathbf{c})\|_2 \\
 & \leq K \sqrt{m} \|\hat{\mathbf{u}}(\mathbf{c}) - \mathbf{u}(\mathbf{c})\|_\infty
 \end{aligned}$$

<sup>3</sup>Note that  $K$  can depend on  $\eta$ , but we are only interested in  $\eta$ , which is small relative to the game parameters, so this is negligible.



Therefore, for (12):

$$\begin{aligned}
& \sum_{t \in \mathcal{T}} U_i^d(c_t, t) |\omega_t(\hat{\mathbf{u}}(\mathbf{c})) - \omega_t(\mathbf{u}(\mathbf{c}))| \\
& \leq \sum_{t \in \mathcal{T}} U_i^d(c_t, t) K \sqrt{m} \|\hat{\mathbf{u}}(\mathbf{c}) - \mathbf{u}(\mathbf{c})\|_\infty \\
& \leq \sum_{t \in \mathcal{T}} U_i^d(c_t, t) K \sqrt{m} \eta \leq K \sqrt{m} \sum_{t \in \mathcal{T}} r_i^d(t) \eta
\end{aligned}$$

Putting it all together, we get:

$$|\hat{U}_i^d(\mathbf{c}) - \hat{U}_i^d(\mathbf{c})| \leq \eta(1 + (1 + K\sqrt{m}) \sum_{t \in \mathcal{T}} r_i^d(t))$$

□

## B Approximate Nash Equilibrium

In order to describe our construction in detail, we will first introduce some definitions and notations. First, since all defenders prefer targets to be protected, we expect that in NE the defenders would utilize all of their resources. We call such a strategy profile an *efficient* strategy profile.

**Definition B.1 (Efficient Strategy Profile).** *A strategy profile  $X$  is efficient if for each defender  $i \in \mathcal{N}$ ,  $\sum_{t \in \mathcal{T}} x_{i,t} = k_i$ .*

Indeed, if a strategy profile  $X$  is not efficient, then a defender can increase the coverage of all targets and gain more utility. However, if a target is already fully covered, this may not be possible. Therefore, for simplicity, in the paper we assume that the game is not  $\delta/b$ -saturated, for  $b = \min_{t \in \mathcal{T}} r^a(t) - p^a(t)$ . This follows the ideas from [Mutzari *et al.*, 2021] (they assume the game is canonical, which is equivalent to assuming that the game is not (0,0) saturated).

**Definition B.2 (Height of a Coverage).** *The height of a coverage vector  $\mathbf{c} \in [0, 1]^m$ , denoted  $\text{height}(\mathbf{c})$ , is the optimal attacker utility it yields,  $\text{height}(\mathbf{c}) := \max_{t \in \mathcal{T}} U^a(\mathbf{c}, t)$ .*

**Definition B.3 (Mini-Max Height).** *We denote the min-max attacker height by*

$$\bar{u} = \min_{\mathbf{c} \in \mathcal{C}_{k_{\mathcal{N}}}} \text{height}(\mathbf{c}) = \min_{\mathbf{c} \in \mathcal{C}_{k_{\mathcal{N}}}} \max_{t \in \mathcal{T}} U^a(\mathbf{c}, t)$$

Note that in the uncoordinated case scenario, there might not exist any strategy profile  $X$  where  $\text{height}(\text{cov}(X)) = \bar{u}$ , because the lack of cooperation introduces inefficiencies.

Lastly, we formally define the *preference profile* induced by a coverage vector, which captures the heterogeneous preferences of the defenders.

---

**Algorithm 2** ALLOC

---

**Input:**  $u, \tilde{u} \in (\bar{p}, \bar{r})$ ;**Parameter:** a canonical SSG  $G$ ;**Output:** a strategy profile  $X = \langle x_{i,t} \rangle$ ;

- 1: Set  $x_{i,t} \leftarrow 0$ ,  $\Delta_t \leftarrow \overline{\text{cov}}_t(u)$  and  $c_t \leftarrow 0$  for all  $i, t$ ;
  - 2: **for**  $i \in G.N$  **do**
  - 3:   **for**  $t$  in  $G.J_i(\tilde{u}) = t_{i1}, \dots, t_{im}$  **do**
  - 4:      $x_{i,t} \leftarrow \min \{ \Delta_t, G.k_i - \sum_{t' \in \mathcal{T}} x_{i,t'} \}$ ;
  - 5:      $c_t \leftarrow 1 - (1 - c_t)(1 - x_{i,t})$ ;
  - 6:     **if**  $c_t = 1$  **then**  $\Delta_t \leftarrow 0$  **else**  $\Delta_t \leftarrow \frac{\overline{\text{cov}}_t(u) - c_t}{1 - c_t}$ ;
  - 7:     **if**  $x_{i,t} > 0$  **then**  $t^* \leftarrow t$ ;
  - 8: **return**  $X, t^*$ ;
- 

**Definition B.4 (Induced Preference Profile).** A preference profile  $J$  induced by a coverage vector  $\mathbf{c}$  is a list of binary relations  $\langle \preceq_i^J \rangle_{i \in \mathcal{N}}$  such that  $t_1 \preceq_i^J t_2$  iff  $U_i^d(\mathbf{c}, t_1) \leq U_i^d(\mathbf{c}, t_2)$ . For any  $\bar{p} < u < \bar{r} := \max_{t \in \mathcal{T}} r^a(t)$ , we define  $\overline{\text{cov}}(u)$  to be the coverage vector such that on every target  $t \in \mathcal{T}$ ,  $U^a(\overline{\text{cov}}(u), t) = \min(u, p^a(t))$ . Overloading notations, the preference profile induced at  $u$  is  $J(u) := J(\overline{\text{cov}}(u))$ .

We are now ready to describe [Gan *et al.*, 2018]’s procedure, ALLOC, for canonical games, depicted in Algorithm 2. The algorithm gets as an input an attacker utility value  $u$ . In a greedy fashion, each defender  $i$  in his turn allocates resources to targets according to his induced preference profile  $J(u)$ ,  $t_{i1} \preceq_i t_{i2} \preceq_i \dots \preceq_i t_{im}$ , covering each target so that the attacker’s utility on it gets down to  $u$ , or until he runs out of resources.

The main observation of [Gan *et al.*, 2018] for canonical games can be phrased as follows: there exists an input  $u^*$  such that  $\text{ALLOC}(u^*)$  outputs an efficient strategy profile  $X$  and  $\text{cov}(X) = \overline{\text{cov}}(u^*)$ . They also describe how to find  $u^*$  efficiently. We denote by  $t^*$  the last target that is covered in  $\text{ALLOC}(u^*)$ . Next, if the last target is covered up to height  $u^* - A\zeta$  for  $A$  being dependent regarding the game parameters only, then the resulting strategy profile  $X$  is an  $\zeta$ -NSE, which means it’s an  $\zeta$ -NE assuming pessimistic defender beliefs. Provably, this works because any of the following deviation attempts fail:

1. A defender cannot induce the attacker to attack targets to which he didn’t allocate resources. This is because it would require him to overly cover all of the other targets, which would require too many resources for the efficiency of  $X$ .
2. A defender can induce the attacker to attack a target he covered by reducing its coverage, but by construction such targets give him less utility than  $t^*$ .
3. A defender can increase the coverage of the target  $t^*$ , but it is possible to choose an  $A$  such that this will not give him  $> \zeta$  utility gain.

Although using the same construction ensures that deviation attempts of type 2 and 3 still fail in our model, deviation type 1 becomes possible, since equally covered

---

**Algorithm 3**  $\text{CALC}_\zeta\text{-NE}$ 

---

**Parameter:** (1)  $\delta, \epsilon > 0$ ;  
(2) a  $\delta$ -canonical game  $\mathcal{G} = (G, \omega)$ ;  
**Output:** a  $\zeta$ -NE  $X = \langle x_{i,t} \rangle, \zeta$ ;

```
1: right  $\leftarrow G.\bar{r}$ , left  $\leftarrow G.\bar{p}$ ;  
2: while right - left  $> \delta^2$  do  
3:   u  $\leftarrow$  (left + right) / 2;  
4:   X, t = ALLOC(u,u);  
5:   if  $U_i^d(X, t) > u$  then left  $\leftarrow$  u else right  $\leftarrow$  u;  
6:    $\zeta \leftarrow G.B * \epsilon + G.C * \delta$ ;  
7: return ALLOC(u - G.A *  $\delta$ ,u),  $\zeta$ ;
```

---

targets can both be attacked with significant probability. The lack of pessimistic belief forces us to make some adaptations to [Gan *et al.*, 2018] in order to make it work (see Algorithm 3).

## B.1 Proof of Theorem 6.1

In this subsection we formally proof Theorem 6.1.

*Proof.* We divide into cases. Let  $\bar{p} := \max_{t \in \mathcal{T}} p^a(t)$ , and  $\bar{t}$  will be some target with attacker penalty  $\bar{p}$ . In the first case, calling procedure ALLOC with input  $(u < \bar{p} + \delta, \cdot)$  gives a level coverage. In this case, increasing the attacker utility on target  $\bar{t}$  by  $\delta$  will make any other target be attacked with probability  $< \epsilon$ , and hence  $\bar{t}$  is attacked with probability  $\geq 1 - (m-1)\epsilon$ . This is in contradiction to the assumption that the game is not  $(\delta/b, \epsilon)$ -saturated, for  $b = \min_{t \in \mathcal{T}} (r^a(t) - p^a(t))$ .

Next, since  $\text{cov}(\text{ALLOC}(u, u))$  is continuous with  $u$ , there exists some height  $u^* \geq \bar{p} + \delta$  for which  $X^* = \text{ALLOC}(u^*, u^*)$  forms a level coverage  $\mathbf{c}^* = \text{cov}(X^*)$  of height  $u^*$ . This means that for any target  $t \in \mathcal{T}$ , if  $u^* < p^a(t)$  then  $c_t^* = 0$ , and otherwise  $U^a(\mathbf{c}^*, t) = u^*$ . This  $u^*$  can easily be found using binary search up to any fixed precision.

For the second case, denote by  $a = \max_{t \in \mathcal{T}} (r^a(t) - p^a(t))$ , and  $g := \min_{t \in \mathcal{T}} (u^* - p^a(t)) \geq \delta > 0$ . Note that  $g \geq \delta$  only for this second case. Fix the following values for  $A, B, C$ :

- $A := \frac{4ma^2}{bg}$

- 

$$B := \max_{i \in \mathcal{N}} (r_i^d(t^*) - p_i^d(t^*)) \cdot \left( \frac{8ma^3 \sum_{t \in \mathcal{T}} (r^a(t) - p^a(t))}{bg^2} + A \right)$$

- $C := \max_{i \in \mathcal{N}, t \in \mathcal{T}} r_i^d(t^*) + \max_{i \in \mathcal{N}} \sum_{t \in \mathcal{T}} r_i^d(t)$

We also fix  $\epsilon_0 := 1/m$  and  $\delta_0 := \frac{bg^2}{8ma^2}$  and assume that  $\delta < \delta_0$  and  $\epsilon < \epsilon_0$ . Consider  $X = \text{ALLOC}(u^* - A\delta, u^*)$  and denote by  $\mathbf{c} = \text{cov}(X)$  and by  $\mathbf{u} = (U^a(\mathbf{c}, t))_{t \in \mathcal{T}}$ . We divide further into two cases, where in the first case,  $g > A\delta$ . For  $\delta$  sufficiently small, there is only a single target  $t^*$  that yields an attacker utility greater than  $u^*$ , which is also the last target covered in  $\text{ALLOC}(u^*, u^*)$ . For any other target  $t$ , if  $p^a(t) > u^* - A\delta$  then  $c_t = 0$ , and otherwise  $u^t = u^* - A\delta$ . We argue that  $X$  is a  $\zeta$ -NE for  $\zeta = B\delta + C\epsilon$  for some constants  $B, C$  dependant only on  $G$ .

Since  $A > 1$ , this means that with probability  $\geq 1 - m\epsilon$ , target  $t^*$  is being attacked and each defender gets a utility of  $U_i^d(c_{t^*}, t^*)$ . Let  $i \in \mathcal{N}$  be some defender who deviates with some strategy  $x'_i$ . Let  $X' = \langle X_{-i}, x'_i \rangle$  be the resulted strategy profile and  $\mathbf{c}', \mathbf{u}'$  the resulted coverage vector and attacker utility vector respectively. Let  $S$  be the subset of targets not covered by defender  $i$  in  $X$ , that is,  $t \in S \iff x_{i,t} = 0$ . Note that  $\text{BR}(X) = \{t^*\}$  by construction. Also note that, the targets covered by defender  $i$  are less favorable for him by construction at height  $u^*$ , than target  $t^*$ . That is  $\forall t \notin S : t \preceq_i t^*$ , meaning  $\forall t \notin S : U_i^d(u^*, t) \leq U_i^d(u^*, t^*)$ .

If the target  $t^*$  is the most favorable target for defender  $i$  at height  $u^*$ , then the maximum defender  $i$  can do is increase the coverage of target  $t^*$  while keeping the probability the attacker will attack target  $t^*$  large. Since  $X$  is efficient, meaning utilizes all the resources of all defenders, we know that the resulted height of  $X$  is at most  $\text{height}(X) \geq u^* - A\delta$ , therefore we can bound the utility after such a deviation by  $U_i^d(\text{cov}(u^* - A\delta), t^*) + \mathcal{O}(\epsilon)$ . Formally:

$$\begin{aligned}
U_i^d(X') &\leq U_i^d(\text{cov}(u^* - A\delta), t^*) + \epsilon \sum_{t \in \mathcal{T}} r_i^d(t) \\
&= U_i^d(X, t^*) + \frac{r_i^d(t^*) - p_i^d(t^*)}{r^a(t^*) - p^a(t^*)} (\delta' + A\delta) + \epsilon \sum_{t \in \mathcal{T}} r_i^d(t) \\
&\leq U_i^d(X) + \max_{t \in \mathcal{T}} r_i^d(t) \epsilon + \frac{r_i^d(t^*) - p_i^d(t^*)}{r^a(t^*) - p^a(t^*)} (\delta' + A\delta) \\
&\quad + \epsilon \sum_{t \in \mathcal{T}} r_i^d(t)
\end{aligned} \tag{13}$$

Where  $U^a(X, t^*) - u^* = \delta'$  and the second inequality follows directly from Axiom 4 because  $A\delta + \delta' > \delta$ , which is indeed the case since we take  $A > 1$ , since  $a \geq b$  and  $a \geq g$ .

Next, assume that there are targets in  $S \setminus \{t^*\}$  that are more favorable for defender  $i$ . We will show that for any strategy he takes, in the resulting coverage vector  $\mathbf{c}'$ , there will be at least one target  $t' \notin S$  such that for every target  $s \in S$ ,  $U^a(\mathbf{c}', t') > U^a(\mathbf{c}', s) + \delta$ , and hence all targets in  $t'$  will be attacked with probability at most  $\epsilon$ .

Indeed, assume the contrary. If defender  $i$  tries to minimize

$$\Delta(\mathbf{c}') := \min_{t \notin S, s \in S} U^a(\mathbf{c}', t) - U^a(\mathbf{c}', s),$$

he will allocate all of his resources to  $\mathcal{T} \setminus S$ , resulting with equal attacker utility on all of those targets, utilizing all of his resources. We will denote this utility by  $u_-$ . Indeed, otherwise, he could simply shift resources from the more protected targets to the less ones in  $\mathcal{T} \setminus S$ , and reduce  $\Delta$ . Therefore, we know that:

1.  $\sum_{j \notin S} x_{ij} = \sum_{j \notin S} x'_{ij} = k_i$
2.  $\forall j \in \mathcal{T} \setminus \{t^*\} : 1 - (1 - x_{ij})(1 - c_{-j}) = \frac{r^a(j) - (u^* - A\delta)}{r^a(j) - p^a(j)}$
3.  $1 - (1 - x_{i,t^*})(1 - c_{-t^*}) = \frac{r^a(t^*) - (u^* + \delta')}{r^a(t^*) - p^a(t^*)}$
4.  $\forall j \notin S : 1 - (1 - x'_{ij})(1 - c_{-j}) = \frac{r^a(j) - u_{=}}{r^a(j) - p^a(j)}$

We want to claim that  $u^* + A\delta - u_{=} > \delta$ . Let's simplify:

1.  $\sum_{j \notin S} x_{ij} = \sum_{j \notin S} x'_{ij} = k_i$
2.  $\forall j \in \mathcal{T} \setminus \{t^*\} : 1 - x_{ij} = \frac{1}{1 - c_{-j}} \cdot \frac{u^* - A\delta - p^a(j)}{r^a(j) - p^a(j)}$
3.  $1 - x_{i,t^*} = \frac{1}{1 - c_{-t^*}} \cdot \frac{u^* + \delta' - p^a(t^*)}{r^a(t^*) - p^a(t^*)}$
4.  $\forall j \notin S : 1 - x'_{ij} = \frac{1}{1 - c_{-j}} \cdot \frac{u_{=} - p^a(j)}{r^a(j) - p^a(j)}$

Therefore, subtracting  $\mathbf{c} - \mathbf{c}'$ , we get that:

1.  $\sum_{j \notin S} x_{ij} = \sum_{j \notin S} x'_{ij} = k_i$
2.  $\forall j \notin S \setminus \{t^*\} : x_{ij} - x'_{ij} = \frac{1}{1 - c_{-j}} \cdot \frac{u_{=} - u^* + A\delta}{r^a(j) - p^a(j)}$
3.  $x_{i,t^*} - x'_{i,t^*} = -\frac{1}{1 - c_{-t^*}} \cdot \frac{u^* + \delta' - u_{=}}{r^a(t^*) - p^a(t^*)}$

Next, summing up all the equations in 2 and equation 3, and substituting equation 1, we get:

$$\begin{aligned} & \sum_{j \in \bar{S} \setminus \{t^*\}} \frac{1}{1 - c_{-j}} \cdot \frac{u_{=} - u^* + A\delta}{r^a(j) - p^a(j)} \\ & - \frac{1}{1 - c_{-t^*}} \cdot \frac{u^* + \delta' - u_{=}}{r^a(t^*) - p^a(t^*)} = k_i - k_i = 0 \end{aligned} \quad (14)$$

Next, note that  $\Delta = u_{=} - (u^* - A\delta)$ . Therefore:

$$\Delta \cdot \sum_{j \notin S} \frac{1}{1 - c_{-j}} \cdot \frac{1}{r^a(j) - p^a(j)} = \frac{1}{1 - c_{-t^*}} \cdot \frac{\delta' + A\delta}{r^a(t^*) - p^a(t^*)}$$

Therefore, simplifying we get:

$$\Delta \geq \frac{b(\delta' + A\delta)}{a(1 - c_{-t^*}) \cdot \sum_{j \notin S} \frac{1}{1 - c_{-j}}}$$

Next, we remember that  $0 \leq c_{-j} \leq \frac{r^a(j) - \bar{u} + A\delta}{r^a(j) - p^a(j)}$ , since we covered up to height  $u' = \bar{u} - A\delta$ . Therefore,  $1 \geq 1 - c_{-j} \geq \frac{\bar{u} - A\delta - p^a(j)}{r^a(j) - p^a(j)}$ , and we get  $1 \leq \frac{1}{1 - c_{-j}} \leq \frac{r^a(j) - p^a(j)}{\bar{u} - A\delta - p^a(j)} \leq \frac{a}{g - A\delta} \leq \frac{2a}{g}$  for  $A\delta \leq A\delta_0 = g/2$ . Therefore, we get that:

$$\Delta \geq \frac{bg(\delta' + A\delta)}{2a^2|T|}$$

Now we only need to verify  $A\delta, \delta'$  are large enough so that  $\frac{bg(\delta' + A\delta)}{2ma^2} > \delta$ . We can do that by taking  $A\delta = \frac{4ma^2}{bg}\delta \in \mathcal{O}(\delta)$ . Now that  $\Delta$  is at least  $\delta$ , targets outside of  $S$  are attacked with probability at most  $\epsilon$ .

Lastly, for the case where  $g > A\delta$ , we have to show that  $\delta' \in \mathcal{O}(\delta)$  in order to bound the advantage of deviation. Intuitively, since the height  $u^* + \delta'$  is a continuous function of the height of the rest of the targets, namely  $u^* - A\delta$ , if  $A\delta$  is small enough then  $\delta'$  is also small enough, therefore the utility loss from this NLE solution is  $\mathcal{O}(\delta + \epsilon)$ . We would want to give a concrete bound however, so let's look at the resource count.

In order to decrease an attacker utility by  $A\delta$  one needs to increase the coverage of this target by  $(r^a(t) - p^a(t))A\delta$ . The cost of increasing the coverage by  $(r^a(t) - p^a(t))A\delta$  from a value of  $c_t$  by a single defender is  $\frac{1}{1-c_t}(r^a(t) - p^a(t))A\delta \leq \frac{2a(r^a(t) - p^a(t))}{g}A\delta$ . However, a single defender may not have that amount of resource left by this stage. In worst case, we may assume that all of the defenders allocate this same amount of resource, which will always be enough to decrease the coverage by  $A\delta$ .

Therefore, the maximal resource loss we get from increasing the coverage to height  $u^* - A\delta$  is  $m \cdot \frac{2a \sum_t (r^a(t) - p^a(t))}{g} A\delta$ , which will give an attacker utility decrease of  $\delta' \leq (r^a(t^*) - p^a(t^*)) \frac{2am \sum_t (r^a(t) - p^a(t))}{g} A\delta$ .

For the last part of the proof, we have that  $\delta \leq g < A\delta$ . For this case, calling ALLOC with input  $(u^* - A\delta, u^*)$ , will lead to cases where some of the targets get fully covered and still the height  $u^* - A\delta$  is not reached. Denote the set of those targets by  $M$ . Then by the assumption that the game is not saturated, by definition defender  $i$  cannot induce an attack on targets in  $S \cap M$ , because they will stay fully covered. As for targets in  $S \setminus M$ , again defender  $i$  will not have enough security resources to induce an attack on them. That is because now targets in  $M \setminus S$  free him only less security resources to utilize. □

## C Approximate $\alpha$ -Core

### C.1 Proof of Resistance to Subset Deviations

In this subsection we prove the first step of our construction: resistance to deviations of coalitions other than the grand coalition. The adaption of ALLOC to the cooperative case is provided in Algorithm GC-ALLOC is explicitly described in 4.

**Theorem C.1.** *There exist  $A, B, C, \delta_0, \epsilon_0 > 0$  s.t. for any  $(\delta, \epsilon)$ -ABF  $\omega$  with  $\epsilon < \epsilon_0$ ,  $\delta < \delta_0$ , the coalition structure  $\mathcal{CS}$  generated by Algorithm GC-ALLOC with input  $(u = \bar{u} - A\delta, \bar{u})$ , is resistant to any deviation of a coalition  $D \subset \mathcal{N}$ . That is,  $D$  has no  $\zeta = B\epsilon + C\delta$  successful deviation from  $\mathcal{CS}$ .*

---

**Algorithm 4** GC-ALLOC

---

**Input:**  $u, \tilde{u} \in [\bar{p}, \bar{r}]$ ;**Parameter:** (1) a canonical game  $\langle \mathbf{r}^a, \mathbf{p}^a, \mathbf{r}^d, \mathbf{p}^d, \mathbf{k} \rangle$ ;  
(2)  $\delta, \epsilon > 0$  and a  $(\delta, \epsilon)$ -ABF  $\omega$ ;**Output:** a coalition structure  $\mathcal{CS} = \langle \mathcal{N}, X \rangle$ 

- 1: Set  $x_{i,t} \leftarrow 0$ ,  $\Delta_t \leftarrow \overline{cov}_t(u)$  and  $c_t \leftarrow 0$ , for all  $i, t$ ;
  - 2: **for**  $i \in \mathcal{N}$  **do**
  - 3:   **for**  $t = t_{i1}, \dots, t_{im}$  in  $J_i(\tilde{u})$  **do**
  - 4:      $x_{i,t} \leftarrow \min \{ \Delta_t, k_i - \sum_{t' \in \mathcal{T}} x_{i,t'} \}$ ;
  - 5:      $c_t \leftarrow c_t + x_{i,t}$ ;
  - 6:     **if**  $c_t = 1$  **then**  $\Delta_t \leftarrow 0$  **else**  $\Delta_t \leftarrow \overline{cov}_t(u) - c_t$ ;
  - 7: **return**  $X$ ;
- 

*Proof.* We divide into cases. Let  $\bar{p} := \max_{t \in \mathcal{T}} p^a(t)$ , and  $\bar{t}$  will be some target with attacker penalty  $\bar{p}$ . In the first case, calling procedure ALLOC with input  $(u < \bar{p} + \delta, \bar{u})$  gives a level coverage. In this case, increasing the attacker utility on target  $\bar{t}$  by  $\delta$  will make any other target be attacked with probability  $< \epsilon$ , and hence  $\bar{t}$  is attacked with probability  $\geq 1 - (m-1)\epsilon$ . This is in contradiction to the assumption that the game is not  $(\delta/b, \epsilon)$ -saturated, for  $b = \min_{t \in \mathcal{T}} (r^a(t) - p^a(t))$ .

For the second case, denote by  $a = \max_{t \in \mathcal{T}} (r^a(t) - p^a(t))$ , and  $g := \min_{t \in \mathcal{T}} (u^* - p^a(t)) \geq \delta > 0$ . Note that  $g \geq \delta$  only for this second case. Fix the following values for  $A, B, C, \delta_0, \epsilon_0$ :

- $A := 1 + \frac{\sum_t (r^a(t) - p^a(t))^{-1}}{\min_t (r^a(t) - p^a(t))^{-1}}$
- $B := \max_{i \in \mathcal{N}} (r_i^d(t^*) - p_i^d(t^*)) \cdot [(r^a(t^*) - p^a(t^*))^{-1} + \sum_t (r^a(t) - p^a(t))^{-1} + \frac{(\sum_t (r^a(t) - p^a(t))^{-1})^2}{r^a(t^*) - p^a(t^*)}]$
- $C := \max_{i \in \mathcal{N}} \sum_{t \in \mathcal{T}} r_i^d(t)$
- $\delta_0 = (\sum_t (r^a(t) - p^a(t))^{-1})^{-1}$
- $\epsilon_0 = 1/m$

We divide further into two cases, where in the first case,  $g > A\delta$ . For  $\delta$  sufficiently small, there is only a single target  $t^*$  that yields an attacker utility greater than  $\bar{u}$ , which is also the last target covered in  $\text{ALLOC}(\bar{u} - A\delta, \bar{u})$ . For any other target  $t$ , if  $p^a(t) > \bar{u} - A\delta$  then  $c_t = 0$ , and otherwise  $u^t = \bar{u} - A\delta$ . We argue that  $X$  is a  $\zeta$ -NE for  $\zeta = B\delta + C\epsilon$  for some constants  $B, C$  dependant only on  $G$ .

In contradiction, let  $(D, \mathbf{c}'_D)$  be an  $\epsilon'$ -successful deviation. Let  $T_D$  be the set of targets such that  $\mathbf{c}_D(t) > 0$ , and  $T_R = \overline{T_D}$ . For each target  $t \notin T_D$ , we fix the revenge takers coverage to the value that will give the attacker a utility of  $\bar{u} - \delta$ . As  $\bar{u}$  is the mini-max height, this means that after the deviation targets on  $\overline{T_D}$  will be attacked with probability at most  $\epsilon$ . Indeed, if the contrary is true, then the height of the coverage must be smaller than  $\bar{u}$ , the height of the mini-max coverage, a contradiction.

Next, observe the resulted coverage vector, and define  $T_\epsilon$  to be the set of targets that have an attack probability of at least  $\epsilon$ . Note that  $T_\epsilon \subseteq T_D$  as mentioned. Let  $t_{\text{BR}} \in \text{BR} \subseteq T_\epsilon \subseteq T_D$  be some target.

Next, consider two cases. In the first,  $|T_R| \geq 1$ . In this case, the revengers will use the resources left to make any target in  $T_D$  other than  $t_{\text{BR}}$  have an attacker utility smaller than  $U^a(\mathbf{c}'_D, t_{\text{BR}}) - \delta$ . We will now verify that because of the choice of  $A$ , they have enough resources for the revengers to do that.

For each target  $t \in T_R$ , the coverage  $c_R(t)$  required to get to an attacker utility  $u$  is:  $c_R(t) = \frac{r^a(t) - u}{r^a(t) - p^a(t)}$ . Therefore, since  $(\bar{u} - \delta) - (\bar{u} - A\delta) = \frac{\sum_t (r^a(t) - p^a(t))^{-1}}{\min_t (r^a(t) - p^a(t))^{-1}} \delta$ , the resources that the revengers save from decreasing the coverage of the targets in  $T_R$  is:

$$\Delta K_R^+ = \sum_{t' \in T_R} (r^a(t') - p^a(t'))^{-1} \frac{\sum_t (r^a(t) - p^a(t))^{-1}}{\min_t (r^a(t) - p^a(t))^{-1}} \delta$$

And since there is at least one target in  $T_R$ , it is lower bounded by:

$$\begin{aligned} \Delta K_R^+ &\geq \min_{t'} (r^a(t') - p^a(t'))^{-1} \cdot \frac{\sum_t (r^a(t) - p^a(t))^{-1}}{\min_t (r^a(t) - p^a(t))^{-1}} \delta \\ &= \sum_t (r^a(t) - p^a(t))^{-1} \delta \end{aligned} \quad (15)$$

On the other hand, the amount of resource required to decrease all targets in  $T_D \setminus \{t_{\text{BR}}\}$  to have an attacker utility smaller than  $U^a(\mathbf{c}_D, t_{\text{BR}}) - \delta$ . However, this can be upper bounded by the cost of increasing all of the targets by  $\delta$  attacker utility:

$$\Delta K_R^- \leq \sum_t (r^a(t) - p^a(t))^{-1} \delta$$

Therefore,  $\Delta K_R^+ \geq \Delta K_R^-$ , and there are enough security resources to do so.

In the second case,  $T_R = \emptyset$ , therefore  $T_D = T$ . This means that the deviators have allocated resources to every target. Then once again, if the revengers allocate resources to every target  $t$  other than  $t_{\text{BR}}$  in order to make its attack probability less than  $\epsilon$ , this will require at most  $\Delta K_R^- \leq \sum_t (r^a(t) - p^a(t))^{-1} \delta$  resources, and therefore for  $\delta \leq \delta_0 = (\sum_t (r^a(t) - p^a(t))^{-1})^{-1}$ , this requires  $\Delta K_R^- \leq 1$  resources, which we assume the revengers have.

In any case, we see that the revengers can induce the attacker to attack a single target with probability greater than  $\epsilon$ , and that this target is in  $T_D$ . However, consider a deviator  $i \in D$  who allocated resources to target  $t$ . Therefore, if we denote by  $\mathcal{CS}'$  the coalition structure of the deviation and revenge coalitions described above, then the utility of defender  $i$  can be bounded by:

$$U_i^d(\mathcal{CS}') \leq U_i^d(\mathcal{CS}', t_{\text{BR}}) + \epsilon \sum_t r_i^d(t)$$

However, since  $t_{\text{BR}} \in \text{BR}$  and  $\bar{u}$  is the height of the level coverage, then target  $t_{\text{BR}}$  cannot be covered such that the attacker utility on it is lower than  $\bar{u}$ . Therefore, by construction and since defender  $i$  covered target  $t_{\text{BR}}$  before target  $t^*$ , we have:

$$U_i^d(\mathcal{CS}', t_{\text{BR}}) \leq U_i^d(\overline{\mathcal{CS}}, t^*)$$



where  $\overline{\mathcal{CS}}$  is a level coverage of height  $\bar{u} - \delta$ . Finally:

$$U_i^d(\overline{\mathcal{CS}}, t^*) = U_i^d(\mathcal{CS}, t^*) + \frac{r_i^d(t^*) - p_i^d(t^*)}{r^a(t^*) - p^a(t^*)}(\delta + \delta')$$

Where  $\bar{u} + \delta'$  is the height of  $\mathcal{CS}$ , meaning the utility of the attacker on target  $t^*$  in  $\mathcal{CS}$ . Putting this altogether, we get that:

$$U_i^d(\mathcal{CS}') - U_i^d(\mathcal{CS}) \leq \epsilon \sum_t r_i^d(t) + \frac{r_i^d(t^*) - p_i^d(t^*)}{r^a(t^*) - p^a(t^*)}(\delta + \delta')$$

All that is left is compute  $\delta'$ . We know that  $\mathcal{CS}$  utilizes all  $K$  resources, and so is a level coverage of height  $\bar{u}$ . Therefore:

$$\begin{aligned} K &= \sum_t \frac{r^a(t) - \bar{u}}{r^a(t) - p^a(t)} \\ &= \sum_{t \neq t^*} \frac{r^a(t) - (\bar{u} - \delta')}{r^a(t) - p^a(t)} + \frac{r^a(t^*) - (\bar{u} + \delta')}{r^a(t^*) - p^a(t^*)} \end{aligned} \quad (16)$$

Simplifying this gives:

$$\delta' = \delta' (r^a(t^*) - p^a(t^*)) \sum_{t \neq t^*} (r^a(t) - p^a(t))^{-1}$$

Where  $\delta' = \delta + \frac{\sum_t (r^a(t) - p^a(t))^{-1}}{\min_t (r^a(t) - p^a(t))^{-1}} \delta$

Therefore,

$$\begin{aligned} \delta'/\delta &\leq (r^a(t^*) - p^a(t^*)) \sum_t (r^a(t) - p^a(t))^{-1} \\ &\quad + \left( \sum_t (r^a(t) - p^a(t))^{-1} \right)^2 \end{aligned} \quad (17)$$

Thus, there are no  $\zeta$  successful deviations where  $K_D < D$ , if

$$\begin{aligned} \zeta &\geq \epsilon \sum_t r_i^d(t) \\ &\quad + \delta \max_i (r_i^d(t^*) - p_i^d(t^*)) \cdot \left[ \begin{aligned} &(r^a(t^*) - p^a(t^*))^{-1} + \\ &\sum_t (r^a(t) - p^a(t))^{-1} + \\ &\frac{(\sum_t (r^a(t) - p^a(t))^{-1})^2}{r^a(t^*) - p^a(t^*)} \end{aligned} \right] \end{aligned} \quad (18)$$

This completes the proof for  $g > A\delta$ . Lastly, exactly by the same argument in the proof of Theorem 6.1, there are no  $\zeta$  successful deviations in the case where  $g \leq A\delta$ .  $\square$

## C.2 Resistance to Grand Coalition Deviations

In order to obtain resistance to the grand coalition, by definition, it suffices that the solution is Pareto-efficient (see Definition C.2).

**Definition C.2 (Pareto-efficient Element).** Let  $m \in \mathbb{N}$  and  $S \subseteq \mathbb{R}_+^m$  be some set. A vector  $\mathbf{u} \in S$  is Pareto efficient in  $S$  if there doesn't exist another vector  $\mathbf{u}' \in S$  such that  $u'_t \geq u_t$  on each coordinate and there exist  $1 \leq t \leq m$  s.t.  $u'_t > u_t$ . We denote the subset of Pareto efficient elements by  $S^{PE}$ .

**Lemma C.3.** Let  $\mathcal{U}$  be the set of possible utility vectors for defenders in  $\mathcal{N}$ . That is,  $\mathbf{u} \in \mathcal{U}$  if there exist a coverage vector  $\mathbf{c} \in \mathcal{C}_{k_{\mathcal{N}}}$  such that  $\mathbf{u} = (U_i^d(\mathbf{c}))_{i \in \mathcal{N}}$ . Let  $\mathbf{u} \in \mathcal{U}^{PE}$ , and let  $\mathbf{c} \in \mathcal{C}_{\mathcal{N}}$  be a coverage vector inducing  $\mathbf{u}$ . Then  $\mathcal{CS} = \langle \mathcal{N}, \mathbf{c} \rangle$  is a coalition structure resistant to any successful deviation of the grand coalition  $D = \mathcal{N}$ .

*Proof.* Assume in contradiction that there exist a successful deviation of  $(\mathcal{N}, \mathbf{x}_D)$ . Denote by  $\mathbf{u}' = (U_i^d(\mathbf{x}_D))_{i \in \mathcal{N}}$ . Then by definition of a successful deviation, for every deviator  $i \in \mathcal{N}$ ,  $u'_i > u_i$ , contradicting the Pareto efficiency of  $\mathbf{u}$ , since  $\mathbf{u}' \in \mathcal{U}_{\mathcal{N}}$  as well.  $\square$

## C.3 Inducing Arbitrary Attack Distributions

Reasonably, a coalition structure in the core should still result with a coverage vector  $\mathbf{c}$  with  $\text{height}(\mathbf{c}) \approx \bar{u}$ , by the same logic that strategy profiles in NE should be efficient. Namely, if  $\text{height}(\mathbf{c}) \gg \bar{u}$ , the defenders may increase the height of the coverage and improve their utility, since all of the targets are more protected.

However, due to the presence of our ABF, allegedly it may be the case that some attack distribution  $\omega(\mathbf{c}) = p$  that can be generated at  $\text{height}(\mathbf{c}) = u$ , cannot be approximated by any coverage vector of height  $\bar{u} \leq \text{height}(\mathbf{c}) \leq \bar{u} + \mathcal{O}(\delta)$ . We claim that this is never the case. Formally, we prove Proposition 7.1.

*Proof.* For any subset  $S \subseteq \mathcal{T}$ , we denote by  $\omega_S(\mathbf{u}) = \sum_{t \in S} \omega_t(\mathbf{u})$  the probability that the attack will take place in  $S$ . For every  $t \in S$ , we denote by  $\omega_{|S,t}(\mathbf{u}) := \omega_t(\mathbf{u})/\omega_S(\mathbf{u})$  the probability that the attacker will attack target  $t$  given that the attack will take place in  $S$ . We denote for every  $m' \in \mathbb{N}$  the set  $[m'] := \{1, \dots, m'\}$ . We also denote by  $p_S := \sum_{t \in S} p_t$  and  $p_{|S,t} := p_t/p_S$ .

W.L.O.G assume that  $p_1 \leq \dots \leq p_m$ . We prove by induction on  $m$  the following: for every  $1 \leq m' \leq m$ , there exists a utility vector  $\mathbf{u}$  such that:

1.  $\|\omega_{|[m'],t}(\mathbf{u}) - p(t|[m'])\|_{\infty} \leq 0.5m'(m' - 1)\epsilon$
2.  $\forall t \in [m'] : u \leq u_t < u + 2m'\delta$
3.  $|\omega_{|[m'],m'} - p_{|[m'],m'}| \leq (m' - 1)\epsilon$
4.  $\forall t \in [m'] : |u_t - (u + m\delta)| < m'\delta$

If  $m' = 1$  the claim is trivial taking  $\mathbf{u} = (u + m'\delta, 0, \dots, 0)$ . Assume the claim is true for  $m' - 1$ , and let  $\mathbf{u}$  be such a utility vector.

By Axiom 3, if we extend  $\mathbf{u}'$  to a utility vector  $\mathbf{u}' = (u_1, \dots, u_{m'-1}, u_{m'}, 0, \dots, 0)$ , we still have that  $\omega|_{[m'-1],t}(\mathbf{u}') = \omega|_{[m'-1],t}(\mathbf{u})$  for every target  $t \in [m' - 1]$ .

If we take  $u_{m'} = u_{m'-1} + \delta$  then  $\omega|_{[m'],m'}(\mathbf{u}') > 1 - (m' - 1)\epsilon$ , as all the rest of the targets are attacked w.p.  $< \epsilon$ .

If  $u_{m'} = u_{m'-1} - \delta$ , then by Definition 4  $\omega|_{[m'],m'}(\mathbf{u}') < \epsilon$ .

Therefore, by continuity, there exists a value  $u_{m'}$  satisfying  $|u_{m'} - u_{m'-1}| < \delta$  such that  $\omega|_{[m'],m'}(\mathbf{u}')$  is any value in the range  $[\epsilon, 1 - (m' - 1)\epsilon]$ . Therefore, there exists such a value satisfying  $|\omega|_{[m'],m'}(\mathbf{u}') - p|_{[m'],m'}| < (m' - 1)\epsilon$ , as desired. Since by induction hypothesis  $|u_{m'-1} - (u + m\delta)| < (m' - 1)\delta$ , by triangular inequality we have that  $|u_{m'} - (u + m\delta)| < m'\delta$  as desired. Since we didn't change the values of the rest of the targets in  $[m']$ , the only part left to check is that they also have the desired probabilities. Let  $t < m'$  be some target. Then:

$$\begin{aligned}
|\omega|_{[m'],t}(\mathbf{u}') - p|_{[m'],t}| &= \left| \frac{\omega|_{[m'-1],t}(\mathbf{u}')}{\omega|_{[m'-1]}(\mathbf{u}')} \omega|_{[m'-1]}(\mathbf{u}') \right. \\
&\quad \left. - p|_{[m'],t}| \right| \\
&\leq \left| \frac{\omega|_{[m'-1],t}(\mathbf{u})}{\omega|_{[m'-1]}(\mathbf{u})} - p|_{[m'-1],t}| \cdot \omega|_{[m'-1]}(\mathbf{u}) \right. \\
&\quad \left. + |\omega|_{[m'-1]}(\mathbf{u}) - p|_{[m'-1]}| \cdot p|_{[m'-1],t}| \right| \\
&\leq 0.5(m' - 1)(m' - 2)\epsilon + (m' - 1)\epsilon \\
&= 0.5m'(m' - 1)\epsilon
\end{aligned} \tag{19}$$

Therefore, we get that  $\|\omega(\mathbf{u}') - p(\cdot|t \leq m')\|_\infty \leq 0.5m'(m' - 1)\epsilon$ . Therefore, by induction, the claim is also true for  $m' = m$ , which completes the proof.  $\square$

This proves that if  $\overline{\text{cov}}(v)$  is feasible, then by decreasing the coverage of each target by  $\mathcal{O}(\delta)$ , it is possible to  $\mathcal{O}(\epsilon)$  approximate any attack distribution  $p$ . This will result with an  $\mathcal{O}(\delta + \epsilon)$  utility loss for each defender. In what follows we combine all of our observations to construct stable solutions.

## C.4 A Construction of an $\alpha$ $\zeta$ -Core Solution

By this stage, we have everything needed to construct a solution to the  $\alpha$   $\zeta$ -core. We first introduce an approximation-preserving reduction into a simpler problem which can be solved accurately. To do so we again overload notations and define the utility when the attack probability is specified:

**Definition C.4 (Defender Utility for Specified Attack Distribution).** Let  $p$  be some probability distribution over  $m$  targets, let  $i \in \mathcal{N}$  be some defender and  $\mathbf{c} \in \mathcal{C}_{k_i}$ . Then  $U_i^d(\mathbf{c}, p) := \sum_{t \in \mathcal{T}} p(t) \cdot U_i^d(\mathbf{c}, t)$ .

**Lemma C.5.** Fix  $\mathcal{CS}_{GC}$  to be the grand coalition structure with min-max height. There exist  $A, B, \delta_0, \epsilon_0 > 0$  s.t. for every  $\delta < \delta_0, \epsilon < \epsilon_0$ , the following is true:

- $\forall$  probability distribution  $p, \exists$  coalition structure  $\mathcal{CS}$  s.t.  $\forall i \in \mathcal{N} : U_i^d(\mathcal{CS}_{GC}, p) \leq U_i^d(\mathcal{CS}) + A\epsilon + B\delta$ .

- $\forall$  coalition structure  $\mathcal{CS}$ ,  $\exists$  probability distribution  $p$  s.t.  $\forall i \in \mathcal{N} : U_i^d(\mathcal{CS}) \leq U_i^d(\mathcal{CS}_{GC}, p) + A\epsilon + B\delta$ .

*Proof.* Note that the proof is constructive, and the reductions are efficient. Basically, the first direction follows from Proposition 7.1, and it is efficient due to Axiom 2. As for the second direction, take  $p := \omega(\text{cov}(\mathcal{CS}))$ . Any target  $t \in \mathcal{T}$  with  $U^a(\mathcal{CS}, t) < \bar{u} - \delta$  is attacked with probability less than  $\epsilon$ , therefore it would contribute a utility loss bounded by  $r_i^d(t)\epsilon$ . For the rest of the targets, since  $\text{height}(\mathcal{CS}_{GC}) = \bar{u}$ , the utility loss is bounded by  $\delta$ .  $\square$

**Theorem C.6.** *There exists a game dependent  $A, B, \delta_0, \epsilon_0 > 0$  such that for every  $0 < \delta < \delta_0, 0 < \epsilon < \epsilon_0$  and every  $(\delta, \epsilon)$ -ABF, the  $\alpha \zeta = (A\epsilon + B\delta)$ -core is non-empty, and there exists an efficient algorithm that finds stable solutions in it.*

The main idea is the following: we start with a  $\mathcal{CS}$  resistant to any deviation set  $D \subseteq \mathcal{N}$  with  $k_D < k_{\mathcal{N}}$  resources. Denote by  $\mathbf{u}^0 := (U_i^d(\mathcal{CS}))_{i \in \mathcal{N}_+}$ . We then search for a Pareto efficient solution in  $\mathcal{U}_{Pr} := \{(U_i^d(\mathcal{CS}_{GC}, p))_{i \in \mathcal{N}_+} | p \text{ is a probability distribution}\}$  which satisfies the constraints  $U_i^d(\mathcal{CS}_{GC}, p) \geq \mathbf{u}_i^0$  for each defender  $i \in \mathcal{N}_+$ . We know such a solution  $\mathcal{CS}'$  exists, since  $p = \mathbb{1}_{t^*}$  satisfies the constraints. There are three key steps left to finish the proof:

- The solution  $\mathcal{CS}'$  is resistant to any  $\epsilon'$  successful deviation of any  $D \subseteq \mathcal{N}$  defenders with  $k_D < k_{\mathcal{N}}$  resources.  $\epsilon'$  is the same as in Theorem C.1. This follows directly from the fact that  $\mathcal{CS}'$  realizes the constraints and because  $\mathcal{CS}$  is resistant to such deviations.
- The solution  $\mathcal{CS}'$  is resistant to any deviation set  $D \subseteq \mathcal{N}$  of  $k_D = k_{\mathcal{N}}$  resources. This follows directly from the Pareto efficiency of the solution.
- The solution  $\mathcal{CS}'$  can be found efficiently. In order to see that, denote for each  $t \in \mathcal{T}$  by  $\mathbf{u}_t^d := (U_i^d(\mathcal{CS}_{GC}, t))_{i \in \mathcal{N}_+}$ . Then  $\mathcal{U}_{Pr}$  is a polytope with vertices at  $\{\mathbf{u}_t^d | t \in \mathcal{T}\}$  (or a subset of that). Furthermore, the constraints for the defenders are linear with the variables  $p(t)$  which span this polytope. Therefore, if we take any vector  $\alpha := (\alpha_i)_{i \in \mathcal{N}_+}$  with all entries  $> 0$ , then maximizing  $\langle \alpha, \mathbf{u} \rangle$  for  $\mathbf{u} \in \mathcal{U}_{Pr}$  under the specified constraints is an LP optimization problem with  $\mathcal{O}(m+n)$  constraints and  $m$  variables, therefore it can be solved efficiently, and yield the resulting solution.

The last step is to perform our approximation reduction step using Lemma C.5. In this step we gain an additional  $\mathcal{O}(\delta + \epsilon)$  utility loss, and the constants add-up.

## D Examples

### D.1 Example of an ABF

As mentioned, a well studied private case for an attacker ABF is the *soft-max* function:

$$\text{soft-max}_t(\mathbf{u}) := \frac{e^{u_t}}{\sum_{t' \in \mathcal{T}} e^{u_{t'}}$$

In the domain of security games it is often called *quantal response*, as the more utility the attacker gets from a target, the more the likelihood of being attacked increases, respectively.

**Proposition D.1.** *The quantal response function  $\omega$  is an ABF.*

*Proof.*

**Probability Distribution** Let  $\mathbf{u}$  be some attacker utility vector and let  $t \in \mathcal{T}$ . Denoting by  $A = \sum_{t \in \mathcal{T}} e^{u_t}$ :

$$\sum_{t \in \mathcal{T}} \omega_t(\mathbf{u}) = \sum_{t \in \mathcal{T}} e^{u_t} / A = A / A = 1$$

and also  $0 \leq \omega_t(\mathbf{u}) \leq 1$  since  $\exp(x) \geq 0$  for any  $x \in \mathbb{R}$ .

**Monotonous** Let  $\delta > 0$ ,  $t \in \mathcal{T}$  and define  $\mathbf{u}'$  such that  $u'_t = u_t + \delta$  and  $u_{t'} = u_{t'}$  for any other  $t' \in \mathcal{T}$ . Denote by  $A' = \sum_{t \in \mathcal{T}} e^{u'_t}$ , then  $A' > A$ . Also denote by  $B = \sum_{t \neq t' \in \mathcal{T}} e^{u_t} = \sum_{t \neq t' \in \mathcal{T}} e^{u'_t}$ . Then:

$$\omega_t(\mathbf{u}) = 1 - (1 - \omega_t(\mathbf{u})) = 1 - B/A < 1 - B/A' = \omega_t(\mathbf{u}')$$

**Independence** Let  $\emptyset \neq S \subseteq \mathcal{T}$  and let  $t \in S$ . Let  $\mathbf{u}'$  be some attacker utility vector such that for each  $t' \notin S$   $u'_{t'} = u_{t'}$ . Then, denoting by  $C = \sum_{t \in S} e^{u_t} = \sum_{t \in S} e^{u'_t}$ :

$$\frac{\omega_t(\mathbf{u})}{\sum_{t' \in S} \omega_{t'}(\mathbf{u})} = \frac{e^{u_t}/A}{C/A} = \frac{e^{u_{t'}}/A'}{C/A'} = \frac{\omega_t(\mathbf{u}')}{\sum_{t' \in S} \omega_{t'}(\mathbf{u}')}$$

□

**Continuously Differentiable** Holds as a summation, division and composition of continuously differentiable functions.

Note that this would have also worked for any  $\omega_t(\mathbf{u}) = \frac{f(u_t)}{\sum_{t' \in \mathcal{T}} f(u_{t'})}$  where  $f$  is a one-dimensional monotonic continuously differentiable function into  $(0, \infty)$ .

Next we will construct a  $\delta, \epsilon$  ABF.

**Proposition D.2.** *Let  $\delta, \epsilon > 0$ . Then there exists a  $(\delta, \epsilon)$ -ABF.*

*Proof.* Consider  $\omega(\mathbf{u}) := \text{soft-max}(n\mathbf{u})$  for some  $n \in \mathbb{N}$  to be determined later. By Proposition D.2  $\omega$  is an ABF. Consider two targets  $t, t'$  where  $u_{t'} < u_t - \delta$ , then:

$$\omega_{t'}(\mathbf{u}) < \frac{e^{n(u_t - \delta)}}{\sum_{s \in \mathcal{T}} e^{nu_s}} < e^{-n\delta}$$

Therefore, we want  $e^{-n\delta} < \epsilon$ , and so  $n > \frac{1}{\delta} \ln \frac{1}{\epsilon}$ .

□

## D.2 Example of an SSG with no NSE

**Example D.3 (NSE may not Exist).** Consider a single defender SSG with two targets. The attacker is indifferent to the targets. Assume for instance that the reward and the penalty of the attacker on both targets are 1, 0 respectively. The defender has two security resources, and prefers target  $t_1$  over  $t_2$ . Assume for instance that  $r^d(t_1) = 3 > p^d(t_1) = 2 > r^d(t_2) = 1 > p^d(t_2) = 0$ . Then, assuming the attacker breaks ties in opposition to the defender, or even attacks randomly over BR, there is no NE for this game.

*Proof.* First, observe that the defender can get a utility arbitrarily close to  $r^d(t_1) = 3$ . Indeed, let  $0 < \epsilon < 1$  and consider the coverage vector  $(1 - \epsilon, 1)$ , which is feasible for the defender. Then  $\text{BR} = \{t_1\}$ , therefore the utility of the defender will be  $3 - \epsilon$ .

We claim that there is no strategy for the defender to gain a utility  $\geq 3$ . If that is the case, then there is no NE, since the defender can always improve his utility to be a little closer to 3.

Indeed, assume in contradiction that the defender gets a utility  $\geq 3$  playing some strategy  $\mathbf{c}$ . Let  $p$  be the probability that the attacker takes target  $t_1$ . Then  $3 \leq U^d(\mathbf{c}) = p \cdot U^d(\mathbf{c}, t_1) + (1 - p) \cdot U^d(\mathbf{c}, t_2) \leq p \cdot r^d(t_1) + (1 - p) \cdot r^d(t_2) = 2p + 1 \leq 3$ . Therefore,  $p$  must be equal to 1 and so  $c_1 < c_2 \leq 1$ . However, in that case, we get that  $3 = U^d(\mathbf{c}, t_1) = c_1 \cdot r^d(t_1) + (1 - c_1) \cdot p^d(t_1) = 2 + c_1 < 3$ , which is a contradiction.  $\square$

Note that the fact that the defender can get a utility arbitrarily close to 3, but cannot get exactly 3, is due to the fact that the utility is not continuous. Specifically, we took advantage of the fact that:

$$3 = \lim_{c_1 \rightarrow 1^-} U^d((c_1, 1)) \neq U^d((1, 1)) = 1$$

In the case that the attacker plays randomly over equally appealing targets in BR, we would still get  $U^d((1, 1)) = 0.5 \cdot 3 + 0.5 \cdot 1 = 2 \neq 1$ , and a similar argument works.

## D.3 Example of a Robust Solution

Consider the game described in Example D.3. For single defender games, due to the phenomena described in that example, it is well accepted that the attacker breaks ties in favor of the defender. A NE under this assumption is called SSE; it is known to exist and can efficiently be found using LP. However, as we will now demonstrate, it is extremely sensitive and is not robust.

We begin by demonstrating how exactly this optimistic tie-breaking fails to form robust solutions even for single defender games, and then show how our model for SSG handles this problem.

**Proposition D.4.** Consider the game described in Example D.3, assuming that the attacker breaks ties in favor of the defender. Let  $0 < \delta < 1$ . Then any NE of this game is not  $\delta$ -robust.

*Proof.* We have already shown that the utility of the defender is bounded by 3, and now that we assume optimistic tie breaking there is also a unique strategy that achieves that, namely  $\mathbf{c}_{NE} = (1, 1)$ , which is therefore the only NE. Then if we assume, for instance, that the parameters of the attacker on target  $t_2$  slightly change to be  $r^a(t_2) = \delta/2$ ,  $p^a(t_2) = \delta/2$  and the parameters of target  $t_1$  stay the same, then in the resulting game the defender gets a utility of  $\delta/2$ , and therefore it's not NE as the defender can deviate and get a utility of 3.  $\square$

This formally explains the more intuitive argument, that when a solution assumes optimistic tie-breaking, a slight change in the attacker's parameters can lead to major utility loss for the defenders, altering the NE. This happens due to the non-continuity of the defender's utility at the equilibrium point. However, when the attacker follows an ABF and when the two targets are covered equally, they are attacked with some probabilities  $p, q$  (50%-50% in quantal response for example), and slightly changing the attacker's parameters ( $\delta$  parameter) will only slightly change the attacker's probability of attacking each target ( $\epsilon$  parameter). This is the property that leads to robustness.

Next we are going to demonstrate a robust approximate NE of the above game following our constructive proof for Theorem 6.1. To make the game canonical for the sake of simplicity, assume that the defender has a single resource. Therefore,  $\bar{u} = 0.5 \in (0, 1] = (\bar{p}, \bar{r})$ .

Following the constants of the proof, we have  $\delta_0 = 1/64$  (maximal value of  $\delta$  such that the game is  $\delta$ -canonical),  $\epsilon_0 = 1/2$ . With respect to the parameters of the game, observe that these values are not that small, as solving with an accuracy of 2/3 of a decimal point for game parameters in the range of  $[0, 10]$  seems both common and feasible. Therefore, assume that the attacker's ABF for this game has parameters  $\delta = 0.01 < \delta_0$  and  $\epsilon = 0.1 < \epsilon_0$ . For example, consider  $\omega(\mathbf{u}) = \text{soft-max}(n\mathbf{u})$  with  $n = 250$  (see Proposition D.2). However, the proof itself doesn't require knowing the exact ABF  $\omega$ , it is enough to know its parameters.

Therefore, we know from the proof that the solution would look something like  $\mathbf{c}_0 = (0.5 - \alpha, 0.5 + \alpha)$  which, assuming  $\alpha > \delta/2$ , will induce the attacker to attack target  $t_1$  with a probability of at least  $1 - \epsilon$ . Therefore the utility of the defender is at least  $(1 - \epsilon)(2 + (0.5 - \alpha)) = 2.5 - 2.5\epsilon - \alpha + \alpha\epsilon > 2.5 - 2.5\epsilon - \alpha$ , which we will denote by  $U_0$ .

Based on this theorem, we can bound the utility of a defender in this game by  $U < \max 2 + (0.5 + \delta/2), 3\epsilon + 1 = 2.5 + \delta/2$ .

For this example game it is intuitive: either  $c_1 < c_2 - \delta$ , in which case  $c_1$  is limited to  $0.5 + \delta/2$ , or the opposite is true, which would mean that target  $t_1$  is attacked with a probability of at most  $\epsilon$ . We can bound the  $\omega_{t_1}(\mathbf{c})$  by 1 in the first case, and  $U_1^d(\mathbf{c}, t_1)$  by  $r_1^d(t_1)$  for the second case.

Therefore, taking for instance  $\alpha = \delta$  as the theorem suggests,  $\mathbf{c}_0$  is a  $\zeta = 2.5\epsilon + 1.5\delta = 0.265$ -NE. Thus, following Theorem 5.2,  $\mathbf{c}_0$  is a  $\delta = 0.01$ -robust  $\zeta' = 0.665$ -NE.

Also note that as it is a single defender game,  $\mathbf{c}_0$  is also a  $\delta$ -Robust  $\zeta'$  solution to the  $\alpha$ -Core (and also  $\beta, \gamma$ -Cores which coincide for a single defender).

**Example D.5 (Grand Coalition Deviation of Type 2).** Consider the SSG depicted in 2. Let us now use procedure GC-ALLOC with input  $(\bar{u} - A\delta, \bar{u})$ . We claim that the grand coalition has a deviation of type 2 from the resulted strategy profile.

Defender 1 will first cover target  $t_1$ , which is the less favorable target for him, up to height  $u$ . He will then cover part of target  $t_3$ . Afterwards, defender 2 will cover target  $t_2$ , which is less favorable for him, and then he will use the rest of his resources to cover target  $t_3$ , which will result with an attacker height  $\bar{u} + \delta'$ .

We take  $A\delta > \delta$  so that deviation of type 1 will yield up to  $\mathcal{O}(\epsilon)$  utility gain for the defenders, but small enough such that 3 will yield up to  $\mathcal{O}(\delta)$  utility gain for the defenders. For example, we can take  $A\delta = 2\delta$ .

Player	$t_1$	$t_2$	$t_3$	$K_i$
Defender 1	1,1	300,300	2,2	1
Defender 2	300,300	1,1	2,2	1
Attacker	1,0	1,0	1,0	

Table 2: Example of an SSG where using GC-ALLOC doesn't yeild a stable solution due to a deviation of type 2

Unfortunately, the grand coalition still has a deviation of type 2 with  $\Omega(1)$  utility gain for both defenders. In the original  $\mathcal{CS}$ , both defenders get a utility close to 2. Consider a deviation of the grand coalition, with a coverage of  $(0, 0, 1)$ . The utility that each defender gets now is at least  $0.5 * 300 = 150 \gg 2$ . Therefore, this is a successful deviation.

Intuitively, the reason the method for the original game for finding the core doesn't work is that in this case defender 1 covers a target that defender 2 prefers over the chosen target, and vice versa. Together, they can induce the attacker to attack these preferred targets, each with probability 0.5. This wasn't possible in the NLE for the non-cooperative game by definition. Also, in the original model, every defender assumed that the attacker attacks the target that is worst for him, so this wouldn't be a successful deviation there either.

Since the reward and penalties on each target are the same in this game, it is convenient to specify which solutions are in the core by the distribution of the attacker. We claim that the distribution  $(p, 1 - p, 0)$  is in the  $\epsilon$ -core for any  $\frac{2}{300} \leq p \leq 1 - \frac{2}{300}$ . Indeed, in such a strategy the defenders get utilities  $300p, 300(1 - p) \geq 2$  respectively. A deviation by a single defender  $i$  is not successful, since the other defender can block the target  $t_{3-i}$  and the utility of defender  $i$  is bounded by 2. A successful deviation of the grand coalition means that  $u_1 > 300p, u_2 > 300(1 - p)$ . If the coverage vector of this deviation is  $(q_1, q_2, q_3)$  then we must have  $q_1 > p, q_2 > (1 - p)$  and therefore  $q_1 + q_2 + q_3 > 1$ , which is impossible.

It seems as if the solutions we find are a subset of the Pareto efficient defender utility vectors, specifically the Pareto efficient solutions where no subset of defenders can deviate. This observation will turn out to be important for constructing an element in the core.



#### D.4 An example of an SSG with an Empty Robust $\gamma$ -Core

In  $\gamma$ -core (see Definition D.6), we assume that, after deviation, the rest of the defenders are singletons and cannot cooperate, as suggested in [Chander, 2010]. There are usually further limitations on the revenge takers' response, such as that each revenge taker best responds to a deviation [Chander and Tulkens, 2006] or that the resulting strategy profile forms a NE [Chander, 2007]. However, without any of these assumptions, we show that the  $\gamma$ -core of the game may be empty.

**Definition D.6 (Robust  $\gamma$ -Core).** *Let  $\mathcal{CS}$  be some coalition structure. A deviation  $(D, \mathbf{x}_D)$  from  $\mathcal{CS}$  is  $\gamma$ -successful if for every non-cooperative strategy profile of the rest of the defenders  $Y = (\mathbf{y}_r)_{r \in R}$ , the following holds for every deviator  $i \in D$ :*

$$U_i^d(\langle \mathbf{x}_D, Y \rangle) > U_i^d(\mathcal{CS})$$

*$\mathcal{CS}$  is in the  $\gamma$ -core if it has no  $\gamma$  successful deviations.*

We demonstrate that the  $\gamma$ -core may be empty.

**Example D.7 (Empty  $\gamma$ -Core).** *The SSG depicted in Table 1 has an empty  $\gamma$ -core, assuming  $(\frac{7-3\sqrt{5}}{12}, \frac{1}{15})$ -ABF.*

Basically, defenders 1, 2 can always  $\gamma$ -deviate and get a utility  $\geq 4$  each, as can defenders 3, 4, but no coalition structure can yield all defenders  $\geq 4$  utility. Here is a more formal and detailed explanation:

*Proof.* If defenders 1, 2 decide to play  $(2/3, 2/3, 2/3, 0, 0, 0)$  which can be implemented as  $K_{\{1,2\}} = 2$ , then no matter what defenders 3, 4 decide to do, a target  $t \in t_1, t_2, t_3$  is attacked w.p. at most  $3\epsilon$ . Indeed, if defenders 3, 4 want to incentivize the attacker to attack targets in  $t_1, t_2, t_3$ , they need to cover targets  $t_4, t_5, t_6$  and result with a coverage with a diameter smaller than  $\delta$ . The best they can do non-cooperatively with their  $K_{3,4} = 2$  resources is to play  $(0, 0, 0, \alpha, \alpha, \alpha)$  where  $\alpha = 1 - \alpha^2$ . This leads to an attacker height of  $1 - \alpha = \frac{3-\sqrt{5}}{2}$ . However,  $\delta < \frac{3-\sqrt{5}}{2} - \frac{1}{3}$ .

Therefore, after deviation, defenders 1, 2 will get at least  $(1 - 3\epsilon) \cdot p_1^d(t_4) = 5 \cdot (1 - 3\epsilon) = 4$ .

On the other hand, by symmetry, it is easy to see that defenders 3, 4 can get at least a utility of 4 after deviation as well, assuming defenders 1, 2 will respond without coordination. Therefore, if a coalition structure  $\mathcal{CS}$  was in the  $\gamma$ -core, then it must yield a utility  $U_1^d(\mathcal{CS}) = U_2^d(\mathcal{CS}) \geq 4$ , since otherwise, we proved that there exists a deviation of defenders 1, 2. Note that since defenders 1, 2 are of the same type, their utility is always equal. For the same reason,  $U_3^d(\mathcal{CS}) = U_4^d(\mathcal{CS}) \geq 4$ . Denote by  $p$  the probability that the attacker attacks a target in  $\{t_1, t_2, t_3\}$ . On the one hand,  $4 \leq U_1^d(\mathcal{CS}) \leq p \cdot 1 + (1-p) \cdot 6$ , therefore  $p \geq 0.6$ . On the other hand,  $4 \leq U_3^d(\mathcal{CS}) \leq p \cdot 6 + (1-p) \cdot 1$ , therefore  $p \leq 0.4$ , a contradiction. Thus, there always exists a  $\gamma$ -successful deviation, and the  $\gamma$ -core is empty.  $\square$

It is worth mentioning that the  $\gamma$  core for the original MSSG model is also empty, by a similar example.