

Bayesian Repeated Zero-Sum Games with Persistent State, with Application to Security Games

Vincent Conitzer¹, Yuan Deng², and Shaddin Dughmi³

¹ Duke University `conitzer@cs.duke.edu`

² Duke University `ericdy@cs.duke.edu`

³ University of Southern California `shaddin@usc.edu`

Abstract. We study infinitely-repeated two-player zero-sum games with one-sided private information and a persistent state. Here, only one of the two players learns the state of the repeated game. We consider two models: either the state is chosen by nature, or by one of the players. For the former, the equilibrium of the repeated game is known to be equivalent to that of a one-shot public signaling game, and we make this equivalence algorithmic. For the latter, we show equivalence to one-shot team max-min games, and also provide an algorithmic reduction. We apply this framework to repeated zero-sum security games with private information on the side of the defender and provide an almost complete characterization of their computational complexity.

Keywords: Bayesian repeated game · Equilibrium characterization · Equilibrium computation · Computational complexity

1 Introduction

Private information can give one a strategic advantage over other players in a game. However, if play is repeated, then taking advantage of one’s private information through one’s actions risks leaking that information and thereby the advantage. This is nicely illustrated in the movie *The Imitation Game*, in which British intelligence, having cracked the *Enigma* code, strategically decides not to act on some of its information, in order to preserve its informational advantage [16]. Less dramatically, consider a buyer and a seller that interact repeatedly. The seller has a higher-quality and a lower-quality version of the item for sale, and offers these at different prices. The buyer may, at the current prices, prefer the higher-quality version – but worry that choosing this option will reveal her (persistently) high valuation/type, causing the seller to raise prices in the future, and therefore choose the cheaper low-quality version instead.

In equilibrium, to what extent should a party with an informational advantage refrain from acting on this information? This is the question we set out to address in this paper. It is, in its most general form, a challenging question to answer. The state of the game may change over time; there may be a multiplicity of equilibria; the discount factor matters; and so on. Thus, answering

the question in general would require us to simultaneously resolve a number of fundamental questions in (algorithmic) game theory. In this paper, in order to stay focused on the question at hand, we focus on the following special case:

- The state of the game is *persistent*, i.e., it does not change over time (the game is repeated rather than stochastic).
- Only one player has private information, and it does not change.
- The game is two-player and zero-sum.
- Each agent cares about their long-term average payoff.

Even in this setting, it is easy to see that the optimal answer is in general not one of the two extremes – either always exploit information fully, or never use it. Some information may not be actionable for the adversary, in which case one can simply take advantage of it and not worry about revealing it. On the other hand, for other information, it is possible that the adversary would be able to make even better use of it than the initially better-informed player. In that case, the benefits of getting to use the information for one round, without the adversary being able to use it in that particular round, will be completely wiped out by the infinitely many remaining rounds in which the adversary can use the information better.

The technical and conceptual foundations for the study of repeated games of incomplete information with persistent state were laid by [2]. They consider a persistent state of the game drawn by nature from a common prior, and agents who receive private signals regarding this state. [18] provides an in-depth accounting of the special case of this model with two players and zero-sum payoffs. The aforementioned texts reveal that the even-more-special case we consider, that of repeated two-player zero-sum games with one-sided private information, admits an essentially-unique equilibrium (in the sense of payoff equivalence) with an elegant, simple, and instructive characterization which is robust to modeling assumptions. In particular, the equilibrium of the repeated two-player game is equivalent, in a precise technical sense, to the equilibrium of a one-shot *public signaling game* with three players. Moreover, this characterization is robust to how one chooses to model long-term payoffs; say through using a discount factor, taking the limit of the finite repeated game as the number of stages grows to infinity, or considering the infinite game directly. Even mild generalizations of this special case, for example to more players, non-zero-sum payoffs, or incomplete information on both sides, lead to the collapse of this characterization, and such settings are not yet fully understood to the best of our knowledge. This further cements our model as the timely choice for algorithmic study.

1.1 Our Contributions

We examine repeated two-player zero-sum games with one-sided private information from the perspective of algorithmic game theory, both in general and as exemplified by application to the influential domain of security games (see e.g. [20]). We consider both the case when the state is drawn by nature — this is

the classical model in [2,18] — as well as a natural, and to our knowledge novel, variant in which the (typically randomized) state is chosen by one of the players, who is therefore the informed party. We refer to this variant as the *allocation* model of repeated games.

The domain-agnostic part of the paper is organized as follows. For the classical model, where the game state is drawn by nature, we first provide (a) our own exposition of the previously-described equilibrium characterization in terms of one-shot public signaling games, one that is particularly tailored to an algorithmic game theory audience and makes explicit the connection to recent work on public signaling games (e.g. [6,7,8]). Then, we turn to our novel contributions. We provide (b) an efficient reduction to equilibrium computation in the related one-shot public signaling game to make the equilibrium characterization constructive. For the allocation model, where one of the players determines the (persistent) state, we provide (a') a characterization of the equilibrium of the repeated game as equivalent, in a precise technical sense, to the equilibrium of a one-shot three-player *team max-min game*, as first studied by [19]; (b') an efficient reduction to computing the equilibrium of the associated team max-min game. We note that, in both (b) and (b'), the uninformed player's strategy is particularly nontrivial, and involves efficiently solving a related instance of *Blackwell's approachability* [1,4]. We also note that the reductions in (b) and (b') are “reversible”, since both the repeated game and the associated one-shot game share the same game value. Finally, we (c) show that the allocation model is computationally easier than the classical model by way of a polynomial time reduction. We note that this is not reversible, and the complexity relationship is strict, as evidenced by our results for security games which we describe next.

We then examine repeated zero-sum security games with private information on the side of the defender. In the security games we consider, the state is a deployment of “treasures” to “locations”, a defender strategy is a deployment of “defensive resources” to the locations, and the attacker's strategy is a location to attack. Such security games are particularly versatile exemplars for both the classical and allocation models of repeated games with persistent state. The classical model abstracts challenges faced in recent applications to environmental protection [9,10,21,22,23], where the locations of environmental assets (the treasures) are determined by nature and slow to change over time. The allocation model can be applied to armed conflict scenarios in which supply-chain assets (the treasures) must be deployed covertly to locations early on in the conflict, and can not be easily moved from stage to stage.

We show that the classical model of repeated security games is strongly NP-hard even when treasures, locations, and defensive resources are homogeneous. A more nuanced picture emerges for the allocation model of repeated security games: the fully homogeneous case is tractable, as is the case where only the treasures are heterogeneous. The fully heterogeneous case is strongly NP-hard. Remaining cases are either weakly or strongly NP-hard, and we provide an almost complete accounting of the computationally complexity of all combinations.

2 Preliminaries

2.1 One-shot Games

A one-shot two-player zero-sum game of complete information is described by a utility function $\mathcal{U} : S_1 \times S_2 \rightarrow \mathbb{R}$, where S_i is the family of *pure strategies* for player i , and $\mathcal{U}(s_1, s_2)$ is the utility of player 1 when player 1 plays $s_1 \in S_1$ and player 2 plays $s_2 \in S_2$. Implicitly, the utility of player 2 is $-\mathcal{U}(s_1, s_2)$. A *mixed strategy* for player i is $s_i \in \Delta(S_i)$, where $\Delta(S_i)$ is the set of distributions over S_i . A one-shot two-player Bayesian zero-sum game with incomplete information on one side $(\Pi, \{\mathcal{U}^\theta\}_{\theta \in \Theta})$ is given by

- Pure strategy sets S_1 and S_2 for players 1 and 2 respectively;
- A family Θ of *states of nature*;
- For each state $\theta \in \Theta$, a one-shot two-player zero-sum game of complete information \mathcal{U}^θ ;
- A *prior distribution* Π over states of nature Θ .

In such a game, nature draws θ from Θ according to the prior Π and then player 1 learns the state θ while player 2 is uninformed about the state. Both players simultaneously choose their strategies s_i , which results in a utility of $\mathcal{U}^\theta(s_1, s_2)$ to player 1 and $-\mathcal{U}^\theta(s_1, s_2)$ to player 2. Moreover, given a distribution Π over Θ , we denote by \mathcal{U}^Π the game induced by Π such that player 1's payoff is $\mathcal{U}^\Pi(s_1, s_2) = \sum_{\theta \in \Theta} \Pi(\theta) \cdot \mathcal{U}^\theta(s_1, s_2)$. We restrict attention to games where Θ , S_1 , S_2 are finite, or at least compact. All *mixed* Nash equilibria of such a game are payoff equivalent to the Nash equilibrium in which each player employs their maximin mixed strategy [15].

2.2 Bayesian Repeated Games

We now describe the classical model of Bayesian repeated games that we consider, henceforth just *Bayesian repeated games* for convenience. Here, a Bayesian zero-sum game is repeated infinitely many times, with incomplete information on one side. We call the one-shot game the *stage game*, and refer to each iteration as a *stage*. We replicate the standard assumptions made by [2,18], as follows. We assume that the state of nature is *persistent*: it does not change from stage to stage.⁴ Moreover, we assume that players observe each others' pure strategies after each stage, but do not observe the payoffs directly. This assumption is necessary for the model to be interesting: If players can observe the payoffs directly, then the uncertainty in the game is superfluous, as players can eventually reconstruct relevant entries of the game matrix and the state of nature. Obscuring payoffs in this manner can be viewed as abstracting a situation where payoffs are

⁴ If the state of nature is drawn afresh at each stage, then repetition is superfluous for a zero-sum game: the folk theorem and minimax theorem imply that repeating the minimax equilibrium at each stage is the essentially unique equilibrium of the repeated game (up to payoff equivalence).

delayed till the end of the (long, many stage) game. Formally, given a two-player Bayesian zero-sum stage game $G_{\text{repeated}} = (\Pi, \{\mathcal{U}^\theta\}_{\theta \in \Theta})$ as described above, the Bayesian repeated game proceeds as follows:

1. θ is drawn by nature from Π and player 1 learns θ while player 2 does not;
2. The stage game \mathcal{U}^θ is repeated infinitely many times. After each stage, each player observes the pure strategy played by the other player, but does not directly observe the utility gained.

A *history* of play with T stages $H_T = ((s_1^1, s_2^1), (s_1^2, s_2^2), \dots, (s_1^T, s_2^T))$ is a finite sequence, where s_i^t is player i 's pure strategy at stage t . For convenience, we will use the vectorized form without superscript $\vec{s}_i = (s_i^1, \dots, s_i^T)$ to represent the strategy of player i . A pure strategy for player 1 in the repeated game is a function which maps the state θ and an observed history H to player 1's strategy in the next stage of the repeated game, while a pure strategy for player 2 simply maps the observed history H to player 2's strategy in the next stage. A mixed strategy is naturally a distribution over such functions.

2.3 Bayesian Allocation Games

In addition to classical Bayesian repeated games, we introduce a novel variant, the *Bayesian allocation game*, in which the distribution Π of the states is determined by player 1 instead of the nature. Formally, given one-shot games $G_{\text{alloc}} = (\{\mathcal{U}^\theta\}_{\theta \in \Theta})$, the Bayesian allocation game proceeds as follows:

1. Player 1 selects a prior Π over Θ that player 2 cannot observe;
2. θ is drawn by nature from Π and player 1 learns θ while player 2 does not;
3. The stage game \mathcal{U}^θ is repeated infinitely many times. After each stage, each player observes the pure strategy played by the other player, but does not directly observe the utility gained.

In the Bayesian allocation game, in addition to choosing the actions to play at each stage, player 1's strategy also includes a choice of the prior $\Pi \in \Delta(\Theta)$. We emphasize that the prior Π chosen by player 1 is *private*, and thus it cannot be observed by player 2 in the Bayesian allocation game.

2.4 Utility and Equilibrium Model

We consider the utility/equilibrium models deduced from the infinitely-repeated game perspective for agents that are interested in their long-term payoffs. Each player's expected utility is the limit, as $T \rightarrow \infty$, of his average expected utility over the first T stages alone. Though this limit may not exist in general, we can nevertheless define a value and equilibrium as in [2,18]. The max-min value of the game is the supremum over all player 1's mixed strategies, of the infimum over player 2's mixed strategies, of the limit infimum as $T \rightarrow \infty$ of player 1's average expected utility. Player 1's max-min strategy is that attaining this supremum.

We can similarly define the min-max value of the game and Player 2’s min-max strategy. When both the max-min and min-max values are equal we refer to them as the value of the game, and the corresponding max-min and min-max strategies form the equilibrium. For a Bayesian repeated game G_{repeated} and a Bayesian allocation game G_{alloc} , we denote their game value by $\nu_{\text{repeated}}(G_{\text{repeated}})$ and $\nu_{\text{alloc}}(G_{\text{alloc}})$, respectively. Several other natural utility/equilibrium models are equivalent to this one, as we describe in Appendix A.

Example 1. Consider a zero-sum security game with 3 identical locations (denoted by ℓ_A, ℓ_B, ℓ_C) and 2 identical treasures, in which the defender can defend 1 location. The defender determines how to allocate the treasures to the locations (once) and how to defend them (every round). The attacker earns one unit of payoff if she attacks an undefended location with a treasure, and zero otherwise. For comparison, in the one-shot Bayesian allocation game (i.e., if there is only a single round), it is straightforward to verify that the optimal strategy for the defender is to allocate two treasures uniformly at random, and for each realization, defend each of the two locations with a treasure with probability $\frac{1}{2}$, leading to an expected payoff $\frac{1}{3}$ for the attacker. However, it turns out that in the infinitely-repeated version, an optimal strategy (unique up to symmetries) to allocate the treasures for the defender is as follows:

- Allocate a treasure to ℓ_A with probability 1;
- Allocate the remaining treasure to ℓ_B with probability $\alpha = \frac{\sqrt{5}-1}{2} \approx 0.618$ and to ℓ_C with probability $1 - \alpha = \frac{3-\sqrt{5}}{2} \approx 0.382$.

In each stage of the repeated game, the defender defends ℓ_A with probability α (so that the attacker’s utility of attacking this location is $1 - \alpha$), and defends ℓ_B with probability $1 - \alpha$ (so that the attacker’s utility of attacking this location is $\alpha^2 = 1 - \alpha$). The defender never defends ℓ_C (so that the attacker’s utility for attacking this target is also $1 - \alpha$).

The above example illustrates a fundamental difference between a one-shot Bayesian allocation game and its infinitely-repeated counterpart. In the one-shot version, the optimal strategy for the defender correlates the allocation and the defensive strategy, and thus, the game is reduced to a two-player zero-sum normal-form game so that the minimax theorem can be applied. However, in the infinitely-repeated version, we will show that in the equilibrium, the allocation of treasures and the defensive strategy are *independent*, as in the example above. In other words, there exists no benefit for the defender to correlate the allocation and the defensive strategy in the infinitely-repeated Bayesian allocation game. Note that the attacker’s payoff is larger in the infinitely-repeated version as $1 - \alpha = \frac{3-\sqrt{5}}{2} > \frac{1}{3}$. Intuitively, this is because the attacker can observe the defender’s historical defensive actions in the infinitely-repeated game. This is disadvantageous for the defender: either the defensive actions over time give away where the treasures are, or these actions have to be chosen in such a way that they do not, which is a costly constraint. We also emphasize that the game value is an irrational number, demonstrating that the infinitely-repeated Bayesian allocation game cannot be solved by a linear program.

3 Reductions from Repeated Games to One-shot Games

In this section, we discuss the relationship between one-shot games and both our models of infinitely repeated games, so that one can solve the infinitely repeated game by first solving the corresponding one-shot game. The equivalence between classical Bayesian repeated games and public signaling games has already been shown by [2] and [18]; for completeness, we will fully elaborate on this equivalence first in Section 3.1. This will set the stage for our novel results on the equivalence between Bayesian allocation games and team max-min games (Section 3.2), and on the computational complexity of both models (Section 3.3). The omitted proofs in this section are deferred to Appendix B.

3.1 Equivalence between Bayesian Repeated Games and Public Signaling Games (Reproducing Known Results)

We begin with reproducing the known result relating the classical model of Bayesian repeated games to public signaling games [2,18].

Definition 1 (Public Signaling Game [6,7,8]). *Consider a one-shot two-player zero-sum game $G_{\text{signal}} = (\Pi, \{\mathcal{U}^\theta\}_{\theta \in \Theta})$ where players a-priori know nothing about θ besides its prior Π . We consider a credible principal who is privy to the realization of θ . The principal designs a public signaling scheme: a randomized function $\varphi : \Theta \rightarrow \Delta(\Sigma)$ mapping states of nature to an abstract set of signals Σ . The order of events is as follows:*

- The principal commits to φ ;
- The nature draws $\theta \sim \Pi$ and the principal learns θ ;
- The principal invokes the signaling scheme to obtain a signal $\sigma \sim \varphi(\theta)$;
- Both players learn σ , and update their beliefs about the state θ , denoted as $\Pi_{\varphi, \sigma}$, according to the Bayes' rule: $\Pi_{\varphi, \sigma}(\theta) = \frac{\Pr[\varphi(\theta)=\sigma] \cdot \Pi(\theta)}{\sum_{\theta' \in \Theta} \Pr[\varphi(\theta')=\sigma] \cdot \Pi(\theta')}$.
- Players play the equilibrium strategies in the zero-sum game $\mathcal{U}^{\Pi_{\varphi, \sigma}}$.

We assume that the principal designs φ so as to maximize player 1's expected utility, the maximum value of which, denoted by $\nu_{\text{signal}}(G_{\text{signal}})$, is the game value of the public signaling game.

It turns out the equilibrium in Bayesian repeated games corresponds to the solution of the above signaling problem in a precise sense, stated below [2,18].

Theorem 1. $\nu_{\text{repeated}}(G_{\text{repeated}}) = \nu_{\text{signal}}(G_{\text{signal}})$ when $G_{\text{repeated}} = G_{\text{signal}}$.

We will prove Theorem 1 by constructing the equilibrium strategy \bar{s}_1^*, \bar{s}_2^* for player 1 and 2, respectively in the Bayesian repeated game G_{repeated} from the solution of the public signaling game G_{signal} . For convenience, in the Bayesian repeated game, we will refer to player 1 (the informed player) as the *leader* and player 2 (the uninformed player) as the *follower*.

In particular, we will show that in the Bayesian repeated game G_{repeated} , if the leader plays strategy \bar{s}_1^* , then no matter how the follower reacts, the leader can guarantee himself an average utility at least the game value $\nu_{\text{signal}}(G_{\text{signal}})$ in the public signaling game G_{signal} over the first T stages as $T \rightarrow \infty$. On the other hand, if the follower plays strategy \bar{s}_2^* , then no matter how the leader reacts, the follower can guarantee the leader an average utility at most $\nu_{\text{signal}}(G_{\text{signal}})$ over the first T stages as $T \rightarrow \infty$.

Lemma 1. *When $G_{\text{repeated}} = G_{\text{signal}}$, in the Bayesian repeated game G_{repeated} , consider the following strategy for the leader:*

- upon learning the state θ of the nature, the leader invokes the optimal signaling strategy φ of the public signaling game G_{signal} to obtain $\sigma \sim \varphi(\theta)$;
- the leader then discards all information other than σ , i.e., behaves as if his belief is $\Pi_{\varphi, \sigma}$, and plays the maximin strategy in the game $\mathcal{U}^{\Pi_{\varphi, \sigma}}$, i.e., $\operatorname{argmax}_{s_1} \min_{s_2} \mathcal{U}^{\Pi_{\varphi, \sigma}}(s_1, s_2)$, repeatedly.

This strategy can guarantee the leader an average expected utility $\nu_{\text{signal}}(G_{\text{signal}})$.

Although the strategy for the leader is easy to construct from the signaling scheme of the public signaling game, the follower's strategy is not so straightforward. The main difficulty is that there does not exist a credible principal in the repeated game as in the public signaling game, and therefore, the follower is uncertain about whether the leader exactly follows the scheme. In particular, the leader might have incentive to deviate by sending a different signal: conditioned on his type θ , choose σ^* such that $\sigma^* = \operatorname{argmax}_{\sigma \in \Sigma} \mathcal{U}^\theta(s_1^*(\sigma), s_2^*(\sigma))$, where

$$s_1^*(\sigma) = \operatorname{argmax}_{s_1} \min_{s_2} \mathcal{U}^{\Pi_{\varphi, \sigma}}(s_1, s_2) \quad \text{and} \quad s_2^*(\sigma) = \operatorname{argmin}_{s_2} \max_{s_1} \mathcal{U}^{\Pi_{\varphi, \sigma}}(s_1, s_2).$$

In other words, the leader can send a signal σ^* that gives himself the maximum utility conditioned on θ . Therefore, the follower's strategy cannot rely on the possibly non-credible signaling scheme.

To circumvent this difficulty, we will construct an adaptive strategy for the follower, which does not depend on the non-credible signal σ but only depends on the prior Π and the history of play. Our approach relies on the solution of the dual program of the public signaling game. For convenience, given a distribution Π over Θ , let $f(\Pi) = \max_{s_1} \min_{s_2} \mathcal{U}^\Pi(s_1, s_2)$ be the game value of the induced game \mathcal{U}^Π . The problem of computing the optimal public signaling scheme can be formulated as the following linear program with infinitely many variables $x(\Pi')$ for $\Pi' \in \Delta(\Theta)$ [6,7,8]:

$$\begin{aligned} & \max \sum_{\Pi' \in \Delta(\Theta)} x(\Pi') \cdot f(\Pi') \\ & \text{s.t.} \quad \sum_{\Pi' \in \Delta(\Theta)} x(\Pi') \cdot \Pi'(\theta) = \Pi(\theta) \quad \forall \theta \in \Theta \\ & \quad \quad x(\Pi') \geq 0 \quad \quad \quad \forall \Pi' \in \Delta(\Theta) \end{aligned} \tag{1}$$

Intuitively, a signaling scheme can be viewed as a convex decomposition of the prior Π into a collection of posteriors $\{\Pi'\}$ [8,13]. Based on the primal, we can

construct its dual with $|\Theta|$ variables $y(\theta)$ for $\theta \in \Theta$ as follows:

$$\begin{aligned} \min \quad & \sum_{\theta \in \Theta} y(\theta) \cdot \Pi(\theta) \\ \text{s.t.} \quad & \sum_{\theta \in \Theta} y(\theta) \cdot \Pi'(\theta) \geq f(\Pi') \quad \forall \Pi' \in \Delta(\Theta) \end{aligned} \quad (2)$$

Let x^* and y^* be the solution of the primal and the dual, respectively. By strong duality, we have

$$\sum_{\Pi' \in \Delta(\Theta)} x^*(\Pi') \cdot f(\Pi') = \sum_{\theta \in \Theta} y^*(\theta) \cdot \Pi(\theta) = \nu_{\text{signal}}(G_{\text{signal}}).$$

We will interpret y and Π as vectors such that $\vec{y} = (y(\theta_1), \dots, y(\theta_{|\Theta|}))$ and $\vec{\Pi} = (\Pi(\theta_1), \dots, \Pi(\theta_{|\Theta|}))$. The inner product $\langle \vec{y}, \vec{\Pi} \rangle$ is defined as $\sum_{\theta \in \Theta} y(\theta) \cdot \Pi(\theta)$. The next proposition directly follows the feasibility of \vec{y}^* and strong duality:

Proposition 1. *For any prior $\vec{\Pi}$ in the public signaling game, there exists \vec{y}^* such that $\langle \vec{y}^*, \vec{\Pi} \rangle = \nu_{\text{signal}}(G_{\text{signal}})$ and $\forall \Pi' \in \Delta(\Theta)$, $\langle \vec{y}^*, \vec{\Pi}' \rangle \geq f(\Pi')$.*

Hence, if the follower can ensure that for any strategy \vec{s}_1 deployed by the leader, there exists an adaptive mixed strategy \vec{s}_2 for the follower such that,

$$\forall \theta \in \Theta, \lim_{T \rightarrow \infty} \frac{\sum_{t=1}^T \mathcal{U}^\theta(s_1^t, s_2^t)}{T} \leq y^*(\theta), \quad (3)$$

then the average utility of the leader as $T \rightarrow \infty$ would be

$$\lim_{T \rightarrow \infty} \sum_{\theta \in \Theta} \Pi(\theta) \cdot \frac{\sum_{t=1}^T \mathcal{U}^\theta(s_1^t, s_2^t)}{T} \leq \sum_{\theta \in \Theta} \Pi(\theta) \cdot y^*(\theta) = \nu_{\text{signal}}(G_{\text{signal}}).$$

To prove (3), it is equivalent to show that $\mathcal{R}(\vec{y}^*) = \{\vec{v} \mid \vec{v} \leq \vec{y}^*\}$ is approachable.

Definition 2 (Blackwell's Approachability [4]). *Given a convex set \mathcal{R} of vectors of utilities, we say \mathcal{R} is approachable from the perspective of the follower, if for any strategy of the leader \vec{s}_1 , there exists an adaptive strategy \vec{s}_2 for the follower such that $\lim_{T \rightarrow \infty} \text{dist}\left(\frac{1}{T} \sum_{t=1}^T \vec{\mathcal{U}}(s_1^t, s_2^t), \mathcal{R}\right) = 0$ almost surely, where $\vec{\mathcal{U}}(s_1, s_2) = (\mathcal{U}^{\theta_1}(s_1, s_2), \dots, \mathcal{U}^{\theta_{|\Theta|}}(s_1, s_2))$ and $\text{dist}(\vec{u}, \mathcal{R}) = \min_{\vec{v} \in \mathcal{R}} \|\vec{v} - \vec{u}\|$.*

Theorem 2 ([2,18]). $\mathcal{R}(\vec{y}^*) = \{\vec{v} \mid \vec{v} \leq \vec{y}^*\}$ is approachable.

To establish the approachability of $\mathcal{R}(\vec{y}^*)$ of Theorem 2, we first consider a simple case in which the convex set is a halfspace $\mathcal{H}(\vec{\Pi}', b)$ such that $\vec{v} \in \mathcal{H}(\vec{\Pi}', b)$ if and only if $\langle \vec{\Pi}', \vec{v} \rangle \leq b$.

Lemma 2. *A halfspace $\mathcal{H}(\vec{\Pi}', b)$ is approachable if $f(\Pi') \leq b$.*

Theorem 3 ([4]). *A convex set \mathcal{R} is approachable if and only if all halfspaces containing \mathcal{R} are approachable.*

All that remains to show is that all halfspaces containing $\mathcal{R}(\vec{y}^*)$ are approachable.

Lemma 3. *All halfspaces containing $\mathcal{R}(\vec{y}^*) = \{\vec{v} \mid \vec{v} \leq \vec{y}^*\}$ are approachable.*

Proof. Notice that any minimal halfspace containing $\mathcal{R}(\bar{y}^*)$ must cross \bar{y}^* by the construction of $\mathcal{R}(\bar{y}^*)$. Therefore, such a halfspace can be represented by $\mathcal{H}(\vec{\Pi}', \langle \vec{\Pi}', \bar{y}^* \rangle)$ with $\Pi' \in \Delta(\Theta)$. By Proposition 1, $f(\Pi') \leq \langle \vec{\Pi}', \bar{y}^* \rangle$, and therefore, by Lemma 2, $\mathcal{H}(\vec{\Pi}', \langle \vec{\Pi}', \bar{y}^* \rangle)$ is approachable.

Combining Theorem 3 and Lemma 3, we finish the proof of Theorem 2. We can then apply Blackwell's construction [4] to obtain an adaptive strategy for the follower that approaches $\mathcal{R}(\bar{y}^*)$ almost surely.

Intuitively, at stage t , if $\frac{1}{t-1} \sum_{\tau=1}^{t-1} \vec{U}(s_1^\tau, s_2^\tau) \notin \mathcal{R}(\bar{y}^*)$, then the follower first finds a halfspace $\mathcal{H}(\vec{\Pi}', \langle \vec{\Pi}', \bar{y}^* \rangle)$ that separates $\frac{1}{t-1} \sum_{\tau=1}^{t-1} \vec{U}(s_1^\tau, s_2^\tau)$ and $\mathcal{R}(\bar{y}^*)$. Given such a Π' , the follower plays the minimax strategy of $\mathcal{U}^{\Pi'}$ at stage t , and then the distance between the vector of average utilities and $\mathcal{R}(\bar{y}^*)$ will become smaller after stage t . Observe that the follower's strategy can be computed from the prior Π , the game G_{repeated} , and the history of play. In doing so, it guarantees that the expected average utility of the leader is at most $\nu_{\text{signal}}(G_{\text{signal}})$ in the limit, and Proposition 2 follows:

Proposition 2. *In a Bayesian repeated game $G_{\text{repeated}} = (\Pi, \{\mathcal{U}^\theta\}_{\theta \in \Theta})$, given \bar{y}^* satisfying Proposition 1 and an oracle to compute the minimax strategy of the zero-sum game $\mathcal{U}^{\Pi'}$ for all $\Pi' \in \Delta(\Theta)$, there exists an efficient algorithm to construct the follower's optimal strategy.*

We will elaborate the complexity of computing \bar{y}^* in Section 3.3.

3.2 Equivalence between Bayesian Allocation Games and Team Max-min Games

Definition 3 (Team Max-min Game [19]). *In a zero-sum team max-min game $G_{\text{team}} = (\{\mathcal{U}^\theta\}_{\theta \in \Theta})$, in addition to player 1 and 2, there is a player 3 whose set of pure strategies is Θ . Player 1 and player 3 form a team and share the same utility such that when player 1 plays $s_1 \in S_1$, player 2 plays $s_2 \in S_2$, and player 3 plays $\theta \in \Theta$, the utility for both player 1 and player 3 is $\mathcal{U}^\theta(s_1, s_2)$, while the utility for player 2 is $-\mathcal{U}^\theta(s_1, s_2)$. A team max-min equilibrium is a Nash equilibrium that maximizes the team's utility and we denote its game value by $\nu_{\text{team}}(G_{\text{team}})$: $\nu_{\text{team}}(G_{\text{team}}) = \max_{s_1 \in \Delta(S_1), \Pi \in \Delta(\Theta)} \min_{s_2 \in \Delta(S_2)} \mathcal{U}^\Pi(s_1, s_2)$.*

We emphasize that player 1's strategy and player 3's strategy are not allowed to be correlated; otherwise, the team max-min game degenerates to a classic two-player zero-sum game in which player 1 and 3 can be treated as a single player. [19] show that a team max-min equilibrium always exists. It turns out the equilibrium in Bayesian allocation games corresponds to the solution of the above team max-min games in a precise sense, stated below.

Theorem 4. $\nu_{\text{alloc}}(G_{\text{alloc}}) = \nu_{\text{team}}(G_{\text{team}})$ when $G_{\text{alloc}} = G_{\text{team}}$.

To prove Theorem 4, we will construct strategies for players in the Bayesian allocation game from the equilibrium strategies in the team max-min game.

Lemma 4. When $G_{alloc} = G_{team}$, let s_1^*, s_2^*, Π^* be the equilibrium strategies for the team max-min game G_{team} . In the Bayesian allocation game G_{alloc} , consider the following strategy for the leader:

- set the prior Π to be Π^* ;
- repeatedly play strategy s_1^* for every stage.

This strategy can guarantee the leader an average expected utility $\nu_{team}(G_{team})$.

In comparison to the Bayesian repeated games in which the follower knows the prior, the follower does not even know the prior set by the leader in the Bayesian allocation game. To overcome this obstacle, observe that in the Bayesian repeated game, the approachability of a convex set is a property that only depends on the collection of games $(\{\mathcal{U}^\theta\}_{\theta \in \Theta})$ but independent of the prior Π . Motivated by this observation, we show that

$$\mathcal{R}(\nu_{team}(G_{team}) \cdot \vec{1}) = \{\vec{v} \mid \vec{v} \leq \nu_{team}(G_{team}) \cdot \vec{1}\}$$

is approachable where $\vec{1}$ is a vector of all ones. Note that if such a convex set is approachable, for any prior $\Pi \in \Delta(\Theta)$, the average utility of the leader is

$$\lim_{T \rightarrow \infty} \sum_{\theta \in \Theta} \Pi(\theta) \cdot \frac{\sum_{t=1}^T \mathcal{U}^\theta(s_1^t, s_2^t)}{T} \leq \sum_{\theta \in \Theta} \Pi(\theta) \cdot \nu_{team}(G_{team}) = \nu_{team}(G_{team}).$$

Lemma 5. $\mathcal{R}(\nu_{team}(G_{team}) \cdot \vec{1})$ is approachable.

Proof. To show $\mathcal{R}(\nu_{team}(G_{team}) \cdot \vec{1})$ is approachable, by Theorem 3, it suffices to show that any halfspaces containing $\mathcal{R}(\nu_{team}(G_{team}) \cdot \vec{1})$ is approachable. Note that any minimal halfspace containing $\mathcal{R}(\nu_{team}(G_{team}) \cdot \vec{1})$ must cross $\nu_{team}(G_{team}) \cdot \vec{1}$. Therefore, such a halfspace can be represented by

$$\mathcal{H}(\vec{\Pi}, \langle \vec{\Pi}, \nu_{team}(G_{team}) \cdot \vec{1} \rangle) = \mathcal{H}(\vec{\Pi}, \nu_{team}(G_{team})).$$

By Lemma 2, it is enough to show that for all $\Pi \in \Delta(\Theta)$, $f(\Pi) \leq \nu_{team}(G_{team})$. Recall that $\nu_{team}(G_{team})$ is the game value of the team max-min game, and thus,

$$\nu_{team}(G_{team}) = \max_{\Pi} \left[\max_{s_1} \min_{s_2} \mathcal{U}^\Pi(s_1, s_2) \right] = \max_{\Pi} f(\Pi). \quad (4)$$

Thus, $\nu_{team}(G_{team}) \geq f(\Pi)$ for any $\Pi \in \Delta(\Theta)$, which concludes the proof.

3.3 Computational Complexity of the Follower's Optimal Strategy

As demonstrated before, constructing the follower's optimal strategy in Bayesian repeated games requires a solution to the dual program (2). Hence, it is not immediate that one can efficiently construct the follower's optimal strategy if the public signaling game is efficiently solvable. Here, we say an algorithm is efficient if the running time of the algorithm is polynomial in terms of the number of states $|\Theta|$, and the number of pure strategies $|S_1| + |S_2|$.




Fig. 1. The relationships of computational problems, assuming the minimax strategy of \mathcal{U}^I can be computed efficiently for all $I \in \Delta(\Theta)$: the arrows point to problems that are computationally easier.

We manage to show that, when the minimax strategy of \mathcal{U}^I can be computed efficiently for all $I \in \Delta(\Theta)$, in both Bayesian repeated games and Bayesian allocation games, the follower's optimal strategy can be efficiently constructed if the corresponding game values are given. We further show that team max-min game is computationally easier than the public signaling game, and therefore, Bayesian allocation game is computationally easier than the Bayesian repeated game. Figure 1 summarizes the relationships of the computational problems discussed in this section, while the proofs are deferred to Appendix B.4.

4 Bayesian Repeated Security Games

In Section 3, we have shown that Bayesian repeated games can be reduced to public signaling games, while Bayesian allocation games can be reduced to team max-min games. However, it has been shown that both public signaling games and team max-min games are computationally intractable for general zero-sum games and even worse, no FPTAS is possible [5,8]. Particularly, public signaling games do not even admit PTAS [3,17].

Motivated by the applications in the domain of repeated security games, we will concern ourselves with repeated games where the stage game is a *security game* of a particularly simple form. The one-shot complete-information security games are described by a set L of *locations*, a set M of *treasures*, and a set R of *defensive resources*. For convenience, we use \perp to denote a *null* treasure or a *null* defensive resource. $v : L \times (M \cup \perp) \rightarrow \mathbb{R}_{\geq 0}$ is a *location-treasure importance* function such that $v(\ell, m)$ characterizes the utility loss of the defender if location $\ell \in L$ with treasure $m \in M$ allocated is attacked without defense. In addition, there is a *defense-quality* function $q : L \times (M \cup \perp) \times (R \cup \perp) \rightarrow \{0, 1\}$ such that $q(\ell, m, r)$ characterizes the effectiveness of allocating defensive resource $r \in R$ to defend location $\ell \in L$ that hosts treasure m . Note that in our setting, a defensive resource is either 100% effective for a combination of location and treasure or totally useless. For a *null* treasure, we have $v(\ell, \perp) = 0$ for all ℓ , and for a *null* defensive resource, we have $q(\ell, m, \perp) = 0$ for all ℓ and m .

A state of nature is a matching $\theta : L \rightarrow M$ that maps the locations to treasures such that for any $i, j \in L$ with $i \neq j$, $\theta(i) \neq \perp$, and $\theta(j) \neq \perp$, we

have $\theta(i) \neq \theta(j)$. A pure strategy for the defender is also a matching $D : L \rightarrow R$ that maps the locations to the defensive resources such that for any $i, j \in L$ with $i \neq j$, $D(i) \neq \perp$, and $D(j) \neq \perp$, we have $D(i) \neq D(j)$. Finally, a pure strategy for the attacker is a single location $a \in L$ to attack. A mixed strategy is naturally a distribution over such functions. The defender's utility under θ when the defender plays D and the attacker plays a is $\mathcal{U}^\theta(D, a) = -(1 - q(a, \theta(a), D(a))) \cdot v(a, \theta(a))$, while the attacker's utility is simply $-\mathcal{U}^\theta(D, a)$.

We say the treasures are *homogeneous* if for all $m \in M$, $v(\ell, m)$ equals to some constant for all $\ell \in L$; the locations are *homogeneous* if for all $\ell \in L$, $v(\ell, m)$ equals to some constant for all $m \in M$; and the defensive resources are *homogeneous* if $q(\ell, m, r) = 1$ for all $\ell \in L$, $m \in M$, and $r \in R$. If the condition of homogeneity is not satisfied, we say they are *heterogeneous*.

We analyze the complexity of repeated security games under the contexts of both Bayesian repeated games and Bayesian allocation games. In Bayesian repeated games, an algorithm is efficient if its running time is in polynomial of $|\Theta|$, $|L|$, $|M|$, and $|R|$; while in Bayesian allocation games, an algorithm is efficient if its running time is in polynomial of $|L|$, $|M|$, and $|R|$.

The next proposition illustrates that one can efficiently compute the minimax strategy of the security game \mathcal{U}^Π given the marginals of a prior Π .

Proposition 3. *Given the marginals of Π , the optimal strategies for both the defender and the attacker in the security game \mathcal{U}^Π can be computed efficiently.*

The proofs in this section are deferred to Appendix C. However, for our class of security games with a general prior Π , computing the game value of the Bayesian repeated games is computationally intractable.

Theorem 5. *It is strongly NP-hard to compute the game value of the Bayesian repeated games with a security game as the stage game, even when all of treasures, locations, and defensive resources are homogeneous. Moreover, no FPTAS is possible. Consequently, it is strongly NP-hard to compute any representation of the equilibrium which permits computing the game value.*

5 Computational Complexity for Security Bayesian Allocation Games

In this section, we turn to Bayesian allocation games with a security game as the stage game. The proofs in this section are deferred to Appendix D. It turns out that a Bayesian allocation game with a security game as the stage game can be efficiently solved when only the treasures are heterogeneous.

Theorem 6. *There exists an efficient algorithm to compute the game value and the defender's optimal strategy of a Bayesian allocation game with a security game as the stage game, when only the treasures are heterogeneous.*

Moreover, the following lemma illustrates that one can efficiently construct the attacker's strategy efficiently when the game value is given.

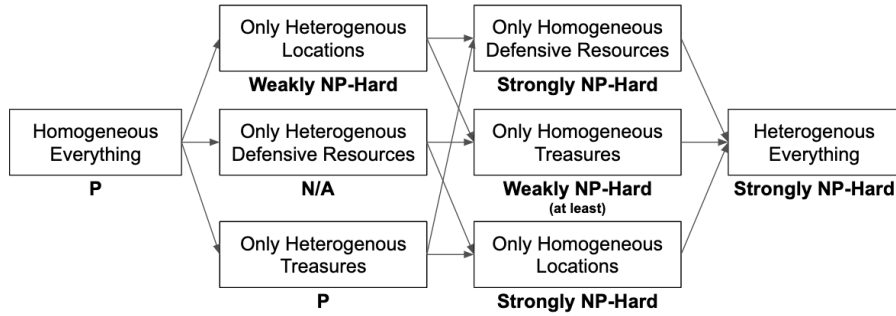


Fig. 2. The computational complexity of Bayesian allocation games with a security game as the stage game: the arrows point to more general versions of the problem.

Lemma 6. *Given the game value of a Bayesian allocation game with a security game as the stage game, there exists an efficient algorithm to compute the attacker’s optimal strategy.*

Therefore, one can efficiently construct both the defender’s optimal strategy and the attacker’s optimal strategy when only the treasures are heterogeneous. However, going beyond, the problem becomes computationally intractable.

Theorem 7. *It is weakly NP-hard to compute the game value of the Bayesian allocation games with a security game as the stage game, when only the locations are heterogeneous. Moreover, there exists a pseudo-polynomial time algorithm that can compute the game value.*

Theorem 8. *It is strongly NP-hard to compute the game value of the Bayesian allocation games with a security game as the stage game, when only the defensive resources are homogeneous, or only the locations are homogeneous.*

There are three other settings that have not been discussed: (1) heterogeneous everything; (2) only treasures are homogeneous; and (3) only defensive resources are heterogeneous. For the setting in which everything is heterogeneous, it is also strongly NP-hard to compute the game value since it is a more general setting than the settings in which only defensive resources are homogeneous or only locations are homogeneous. As for the setting in which only treasures are homogeneous, it is at least weakly NP-hard to compute the game value since it is a more general setting than the case in which only locations are heterogeneous. We leave it as an open question to settle whether it is strongly NP-hard. Finally, for the setting in which only defensive resources are heterogeneous, this setting is not well-defined: since the locations and the treasures are homogeneous, a defensive resource should be either effective or ineffective for any combination of the locations and the treasures. Consequently, the defender can simply eliminate the ineffective defensive resources to focus on effective ones, which reduces the problem to the case in which everything is homogeneous.

References

1. Abernethy, J., Bartlett, P.L., Hazan, E.: Blackwell approachability and no-regret learning are equivalent. In: Proceedings of the 24th Annual Conference on Learning Theory. pp. 27–46 (2011)
2. Aumann, R.J., Maschler, M.: Repeated games with incomplete information. MIT press (1995)
3. Bhaskar, U., Cheng, Y., Ko, Y.K., Swamy, C.: Hardness results for signaling in bayesian zero-sum and network routing games. In: Proceedings of the 2016 ACM Conference on Economics and Computation. pp. 479–496 (2016)
4. Blackwell, D.: An analog of the minimax theorem for vector payoffs. *Pacific Journal of Mathematics* **6**(1), 1–8 (1956)
5. Borgs, C., Chayes, J., Immorlica, N., Kalai, A.T., Mirrokni, V., Papadimitriou, C.: The myth of the folk theorem. *Games and Economic Behavior* **70**(1), 34–43 (2010)
6. Cheng, Y., Cheung, H.Y., Dughmi, S., Emamjomeh-Zadeh, E., Han, L., Teng, S.H.: Mixture selection, mechanism design, and signaling. In: 2015 IEEE 56th Annual Symposium on Foundations of Computer Science. pp. 1426–1445. IEEE (2015)
7. Dughmi, S.: Algorithmic information structure design: a survey. *ACM SIGecom Exchanges* **15**(2), 2–24 (2017)
8. Dughmi, S.: On the hardness of designing public signals. vol. 118, pp. 609–625. Elsevier (2019)
9. Fang, F., Nguyen, T.H.: Green security games: Apply game theory to addressing green security challenges. *ACM SIGecom Exchanges* **15**(1), 78–83 (2016)
10. Fang, F., Stone, P., Tambe, M.: When security games go green: Designing defender strategies to prevent poaching and illegal fishing. In: Twenty-Fourth International Joint Conference on Artificial Intelligence (2015)
11. Grötschel, M., Lovász, L., Schrijver, A.: The ellipsoid method and its consequences in combinatorial optimization. *Combinatorica* **1**(2), 169–197 (1981)
12. Immorlica, N., Kalai, A.T., Lucier, B., Moitra, A., Postlewaite, A., Tennenholtz, M.: Dueling algorithms. In: Proceedings of the forty-third annual ACM symposium on Theory of computing. pp. 215–224 (2011)
13. Kamenica, E., Gentzkow, M.: Bayesian persuasion. *American Economic Review* **101**(6), 2590–2615 (2011)
14. Lee, Y.T., Sidford, A., Vempala, S.S.: Efficient convex optimization with membership oracles. In: Conference On Learning Theory. pp. 1292–1294 (2018)
15. Neumann, J.v.: Zur theorie der gesellschaftsspiele. *Mathematische annalen* **100**(1), 295–320 (1928)
16. Rockmore, D.: What’s missing from “the imitation game”. <https://www.newyorker.com/tech/annals-of-technology/imitation-game-alan-turing> (November 2014), accessed: 2020-01-23
17. Rubinstein, A.: Eth-hardness for signaling in symmetric zero-sum games. *CoRR abs/1510.04991* (2015)
18. Sorin, S.: A first course on zero-sum repeated games, vol. 37. Springer Science & Business Media (2002)
19. von Stengel, B., Koller, D.: Team-maxmin equilibria. *Games and Economic Behavior* **21**(1-2), 309–321 (1997)
20. Tambe, M.: Security and game theory: algorithms, deployed systems, lessons learned. Cambridge university press (2011)
21. Wang, Y., Shi, Z.R., Yu, L., Wu, Y., Singh, R., Joppa, L., Fang, F.: Deep reinforcement learning for green security games with real-time information. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 33, pp. 1401–1408 (2019)

22. Xu, H., Ford, B., Fang, F., Dilkina, B., Plumtre, A., Tambe, M., Driciru, M., Wanyama, F., Rwetsiba, A., Nsubaga, M., et al.: Optimal patrol planning for green security games with black-box attackers. In: International Conference on Decision and Game Theory for Security. pp. 458–477. Springer (2017)
23. Xu, H., Wang, K., Vayanos, P., Tambe, M.: Strategic coordination of human patrollers and mobile sensors with signaling for security games. In: Thirty-Second AAAI Conference on Artificial Intelligence (2018)

Appendix

A Utility and Equilibrium Model

There are three main utility/equilibrium models when agents are interested in their long-term payoffs. All three are equivalent for repeated two-player zero-sum games with one-sided private information, whether in the classical or allocation models.

- Asymptotic finitely-repeated game: Each player’s utility is the average of his utilities from the first T stages of the game. Since the stage game is finite/compact, so is the T -stage game, and there is an essentially unique equilibrium (up to payoff equivalence): the minimax equilibrium. The value of the game is therefore well defined. We take the limit of the equilibrium and the game value as $T \rightarrow \infty$.
- Asymptotic time-discounted game: In the γ -discounted repeated game for $\gamma \in (0, 1)$, each player’s (γ -discounted) utility is the sum, over all stages $t = 1, \dots$, of his utility from the t -th stage game multiplied by $(1 - \gamma)\gamma^{t-1}$. Through an appropriate choice of topology for each player’s mixed strategy space, the γ -discounted game is a compact zero-sum game of incomplete information. Again, the minimax equilibrium is the essentially unique equilibrium, and the value of the game is well defined. We take the limit of the equilibrium and the value of the game as the discount factor γ approaches 1.
- Infinitely-repeated game: Here, each player’s expected utility is the limit, as $T \rightarrow \infty$, of his average expected utility over the first T stages alone. Though this limit may not exist in general, we can nevertheless sometimes define a value and equilibrium as in [2,18]. The max-min value of the game is the supremum over all player 1’s mixed strategies, of the infimum over player 2’s mixed strategies, of the limit infimum as $T \rightarrow \infty$ of player 1’s average expected utility. Player 1’s max-min strategy is that attaining this supremum. We can similarly define the min-max value of the game and Player 2’s min-max strategy. When both the max-min and min-max values are equal we refer to them as the value of the game, and the corresponding max-min and min-max strategies form the equilibrium.

The following theorem can be gleaned from a careful reading of [2,18].

Theorem 9 ([2,18]). *Consider a two-player Bayesian repeated zero-sum game with incomplete information on one side. The value of the asymptotic finitely-repeated game, the value of the asymptotic time-discounted game, and both values of the infinitely repeated game, are all equal. Moreover, the associated equilibria for these three models coincide.*

Though this theorem is stated for the classical model of repeated games, it is easy to see that it applies equally well to the allocation model.

B Omitted Materials in Section 3

B.1 Proof of Lemma 1

Proof. Note that after obtaining the signal σ , the leader's strategy per stage is non-adaptive, denoted by $s_1^*(\sigma)$. Therefore, the follower's strategy can only depend on the game $G_{\text{repeated}} = G_{\text{signal}}$ as well as the signal σ after observing the leader's strategy. Moreover, given each signal $\sigma \in \Sigma$, the follower's belief is exactly updated to $\Pi_{\varphi, \sigma}$. As a result, conditioned on a signal σ , to minimize the leader's utility, the follower's optimal strategy is

$$s_2^*(\sigma) = \operatorname{argmin}_{s_2} \max_{s_1} \mathcal{U}^{\Pi_{\varphi, \sigma}}(s_1, s_2).$$

By the minimax theorem, we have

$$\mathcal{U}^{\Pi_{\varphi, \sigma}}(s_1^*(\sigma), s_2^*(\sigma)) = \max_{s_1} \min_{s_2} \mathcal{U}^{\Pi_{\varphi, \sigma}}(s_1, s_2).$$

Thus, the expected utility of the leader is at least

$$\sum_{\theta \in \Theta} \Pi(\theta) \sum_{\sigma \in \Sigma} \Pr[\varphi(\theta) = \sigma] \cdot \mathcal{U}^{\theta}(s_1^*(\sigma), s_2^*(\sigma)),$$

which equals to

$$\sum_{\sigma \in \Sigma} \left(\sum_{\theta \in \Theta} \Pr[\varphi(\theta) = \sigma] \cdot \Pi(\theta) \right) \cdot \mathcal{U}^{\Pi_{\varphi, \sigma}}(s_1^*(\sigma), s_2^*(\sigma)) = \nu_{\text{signal}}(G_{\text{signal}}),$$

and concludes the proof.

B.2 Proof of Lemma 2

Proof. Let $\vec{v}^{t-1} = \frac{1}{t-1} \sum_{\tau=1}^{t-1} \vec{\mathcal{U}}(s_1^\tau, s_2^\tau)$ be the vector of the leader's average utilities for the first $(t-1)$ stages. If $\vec{v}^{t-1} \in \mathcal{H}(\vec{\Pi}', b)$, then the follower may play an arbitrary strategy at stage t . Otherwise, when $\vec{v}^{t-1} \notin \mathcal{H}(\vec{\Pi}', b)$, we have $\langle \vec{\Pi}', \vec{v}^{t-1} \rangle > b$. Recall that by minimax theorem,

$$f(\Pi') = \min_{s_2} \max_{s_1} \mathcal{U}^{\Pi'}(s_1, s_2) = \min_{s_2} \max_{s_1} \langle \vec{\Pi}', \vec{\mathcal{U}}(s_1, s_2) \rangle \leq b.$$

Therefore, when the follower plays the minimax strategy for the game Π' , i.e., $s_2^* = \operatorname{argmin}_{s_2} \max_{s_1} \mathcal{U}^{\Pi'}(s_1, s_2)$, we have $\langle \vec{\Pi}', \vec{\mathcal{U}}(s_1, s_2^*) \rangle \leq b$ for any strategy s_1 of the leader. Thus, the leader's average utilities for the first t rounds will be $\vec{v}^t = (1 - \frac{1}{t})\vec{v}^{t-1} + \frac{1}{t}\vec{\mathcal{U}}(s_1, s_2^*)$, and we have the distance between \vec{v}^t and $\mathcal{H}(\vec{\Pi}', b)$ is closer than the distance between \vec{v}^{t-1} and $\mathcal{H}(\vec{\Pi}', b)$:

$$\begin{aligned} \langle \vec{\Pi}', \vec{v}^t \rangle - b &= (1 - \frac{1}{t})\langle \vec{\Pi}', \vec{v}^{t-1} \rangle + \frac{1}{t}\langle \vec{\Pi}', \vec{\mathcal{U}}(s_1, s_2^*) \rangle - b \\ &< (1 - \frac{1}{t})(\langle \vec{\Pi}', \vec{v}^{t-1} \rangle - b). \end{aligned}$$

B.3 Proof of Lemma 4

Proof. Since the leader's strategy is independent of the realization of the true state, the follower can learn nothing about the states from the leader's strategy. Therefore, the follower has no information about the realized state (even the prior of the states). Note that in a Nash equilibrium of the team max-min game, fixing player 3's strategy Π^* , player 1 must play the maximin strategy in game \mathcal{U}^{Π^*} since \mathcal{U}^{Π^*} is a zero-sum game. Hence, for every stage, we have $\mathcal{U}^{\Pi^*}(s_1^*, s_2) \geq \max_{s_1} \min_{s_2} \mathcal{U}^{\Pi^*}(s_1, s_2) = \nu_{\text{team}}(G_{\text{team}})$ for all $s_2 \in \Delta(S_2)$.

B.4 Computational Complexity of the Follower's Optimal Strategy

Given a prior $\Pi \in \Delta(\Theta)$, let $\tilde{f}(\Pi) = \nu_{\text{signal}}(G_{\text{signal}})$ with be the game value of the public signaling game $G_{\text{signal}} = (\Pi, \{\mathcal{U}^\theta\}_{\theta \in \Theta})$. We claim that the dual program (2) is equivalent to the following linear program (5):

$$\begin{aligned} \min \quad & \langle \vec{y}, \vec{\Pi} \rangle \\ \text{s.t.} \quad & \langle \vec{y}, \vec{\Pi}' \rangle \geq \tilde{f}(\Pi') \quad \forall \Pi' \in \Delta(\Theta) \end{aligned} \quad (5)$$

in which we replace $f(\Pi')$ with $\tilde{f}(\Pi')$.

Lemma 7. $\tilde{f}(\Pi)$ is concave in Π and the linear program (2) is equivalent to the program (5).

Proof. First notice for all $\Pi' \in \Delta(\Theta)$, we have $\tilde{f}(\Pi') \geq f(\Pi')$ because $x(\Pi') = 1$ and $x(\Pi'') = 0$ for any $\Pi'' \neq \Pi'$ is a feasible solution to the primal (1) corresponding to Π' , i.e., player 1's utility in the public signaling game $G_{\text{signal}} = (\Pi', \{\mathcal{U}^\theta\}_{\theta \in \Theta})$ is at least $f(\Pi')$ if the principal does not signal anything. Hence, any feasible solution of the program (5) is a feasible solution of the program (2). On the other hand, according to the primal (1), $\tilde{f}(\Pi')$ can be written as

$$\tilde{f}(\Pi') = \sum_{\Pi'' \in \Delta(\Theta)} x_{\Pi'}(\Pi'') \cdot f(\Pi'') \quad (6)$$

where $\sum_{\Pi'' \in \Delta(\Theta)} x_{\Pi'}(\Pi'') \cdot \Pi'' = \Pi'$. Therefore, for any Π' and any feasible solution \vec{y} of the program (2), we have that

$$\begin{aligned} \langle \vec{y}, \vec{\Pi}' \rangle &= \sum_{\theta \in \Theta} y(\theta) \cdot \sum_{\Pi'' \in \Delta(\Theta)} x_{\Pi'}(\Pi'') \cdot \Pi''(\theta) \\ &= \sum_{\Pi'' \in \Delta(\Theta)} x_{\Pi'}(\Pi'') \cdot \langle \vec{y}, \vec{\Pi}'' \rangle \\ &\geq \sum_{\Pi'' \in \Delta(\Theta)} x_{\Pi'}(\Pi'') \cdot f(\Pi'') = \tilde{f}(\Pi') \end{aligned}$$

where the inequality follows that \vec{y} is feasible to the program (2). Thus, any feasible solution of the program (2) is a feasible solution of the program (5), which concludes the proof of the lemma.

For the concavity of \tilde{f} , notice that for any feasible solution $\vec{x}_{\Pi'}$ and $\vec{x}_{\Pi''}$ of the primal (1) corresponding to prior Π' and Π'' , respectively, we have that $\alpha \cdot \vec{x}_{\Pi'} + (1 - \alpha) \cdot \vec{x}_{\Pi''}$ is a feasible solution of the primal (1) corresponding to prior $\alpha \cdot \Pi' + (1 - \alpha) \cdot \Pi''$; and therefore, $\tilde{f}(\alpha \cdot \Pi' + (1 - \alpha) \cdot \Pi'') \geq \alpha \cdot \tilde{f}(\Pi') + (1 - \alpha) \cdot \tilde{f}(\Pi'')$.

We are now ready to show that one can efficiently construct the follower's optimal strategy for the Bayesian repeated game if the corresponding public signaling game is efficiently solvable.

Lemma 8. *For a collection of games $\{\mathcal{U}^\theta\}_{\theta \in \Theta}$, if the game value of the public signaling game $G_{\text{signal}} = (\Pi, \{\mathcal{U}^\theta\}_{\theta \in \Theta})$ and the minimax strategy of the zero-sum game \mathcal{U}^Π can be computed efficiently for all $\Pi \in \Delta(\Theta)$, then there exists an efficient algorithm to compute the follower's optimal strategy for the Bayesian repeated game $G_{\text{repeated}} = (\Pi, \{\mathcal{U}^\theta\}_{\theta \in \Theta})$ for any $\Pi \in \Delta(\Theta)$.*

Proof. We will show that (5) admits an efficient separation oracle, and thus, it can be solved efficiently. Given a candidate solution \vec{y} , $\tilde{f}(\Pi')$ is concave in Π' according to Lemma 7 while $\langle \vec{y}, \vec{\Pi}' \rangle$ is linear in Π' , which implies that $h_{\vec{y}}(\Pi') = \langle \vec{y}, \vec{\Pi}' \rangle - \tilde{f}(\Pi')$ is a convex function in Π' . The objective of the separation oracle is to identify a Π' such that $h_{\vec{y}}(\Pi')$ is less than 0, which is indeed a convex optimization problem. Therefore, to implement a separation oracle is equivalent to minimize a convex function $h_{\vec{y}}$ with an access to an efficient evaluation oracle of $h_{\vec{y}}$, which can be done efficiently (see [11,14]). Thus, we can obtain \vec{y}^* satisfying Proposition 1 efficiently. Combining with the condition that the minimax strategy of the zero-sum game $\mathcal{U}^{\Pi'}$ can be computed efficiently for any $\Pi' \in \Delta(\Theta)$, we can construct the follower's strategy efficiently due to Proposition 2.

Lemma 8 suggests that the existence of an efficient algorithm for the primal (1) implies the existence of an efficient algorithm for the dual (2). As discussed in Section 3.2, for Bayesian allocation games, $\mathcal{R}(\nu_{\text{team}}(G_{\text{team}}) \cdot \vec{1})$ can be constructed directly from the game value of the team max-min game, and thus, we can then construct the follower's optimal strategy efficiently once the minimax strategy of the zero-sum game \mathcal{U}^Π can be computed efficiently for all $\Pi \in \Delta(\Theta)$.

Proposition 4. *For a collection of games $\{\mathcal{U}^\theta\}_{\theta \in \Theta}$, if the game value of the team max-min game $G_{\text{team}} = (\{\mathcal{U}^\theta\}_{\theta \in \Theta})$ can be computed efficiently, and the minimax strategy of the zero-sum game \mathcal{U}^Π can be computed efficiently for all $\Pi \in \Delta(\Theta)$, then there exists an efficient algorithm to compute the follower's optimal strategy for the Bayesian allocation game $G_{\text{alloc}} = (\{\mathcal{U}^\theta\}_{\theta \in \Theta})$.*

Our next lemma shows that team max-min games are in general computationally easier than public signaling games, and therefore, Bayesian allocation games are in general computationally easier than Bayesian repeated games.

Lemma 9. *For a collection of games $\{\mathcal{U}^\theta\}_{\theta \in \Theta}$, if the game value of the public signaling game $G_{\text{signal}} = (\Pi, \{\mathcal{U}^\theta\}_{\theta \in \Theta})$ and the minimax strategy of the zero-sum game \mathcal{U}^Π can be computed efficiently for all $\Pi \in \Delta(\Theta)$, then there exists an efficient algorithm to compute the equilibrium strategies of the team max-min game $G_{\text{team}} = (\{\mathcal{U}^\theta\}_{\theta \in \Theta})$.*

Proof. Recall that $\nu_{\text{team}}(G_{\text{team}}) = \max_{\Pi} f(\Pi)$ from (4) and moreover, we have $\max_{\Pi} f(\Pi) = \max_{\Pi} \tilde{f}(\Pi)$ since for $\Pi^* \in \operatorname{argmax}_{\Pi} f(\Pi)$, it is sub-optimal to have a solution \tilde{x}_{Π^*} in which there exists $x_{\Pi^*}(\Pi') > 1$ with $f(\Pi') < f(\Pi^*)$ to the primal (1) corresponding to Π^* . Therefore, to compute the game value of the Bayesian allocation game is equivalent to maximize the concave function \tilde{f} with an access to an efficient evaluation oracle of \tilde{f} , which can be done efficiently (see [11, 14]). Once Π^* (the equilibrium strategy of player 3) is identified, the equilibrium strategies of player 1 and 2 in the team max-min game are simply the equilibrium strategies of game \mathcal{U}^{Π^*} .

C Omitted Materials in Section 4

C.1 Proof of Proposition 3

Before we delve into the computational complexity of repeated games, we first illustrate, given a posterior Π , how to compute the equilibrium strategies of \mathcal{U}^Π efficiently. A mixed strategy for the defender can be summarized by its *marginals*, denoted by $d_{\ell,r}$, the probability of defending location $\ell \in L$ with defensive resource $r \in R$. The set of marginals of the defender's strategies is given by the bipartite matching polytope:

$$\mathcal{D} = \left\{ \vec{d} \in [0, 1]^{|L| \times |R|} \mid \forall \ell \in L, \sum_{r \in R} d_{\ell,r} \leq 1 \text{ and } \forall r \in R, \sum_{\ell \in L} d_{\ell,r} \leq 1 \right\}.$$

The marginals of a mixed strategy of the defender suffice to compute the payoffs. Moreover, we can compute marginals and mixed strategies from each other [11]. Henceforth, we refer to the marginals and mixed strategies interchangeably. The defender's utility when he plays $\vec{d} \in \mathcal{D}$ while the attacker plays a mixed strategy $\vec{a} \in \mathcal{A} = \Delta^{|L|-1}$ where $\Delta^{|L|-1}$ is the simplex with $|L| - 1$ dimensions (such that a_ℓ indicates the probability of the attacker attacking ℓ) can be computed as

$$-\sum_{\theta \in \Theta} \Pi(\theta) \cdot \sum_{\ell \in L} a_\ell \cdot v(\ell, \theta(\ell)) \cdot (1 - \sum_{r \in R} d_{\ell,r} \cdot q(\ell, \theta(\ell), r)).$$

As above, a posterior Π can be summarized by marginals $\vec{g}(\Pi) \in \mathcal{G}$ where \mathcal{G} is a bipartite matching polytope

$$\mathcal{G} = \left\{ \vec{g} \in [0, 1]^{|L| \times |M|} \mid \forall \ell \in L, \sum_{m \in M} g_{\ell,m} \leq 1 \text{ and } \forall m \in M, \sum_{\ell \in L} g_{\ell,m} \leq 1 \right\}$$

between L and M , such that $g_{\ell,m}(\Pi) = \sum_{\theta:\theta(\ell)=m} \Pi(\theta)$. Therefore, the defender's utility can be written as

$$-\sum_{\ell \in L} \sum_{m \in M} g_{\ell,m}(\Pi) \cdot a_{\ell} \cdot v(\ell, m) \cdot \left(1 - \sum_{r \in R} d_{\ell,r} \cdot q(\ell, m, r)\right).$$

The bilinear nature of this expression in \vec{d} and \vec{a} , plus the tractability of polytopes \mathcal{D} and \mathcal{A} , imply that the equilibrium can be computed efficiently given a fixed posterior Π since the game \mathcal{U}^{Π} is a tractable *zero-sum dueling game* [12].

C.2 Proof of Theorem 5

Definition 4 (3-Set Cover). *In an instance of 3-SET COVER, we are given a finite set of elements U with $|U| = 3n$ and a family of Z subsets $E_1, \dots, E_Z \subseteq U$ with $|E_i| = 3$ for all $i \in [Z]$. The task is to determine whether there exists a family of n subsets E_{k_1}, \dots, E_{k_n} such that $\cup_{j=1}^n E_{k_j} = U$.*

We reduce from 3-SET COVER. Given an instance of 3-SET COVER, we construct a security game with one location for each subset E_i , and $\alpha \cdot \beta$ locations for each element $e \in U$, where $\alpha = Z^5$ and $\beta = Z^2 \cdot n$. Hence, there are totally $Z + \alpha \cdot \beta \cdot 3n$ locations. We denote by ℓ_{E_i} the location corresponding to subset E_i and $\ell_{e,\alpha',\beta'}$ with $1 \leq \alpha' \leq \alpha$ and $1 \leq \beta' \leq \beta$ one of $\alpha \cdot \beta$ locations corresponding to element $e \in U$. Moreover, let the number of treasures $|M| = 1 + \beta \cdot (3n - 3)$ and the number of defensive resources $|R| = n$. Since all of treasures, locations, and defensive resources are homogeneous, by normalizing the *importance* function to $v(\ell, m) = 1$ for all $\ell \in L$ and $m \in M$, a state of nature can be directly represented by a function θ that maps a location to a binary such that $\theta(\ell) = 1$ if and only if there is a treasure allocated at location ℓ . In such a security game, the attacker's utility is 1 if and only if the attacker attacks a location ℓ such that $\theta(\ell) = 1$ and ℓ is not defended (while the importance function v and the defense-quality function q can be omitted).

We construct $\alpha \cdot Z$ states of nature in total such that there are α states of nature for each subset E_i , each of which is denoted by $\theta_{E_i,\alpha'}$ for $1 \leq \alpha' \leq \alpha$. Note that in our construction, a requirement is that: for each state $\theta_{E_i,\alpha'}$, there are exactly $|M| = 1 + \beta \cdot (3n - 3)$ locations ℓ with $\theta_{E_i,\alpha'}(\ell) = 1$. The state $\theta_{E_i,\alpha'}$ is defined as follows:

- $\theta_{E_i,\alpha'}(\ell_{E_i}) = 1$ and $\theta_{E_i,\alpha'}(\ell_{E_j}) = 0$ for $j \neq i$.
- $\theta_{E_i,\alpha'}(\ell_{e,\alpha',\beta'}) = 1$ for all $1 \leq \beta' \leq \beta$ if $e \notin E_i$;
- $\theta_{E_i,\alpha'}(\ell_{e,\alpha'',\beta'}) = 0$ for all $1 \leq \beta' \leq \beta$ if $e \in E_i$ or $\alpha'' \neq \alpha'$.

Finally, we let the prior Π be the uniform distribution over all the states of nature.

Lemma 10. *If there exists a solution to the instance of 3-SET COVER, then the game value of the constructed public signaling game is at least $-\frac{1}{\alpha} \cdot \left(1 - \frac{1}{Z}\right)$.*

Proof. Let E_{k_1}, \dots, E_{k_n} be such a 3-set cover and let $C = \{k_1, \dots, k_n\}$. It suffices to exhibit a public signaling scheme which induces an expected utility of the defender at least $-\frac{1}{\alpha} \cdot \left(1 - \frac{1}{Z}\right)$. The following is such a scheme φ :

- Deterministically announce a signal \perp for states $\theta_{E_i, \alpha'}$ for all $1 \leq \alpha' \leq \alpha$ if $i \in C$;
- Deterministically announce a signal i for states $\theta_{E_i, \alpha'}$ for all $1 \leq \alpha' \leq \alpha$ if $i \notin C$;

In other words, the signaling scheme groups together all the states related to the subsets in the set cover, and separately signal the states related to other subsets.

Consider the signal \perp , and the corresponding posterior is that $\Pi_{\varphi, \perp}(\theta_{E_i, \alpha'}) = \frac{1}{n \cdot \alpha}$ for $1 \leq \alpha' \leq \alpha$ if $i \in C$; otherwise, $\Pi_{\varphi, \perp}(\theta_{E_i, \alpha'}) = 0$ for $1 \leq \alpha' \leq \alpha$. As a result, the attacker's expected utility of attacking an undefended location ℓ is $\frac{1}{n}$ for $\ell \in \{\ell_{E_i} \mid i \in C\}$ and is 0 for $\ell \in \{\ell_{E_i} \mid i \notin C\}$. As for the remaining locations $\ell_{e, \alpha', \beta'}$, the attacker's expected utility of attacking any of them if undefended is at most $\frac{1}{\alpha} \cdot (1 - \frac{1}{n})$. This is because (1) only a state $\theta_{E_i, \alpha''}$ with $\alpha'' = \alpha'$ allocates a 1, and therefore, at most n states allocates a 1; (2) since $\{E_i\}_{i \in C}$ forms a 3-set cover, there must exist a subset E_i with $i \in C$ such that $e \in E_i$, and therefore, the corresponding state $\theta_{E_i, \alpha'}$ does not allocate a 1. As a result, there are at most $(n - 1)$ states allocating a 1 among $n \cdot \alpha$ states. Therefore, under the signal \perp , the defender can allocate the n defensive resource to the n locations from $\{\ell_{E_i} \mid i \in C\}$, which limits the attacker's utility to be at most $\frac{1}{\alpha} \cdot (1 - \frac{1}{n})$.

For any signal $i \notin C$, a similar calculation can show that the attacker's expected utility of attacking an undefended location ℓ is 1 for $\ell = \ell_{E_i}$ and is 0 for $\ell \in \{\ell_{E_j} \mid j \neq i\}$. As for the remaining locations $\ell_{e, \alpha', \beta'}$, the attacker's expected utility of attacking any of them if undefended is at most $\frac{1}{\alpha}$. Therefore, the defender can defend the sole location ℓ_{E_i} only, leaving the attacker with locations whose importance is at most $\frac{1}{\alpha}$. To sum up, the attacker's expected utility conditioned on signal \perp is at most $\frac{1}{\alpha} \cdot (1 - \frac{1}{n})$ while conditioned on any signal other than \perp is at most $\frac{1}{\alpha}$. Notice that the probability of \perp is n/Z , and thus, the overall expected utility of the attacker is at most $\frac{n}{Z} \cdot \frac{1}{\alpha} \cdot (1 - \frac{1}{n}) + (1 - \frac{n}{Z}) \cdot \frac{1}{\alpha} = \frac{1}{\alpha} \cdot (1 - \frac{1}{Z})$, which concludes the proof.

Lemma 11. *When $\alpha = Z^5$ and $\beta = Z^2 \cdot n$, if there does not exist a solution to the instance of 3-SET COVER, then the game value of the constructed public signaling game is at most $-\frac{1}{\alpha} \cdot (1 - \frac{2}{Z^2})$.*

Proof. We proceed to prove this lemma by showing that, if there does not exist a solution to the instance of 3-SET COVER, then for any posterior Π , the game value of the induced security game is at most $-\frac{1}{\alpha} \cdot (1 - \frac{2}{Z^2})$. We first claim that for any posterior Π , there are at most n subsets E_i such that $\sum_{\alpha'=1}^{\alpha} \Pi(\theta_{E_i, \alpha'}) > \frac{1}{Z^3}$. To prove this claim, notice that if a subset E_i satisfies $\sum_{\alpha'=1}^{\alpha} \Pi(\theta_{E_i, \alpha'}) > \frac{1}{Z^3}$, it implies the attacker's expected utility of attacking an undefended location ℓ_{E_i} is more than $\frac{1}{Z^3}$. Therefore, to limit the attacker's utility to be at most $\frac{1}{\alpha} \cdot (1 - \frac{2}{Z^2})$, the defender must allocate a defensive resource to ℓ_{E_i} with probability d' such that $\frac{1}{Z^3} \cdot (1 - d') < \frac{1}{\alpha} \cdot (1 - \frac{2}{Z^2})$, which indicates that $d' > 1 - \frac{1}{Z}$ when $\alpha = Z^5$. Hence, for more than n subsets E_i satisfying $\sum_{\alpha'=1}^{\alpha} \Pi(\theta_{E_i, \alpha'}) > \frac{1}{Z^3}$, the defender needs at least $(n+1) \cdot (1 - \frac{1}{Z}) > n$ defensive resources, which produces a contradiction.

For convenience, let $C = \{i \mid \sum_{\alpha'=1}^{\alpha} \Pi(\theta_{E_i, \alpha'}) > \frac{1}{Z^3}\}$ and we have $|C| \leq n$. Furthermore, we have

$$\sum_{i \in C} \sum_{\alpha'=1}^{\alpha} \Pi(\theta_{E_i, \alpha'}) = 1 - \sum_{i \notin C} \sum_{\alpha'=1}^{\alpha} \Pi(\theta_{E_i, \alpha'}) \geq 1 - \sum_{i \notin C} \frac{1}{Z^3} \geq 1 - \frac{1}{Z^2}.$$

Hence, there exists a α^* such that $\sum_{i \in C} \Pi(\theta_{E_i, \alpha^*}) \geq \frac{1}{\alpha} \cdot (1 - \frac{1}{Z^2})$. Since there does not exist a 3-set cover, there exists an element e^* that is not covered by $\cup_{i \in C} E_i$. As a result, for any β' with $1 \leq \beta' \leq \beta$, the attacker's expected utility of attacking an undefended location $\ell_{e^*, \alpha^*, \beta'}$ is $\sum_{i \in C} \Pi(\theta_{E_i, \alpha^*}) \geq \frac{1}{\alpha} \cdot (1 - \frac{1}{Z^2})$. To reduce the attacker's expected utility among these β locations, the best possible strategy of the defender is to evenly distribute n defensive resources among these locations such that each location is defended with probability $\frac{n}{\beta}$. However, after doing so, the attacker's expected utility is still at least $\frac{1}{\alpha} \cdot (1 - \frac{1}{Z^2}) \cdot (1 - \frac{n}{\beta}) = \frac{1}{\alpha} \cdot (1 - \frac{1}{Z^2}) \cdot (1 - \frac{1}{Z^2}) \geq \frac{1}{\alpha} \cdot (1 - \frac{2}{Z^2})$.

Combining Lemma 10, Lemma 11, and Theorem 1, we finish the proof of Theorem 5.

D Omitted Materials in Section 5

D.1 Proof of Lemma 6

For convenience, let the game value be ν . Recall the optimal strategy of the attacker (i.e. the follower in Section 3) is an adaptive strategy that approaches $\mathcal{R}(\nu \cdot \vec{1})$, where $\vec{1}$ is a $|\Theta|$ -dimensional vector. For our class of security games, $|\Theta|$ is the number of possible matchings between L and M , which is exponential in the number of treasures $|M|$ and the number of locations $|L|$.

As described in Section 3, it suffices for the attacker to find a prior Π' such that the halfspace $\mathcal{H}(\vec{\Pi}', \nu)$ separates $\frac{1}{t-1} \sum_{\tau=1}^{t-1} \vec{U}(D^\tau, a^\tau)$ and $\mathcal{R}(\nu \cdot \vec{1})$ at stage t . It is equivalent to find a prior Π' such that

$$\sum_{\theta \in \Theta} \Pi'(\theta) \cdot \left(\frac{1}{t-1} \cdot \sum_{\tau=1}^{t-1} \mathcal{U}^\theta(D^\tau, a^\tau) \right) > \nu,$$

while the left-hand side can be written as

$$\sum_{\theta \in \Theta} \Pi'(\theta) \cdot \left(-\frac{1}{t-1} \cdot \sum_{\tau=1}^{t-1} (1 - q(a^\tau, \theta(a^\tau), D^\tau(a^\tau))) \cdot v(a^\tau, \theta(a^\tau)) \right)$$

which equals to

$$\sum_{\theta \in \Theta} \Pi'(\theta) \cdot \sum_{\ell \in L} v(\ell, \theta(\ell)) \cdot \left(-\frac{1}{t-1} \cdot \sum_{\tau: a^\tau = \ell, 1 \leq \tau \leq t-1} (1 - q(\ell, \theta(\ell), D^\tau(\ell))) \right).$$

For convenience, let

$$\bar{v}(\ell, m) = v(\ell, m) \cdot \left(-\frac{1}{t-1} \cdot \sum_{\tau: a^\tau = \ell, 1 \leq \tau \leq t-1} (1 - q(\ell, m, D^\tau(\ell))) \right),$$

and moreover, let

$$g_{\ell, m}(II') = \sum_{\theta: \theta(\ell) = m} II'(\theta)$$

be the marginals of II' . As a result, we have

$$\sum_{\theta \in \Theta} II'(\theta) \cdot \left(\frac{1}{t-1} \cdot \sum_{\tau=1}^{t-1} \mathcal{U}^\theta(D^\tau, a^\tau) \right) = \sum_{\ell, m} g_{\ell, m}(II') \cdot \bar{v}(\ell, m).$$

Therefore, to find such a prior II' , the attacker can solve the a linear program to maximize δ subject to $\sum_{\ell, m} g_{\ell, m} \cdot \bar{v}(\ell, m) \geq \nu + \delta$ and $\vec{g} \in \mathcal{G}$. Finally, the attacker deploys the minimax strategy of $\mathcal{U}^{II'}$ that can be computed efficiently according to Proposition 3.

D.2 Warm-up: Homogeneous Everything

We provide a warm-up for the simplest case in a Bayesian allocation game with a security game as the stage game, where all of treasures, locations, and defensive resources are homogeneous.

Note that when all of treasures, locations, and defensive resources are homogeneous, it suffices to describe the Bayesian allocation game with three integers: the number of locations $|L|$, the number of treasures $|M|$, and the number of defensive resources $|R|$. Moreover, the states of nature can be simply represented by a function θ that maps a location to a binary such that $\theta(\ell) = 1$ if and only if there is a treasure allocated at location ℓ . For convenience, given a posterior II , let the *importance* function $\text{val}^{II}(\ell) = \sum_{\theta \in \Theta} II(\theta) \cdot \theta(\ell)$ be the attacker's expected utility of attacking ℓ if ℓ is not defended. In addition, note that the mixed strategy of the defender can be summarized by a vector $\vec{d} \in [0, 1]^L$ with $\sum_{\ell \in L} d_\ell \leq |R|$, where d_ℓ is the probability that ℓ is defended.

According to Theorem 4 and Lemma 4, the defender's optimal strategy is to identify a prior II^* and then play the maximin strategy of \mathcal{U}^{II^*} repeatedly. For convenience, let \bar{v}_{alloc} be the *negative* of the game value of the Bayesian allocation game. We will make use of Lemma 13, which provides a characterization of the game value of \mathcal{U}^{II} for a given prior II . We prove Lemma 13 in the main body for the more general setting in which only the treasures are heterogeneous.

The next lemma provides characterization for the vector of the attacker's utilities val^{II^*} . It is worth noticing that for any val' with $\text{val}'(\ell) \in [0, 1]$ for all ℓ and $\sum_{\ell \in L} \text{val}'(\ell) = |R|$, there exists a prior II such that $\text{val}^{II} = \text{val}'$.

Lemma 12. val^{II^*} exhibits the following structure: for at least $|L| - 1$ locations ℓ , either $\text{val}^{II^*}(\ell) = \bar{v}_{\text{alloc}}$ or $\text{val}^{II^*}(\ell) = 1$. Moreover, the remaining location, if any, satisfies $\text{val}^{II^*}(\ell) \in (\bar{v}_{\text{alloc}}, 1)$.

Proof. We first show that for all $\ell \in L$, $\text{val}^{H^*}(\ell) \geq \bar{v}_{\text{allloc}}$. For the sake of contradiction, assume that there exists a location ℓ such that $\text{val}^{H^*}(\ell) < \bar{v}_{\text{allloc}}$. Notice that we can move a very small positive mass from every other location ℓ' to ℓ to obtain val' , and it is clear that $\text{Def}(\bar{v}_{\text{allloc}}, \text{val}^{H^*}) > \text{Def}(\bar{v}_{\text{allloc}}, \text{val}')$, which contradicts Lemma 13.

To complete the proof, assume that there exists two locations ℓ' and ℓ'' such that $\bar{v}_{\text{allloc}} < \text{val}^{H^*}(\ell') \leq \text{val}^{H^*}(\ell'') < 1$. Consider a sufficiently small $\delta > 0$ such that $\bar{v}_{\text{allloc}} < \text{val}^{H^*}(\ell') - \delta < \text{val}^{H^*}(\ell'') + \delta < 1$ and we construct val' by moving a mass of δ from location ℓ' to location ℓ'' . Hence, we have

$$\begin{aligned} & \text{Def}(\bar{v}_{\text{allloc}}, \text{val}^{H^*}) - \text{Def}(\bar{v}_{\text{allloc}}, \text{val}') \\ &= \bar{v}_{\text{allloc}} \cdot \left(\frac{1}{\text{val}^{H^*}(\ell') - \delta} + \frac{1}{\text{val}^{H^*}(\ell'') + \delta} - \frac{1}{\text{val}^{H^*}(\ell')} - \frac{1}{\text{val}^{H^*}(\ell'')} \right) > 0, \end{aligned}$$

which contradicts to Lemma 13.

Intuitively, Lemma 12 suggests that the optimal strategy of the defender is to always allocate some treasures to certain locations so that the attacker knows for sure that these locations have treasures. As for the other treasures, the defender should distribute them almost randomly to the remaining locations. See Figure 3 for an illustration.




Fig. 3. An illustration of the optimal scheme for homogeneous everything

To efficiently compute val^{H^*} , we can first enumerate the number of locations K in which $\text{val}^{H^*}(\ell) = 1$, and therefore there are $(|L| - K - 1)$ locations with $\text{val}^{H^*}(\ell) = \bar{v}_{\text{allloc}}$. For the remaining location, we must have $\text{val}^{H^*}(\ell) = |M| - K - (|L| - K - 1) \cdot \bar{v}_{\text{allloc}}$ to allocate $|M|$ units of treasures in total. As a result, according to Lemma 13, the following equation must be satisfied:

$$K \cdot (1 - \bar{v}_{\text{allloc}}) + \left(1 - \frac{\bar{v}_{\text{allloc}}}{|M| - K - (|L| - K - 1) \cdot \bar{v}_{\text{allloc}}} \right) = |R| \quad (7)$$

which is a quadratic equation for \bar{v}_{allloc} that can be solved efficiently.

Example 2. Consider a security game with 3 locations, 2 treasures, and 1 defensive resource. The following is the allocation strategy of the defender:

- Choose one of the locations and allocate a treasure to it.
- For the remaining two locations, allocate the remaining treasure to one location with probability $\alpha = \frac{\sqrt{5}-1}{2} \approx 0.618$ and allocate it to the other location with probability $1 - \alpha \approx 0.382$.

Note that this leads to an importance function of the form $(1 - \alpha, \alpha, 1)$. Then in each stage of the repeated game, the defender defends the location with importance 1 with probability α (so that the attacker’s utility of attacking this location is $1 - \alpha$), and defends the location with importance α with probability $1 - \alpha$ (so that the attacker’s utility of attacking this location is $\alpha^2 = 1 - \alpha$). The defender never defends the location with important $1 - \alpha$.

Example 3. Consider a security game with a large number $|L|$ of locations, $|L|/2$ treasures, and $|L|/4$ defensive resources. As $|L| \rightarrow \infty$, the following scheme of the defender approaches optimality:

- Choose $|L|/3$ locations and allocate $|L|/3$ treasures to these locations.
- For the remaining $2|L|/3$ locations, allocate a treasure to each location with probability $\frac{1}{4}$.

Note that this leads to an importance function in which $|L|/3$ of the locations have expected importance 1, and the remaining locations have expected importance $\frac{1}{4}$. In each stage, the defender defends locations with importance 1 probability $\frac{3}{4}$ each, and never defends locations with importance $\frac{1}{4}$.

D.3 Only Treasures are Heterogeneous

We focus on the setting in which only the treasures are heterogeneous.

When only the treasures are heterogeneous, the *importance* function v can be simplified as $v(m)$ representing the utility loss of the defender if a location with treasure $m \in M$ allocated is attacked without defense. For convenience, let $M = \{1, \dots, |M|\}$ and $L = \{1, \dots, |L|\}$, and we assume $v(m) \leq v(m + 1)$ for all m . Given a prior Π , let its marginals be $g_{\ell,m}(\Pi) = \sum_{\theta:\theta(\ell)=m} \Pi(\theta)$, and moreover, let $\text{val}^\Pi(\ell) = \sum_{m \in M} g_{\ell,m}(\Pi) \cdot v(m)$ be the attacker’s utility of attacking ℓ if ℓ is undefended. Since the defensive resources are homogeneous, the mixed strategy of the defender can still be summarized by a vector $\vec{d} \in [0, 1]^L$ with $\sum_{\ell \in L} d_\ell \leq |R|$.

According to Theorem 4 and Lemma 4, the defender’s optimal strategy is to identify a prior Π^* and then play the maximin strategy of \mathcal{U}^{Π^*} repeatedly. For convenience, let $\bar{\nu}_{\text{alloc}}$ be the *negative* of the game value ν_{alloc} of the Bayesian allocation game. The following lemma provides a characterization of the game value of \mathcal{U}^Π for a given prior Π . Let $(x)^+ = \max(x, 0)$.

Lemma 13. *For a prior Π , $\text{Def}(\lambda, \text{val}^\Pi) = \sum_{\ell \in L} \left(1 - \frac{\lambda}{\text{val}^\Pi(\ell)}\right)^+ \leq |R|$ if and only if the game value of \mathcal{U}^Π is at least $-\lambda$. Moreover, $\text{Def}(\bar{\nu}_{\text{alloc}}, \text{val}^{\Pi^*}) = \min_\Pi \text{Def}(\bar{\nu}_{\text{alloc}}, \text{val}^\Pi) = |R|$.*

Proof. Notice that given a defender's strategy \vec{d} , the optimal strategy of the attacker is to attack a location from $\operatorname{argmax}_{\ell \in L} (1 - d_\ell) \cdot \operatorname{val}^\Pi(\ell)$. As a result, to achieve utility $-\lambda$ for the defender, for a location ℓ with $\operatorname{val}^\Pi(\ell) \leq \lambda$, there is no need to defend it. As for the location ℓ with $\operatorname{val}^\Pi(\ell) > \lambda$, the defender must defend it with probability at least $1 - \frac{\lambda}{\operatorname{val}^\Pi(\ell)}$. Therefore, we can conclude that the game value of \mathcal{U}^Π is at least $-\lambda$ if and only if $\operatorname{Def}(\lambda, \operatorname{val}^\Pi) \leq |R|$. Moreover, for the prior Π^* in the defender's optimal strategy, $\operatorname{Def}(\bar{\nu}_{\text{alloc}}, \operatorname{val}^{\Pi^*}) = |R|$ because the game value of \mathcal{U}^{Π^*} is exactly ν_{alloc} . Finally, note that the game value of \mathcal{U}^Π for any Π is at most ν_{alloc} , which implies that for any Π , $\operatorname{Def}(\bar{\nu}_{\text{alloc}}, \operatorname{val}^\Pi) \geq |R|$. Thus, $\operatorname{Def}(\bar{\nu}_{\text{alloc}}, \operatorname{val}^{\Pi^*}) = \min_\Pi \operatorname{Def}(\bar{\nu}_{\text{alloc}}, \operatorname{val}^\Pi)$.

In order to design our algorithm, we provide structural results for the optimal prior Π^* , particularly for $\vec{g}(\Pi^*)$ and $\operatorname{val}^{\Pi^*}$. Without loss of generality, assume $\operatorname{val}^{\Pi^*}(\ell) \leq \operatorname{val}^{\Pi^*}(\ell+1)$ for all ℓ . We partition the locations into two sets: **High** = $\{\ell \in L \mid \operatorname{val}^{\Pi^*}(\ell) > \bar{\nu}_{\text{alloc}}\}$ and **Low** = $\{\ell \in L \mid \operatorname{val}^{\Pi^*}(\ell) \leq \bar{\nu}_{\text{alloc}}\}$. Notice that the defender needs to allocate defensive resource to a location ℓ if and only if $\operatorname{val}^{\Pi^*}(\ell) > \bar{\nu}_{\text{alloc}}$, i.e., $d_\ell > 0$ if and only if $\ell \in \text{High}$.

We start with showing that there exists an optimal prior Π^* such that the treasures are allocated to locations in **High** in a decreasing order of their importance.

Lemma 14. *There exists an optimal prior Π^* such that for $\ell \in \text{High}$, it satisfies that if $g_{\ell,m}(\Pi^*) > 0$, then $g_{\ell',m'}(\Pi^*) = 0$ for all $\ell' < \ell$ and $m' > m$.*

Proof. For an optimal prior Π^* , if there exists $\ell' < \ell$ and $m < m'$ such that $g_{\ell',m'}(\Pi^*) > 0$ and $g_{\ell,m}(\Pi^*) > 0$, consider the following $\vec{g}(\Pi')$ that is the same as $\vec{g}(\Pi^*)$ except that we swap a sufficiently small $\delta > 0$ of the allocation of treasure m' and m to location ℓ' and ℓ :

$$\begin{aligned} - g_{\ell',m'}(\Pi') &= g_{\ell',m'}(\Pi^*) - \delta \text{ and } g_{\ell',m}(\Pi') = g_{\ell',m}(\Pi^*) + \delta; \\ - g_{\ell,m'}(\Pi') &= g_{\ell,m'}(\Pi^*) + \delta \text{ and } g_{\ell,m}(\Pi') = g_{\ell,m}(\Pi^*) - \delta. \end{aligned}$$

As a result, $\operatorname{val}^{\Pi'}$ is the same as $\operatorname{val}^{\Pi^*}$ except for location ℓ and ℓ' such that

$$\operatorname{val}^{\Pi'}(\ell') = \operatorname{val}^{\Pi^*}(\ell') - \delta \cdot (v(m') - v(m))$$

while

$$\operatorname{val}^{\Pi'}(\ell) = \operatorname{val}^{\Pi^*}(\ell) + \delta \cdot (v(m') - v(m)).$$

Hence,

$$\begin{aligned} & \operatorname{Def}(\bar{\nu}_{\text{alloc}}, \operatorname{val}^{\Pi^*}) - \operatorname{Def}(\bar{\nu}_{\text{alloc}}, \operatorname{val}^{\Pi'}) \\ &= \frac{\bar{\nu}_{\text{alloc}}}{\operatorname{val}^{\Pi^*}(\ell') - \delta \cdot (v(m') - v(m))} + \frac{\bar{\nu}_{\text{alloc}}}{\operatorname{val}^{\Pi^*}(\ell) + \delta \cdot (v(m') - v(m))} \\ & \quad - \frac{\bar{\nu}_{\text{alloc}}}{\operatorname{val}^{\Pi^*}(\ell')} - \frac{\bar{\nu}_{\text{alloc}}}{\operatorname{val}^{\Pi^*}(\ell)} \\ & \geq 0, \end{aligned}$$

where the inequality uses the facts that $v(m') - v(m) \geq 0$ and $\text{val}^{\Pi^*}(\ell') \leq \text{val}^{\Pi^*}(\ell)$. According to Lemma 13, either Π^* is not an optimal prior (when the above formula is a strict inequality) or Π' is also an optimal prior (when the above formula is an equation).

We now focus on $\text{Low} = \left\{ \ell \in L \mid \text{val}^{\Pi^*}(\ell) \leq \bar{v}_{\text{alloc}} \right\}$. Note that we must have

$$\sum_{\ell \in \text{Low}} \sum_{m \in M} g_{\ell, m}(\Pi^*) \cdot v(m) \leq |\text{Low}| \cdot \bar{v}_{\text{alloc}}, \sum_{\ell \in \text{Low}} \sum_{m \in M} g_{\ell, m}(\Pi^*) \leq |\text{Low}|; \quad (8)$$

or otherwise, either there exists a location $\ell \in \text{Low}$ in which $\text{val}^{\Pi^*}(\ell) > \bar{v}_{\text{alloc}}$ or there does not exist enough space to fulfill these treasures. On the other hand, when (8) is satisfied, the defender can simply allocate these treasures uniformly to the locations in Low : concretely, consider $\bar{g}(\Pi')$ that is the same as $\bar{g}(\Pi^*)$ except that $g_{\ell, m}(\Pi') = \sum_{\ell' \in \text{Low}} \sum_{m \in M} g_{\ell', m}(\Pi^*) / |\text{Low}|$ for all $\ell \in \text{Low}$ and $m \in M$. Hence, we denote by $g_{\text{Low}, m}(\Pi^*) = \sum_{\ell \in \text{Low}} g_{\ell, m}(\Pi^*)$ for simplicity. In addition, note that at least one of the inequalities in (8) must be tight; or otherwise, the defender can allocate more treasures to Low without suffering additional utility loss.

Proposition 5. *For the optimal prior Π^* , it satisfies,*

$$\sum_{m \in M} g_{\text{Low}, m}(\Pi^*) \cdot v(m) \leq |\text{Low}| \cdot \bar{v}_{\text{alloc}}$$

and

$$\sum_{m \in M} g_{\text{Low}, m}(\Pi^*) \leq |\text{Low}|.$$

Moreover, at least one of these two inequalities is tight.

The next lemma provides a finer characterization of $g_{\text{Low}, m}(\Pi^*)$.

Lemma 15. *There exists an optimal prior Π^* such that, if $g_{\text{Low}, m}(\Pi^*) < 1$, then for any pair (m', m'') with $m' < m$ and $m'' > m$, either $g_{\text{Low}, m'}(\Pi^*) = 0$ or $g_{\text{Low}, m''}(\Pi^*) = 0$.*

Proof. For an optimal prior Π^* , if there exists $m' < m < m''$ such that they satisfy $g_{\text{Low}, m}(\Pi^*) < 1$, $g_{\text{Low}, m'}(\Pi^*) > 0$, and $g_{\text{Low}, m''}(\Pi^*) > 0$, consider the following $\bar{g}(\Pi')$ that is the same as $\bar{g}(\Pi^*)$ except for treasures m , m' , and m'' in Low and a location $\ell \in \text{High}$ where $g_{\ell, m}(\Pi^*) > 0$:

- $g_{\ell, m}(\Pi') = g_{\ell, m}(\Pi^*) - \delta$ and $g_{\text{Low}, m}(\Pi') = g_{\text{Low}, m}(\Pi^*) + \delta$;
- $g_{\ell, m'}(\Pi') = g_{\ell, m'}(\Pi^*) + c \cdot \delta$ and $g_{\text{Low}, m'}(\Pi') = g_{\text{Low}, m'}(\Pi^*) - c \cdot \delta$;
- $g_{\ell, m''}(\Pi') = g_{\ell, m''}(\Pi^*) + (1-c) \cdot \delta$ and $g_{\text{Low}, m''}(\Pi') = g_{\text{Low}, m''}(\Pi^*) - (1-c) \cdot \delta$;

for a sufficiently small $\delta > 0$ and $0 \leq c \leq 1$ such that $v(m) = c \cdot v(m') + (1-c) \cdot v(m'')$. It is straightforward to verify that $\sum_{m \in M} g_{\text{Low}, m}(\Pi^*) \cdot v(m) = \sum_{m \in M} g_{\text{Low}, m}(\Pi') \cdot v(m)$ and $\sum_{m \in M} g_{\text{Low}, m}(\Pi^*) = \sum_{m \in M} g_{\text{Low}, m}(\Pi')$. Moreover, $\text{val}^{\Pi'}(\ell) = \text{val}^{\Pi^*}(\ell)$ for all $\ell \in \text{High}$, and therefore, $\text{Def}(\bar{v}_{\text{alloc}}, \text{val}^{\Pi'}) = \text{Def}(\bar{v}_{\text{alloc}}, \text{val}^{\Pi^*})$ so that Π' is also an optimal prior according to Lemma 13.

Lemma 15 suggests that for an optimal prior Π^* , there exists two cut-off treasures m^{c1} and m^{c2} such that for any treasure m with $m < m^{c1}$ or $m > m^{c2}$, we have $g_{\text{Low},m}(\Pi^*) = 0$. Moreover, for $m^{c1} < m < m^{c2}$, we have $g_{\text{Low},m}(\Pi^*) = 1$. Combining with Lemma 14, we conclude that $\text{Def}(\bar{v}_{\text{allloc}}, \text{val}^{\Pi^*})$ is fully characterized by $\bar{g}_{\text{Low}}(\Pi^*)$. See Figure 4 for an illustration of the structure of the optimal scheme.




Fig. 4. An illustration of the optimal scheme for only treasures are heterogeneous

To efficiently compute the optimal scheme, the defender first applies binary search for \bar{v}_{allloc} and checks whether there exists a scheme that can achieve $-\bar{v}_{\text{allloc}}$. Given \bar{v}_{allloc} , the defender enumerates $|\text{Low}|$, the size of locations with importance at most \bar{v}_{allloc} , and the (left) cut-off treasure m^{c1} . Note that the value of $g_{\text{Low},m^{c1}}$ uniquely determines $\{g_{\text{Low},m}\}_{m \in \text{Low}}$ that satisfies both Proposition 5 and Lemma 15 (unless there is no feasible solution). Given $\{g_{\text{Low},m}\}_{m \in \text{Low}}$, the defender then allocates the remaining treasures according to Lemma 14 and computes the game value. Therefore, for fixed \bar{v}_{allloc} and m^{c1} , the defender's task is to optimize the game value via $g_{\text{Low},m^{c1}}$, which is a single parameter optimization problem that can be solved efficiently.

D.4 Only Locations are Heterogeneous

When only the locations are heterogeneous, the *treasure-location importance* function such that v can be simplified as $v(\ell)$ representing the utility loss of the defender if location ℓ with a treasure allocated is attacked without defense. Given a prior Π , let $\gamma_\ell(\Pi) = \sum_\theta \Pi(\theta) \cdot \mathbf{1}\{\theta(\ell) \neq \perp\}$ be the expected number of treasures allocated at location ℓ . Note that any $\vec{\gamma} \in [0, 1]^{|L|}$ with $\sum_\ell \gamma_\ell \leq |M|$ can be induced by a prior Π . As a result, the importance function is $\text{val}^\Pi(\ell) = \gamma_\ell(\Pi) \cdot v(\ell)$. Since the defensive resources are homogeneous, the mixed strategy of the defender can still be summarized by a vector $\vec{d} \in [0, 1]^L$ with $\sum_{\ell \in L} d_\ell \leq |R|$. Let \bar{v}_{allloc} be the *negative* of the game value of the Bayesian allocation game and Lemma 13 applies.

Before presenting the hardness reduction, we first provide a structural result of $\vec{\gamma}(\Pi^*)$. Note that for a location ℓ with $v(\ell) < \bar{v}_{\text{allloc}}$, it is always optimal to allocate a unit of treasure to ℓ without defense. Hence, we assume that $v(\ell) \geq \bar{v}_{\text{allloc}}$ for all ℓ for the rest of the discussion.

Lemma 16. $\bar{\gamma}(\Pi^*)$ exhibits the following structure: for at least $|L| - 1$ locations ℓ , either $\gamma_\ell(\Pi^*) = 1$ or $\gamma_\ell(\Pi^*) = \bar{\nu}_{\text{alloc}}/v(\ell)$. Moreover, for the remaining location ℓ , if any, satisfies $\gamma_\ell(\Pi^*) \in (\bar{\nu}_{\text{alloc}}/v(\ell), 1)$.

Proof. First note that for all $\ell \in L$, $\gamma_\ell(\Pi^*) \geq \bar{\nu}_{\text{alloc}}/v(\ell)$; or otherwise, the defender can allocate more treasures to location ℓ without defense up to $\bar{\nu}_{\text{alloc}}/v(\ell)$.

Assume that there exists two locations ℓ' and ℓ'' such that $\bar{\nu}_{\text{alloc}}/v(\ell) < \gamma_\ell(\Pi^*) < 1$. Consider $\bar{\gamma}(\Pi'_\delta)$ that is the same as $\bar{\gamma}(\Pi^*)$ except for location ℓ' and ℓ'' such that $\gamma_{\ell'}(\Pi'_\delta) = \gamma_{\ell'}(\Pi^*) + \delta$ and $\gamma_{\ell''}(\Pi'_\delta) = \gamma_{\ell''}(\Pi^*) - \delta$, for some $\delta \in \Delta$ such that $\Delta = [L, R]$ where

$$L = - \min \left\{ \gamma_{\ell'}(\Pi^*) - \frac{\bar{\nu}_{\text{alloc}}}{v(\ell')}, 1 - \gamma_{\ell''}(\Pi^*) \right\},$$

and

$$R = \min \left\{ \gamma_{\ell''}(\Pi^*) - \frac{\bar{\nu}_{\text{alloc}}}{v(\ell'')}, 1 - \gamma_{\ell'}(\Pi^*) \right\}.$$

Hence, $\text{val}^{\Pi'_\delta}$ is the same as val^{Π^*} except that $\text{val}^{\Pi'_\delta}(\ell') = \text{val}^{\Pi^*}(\ell') + \delta \cdot v(\ell')$ and $\text{val}^{\Pi'_\delta}(\ell'') = \text{val}^{\Pi^*}(\ell'') - \delta \cdot v(\ell'')$. Then, we have

$$\begin{aligned} & h(\delta) \\ &= \text{Def}(\bar{\nu}_{\text{alloc}}, \text{val}^{\Pi'_\delta}) - \text{Def}(\bar{\nu}_{\text{alloc}}, \text{val}^{\Pi^*}) \\ &= \bar{\nu}_{\text{alloc}} \cdot \left(\frac{1}{\text{val}^{\Pi^*}(\ell')} + \frac{1}{\text{val}^{\Pi^*}(\ell'')} - \frac{1}{\text{val}^{\Pi^*}(\ell') + \delta \cdot v(\ell')} - \frac{1}{\text{val}^{\Pi^*}(\ell'') - \delta \cdot v(\ell'')} \right). \end{aligned}$$

Taking the second order derivative of $h(\delta)$ with respect to δ , we have

$$h''(\delta) = -\bar{\nu}_{\text{alloc}} \cdot \left(\frac{2 \cdot (v(\ell'))^2}{(\text{val}^{\Pi^*}(\ell') + \delta \cdot v(\ell'))^3} + \frac{2 \cdot (v(\ell''))^2}{(\text{val}^{\Pi^*}(\ell'') - \delta \cdot v(\ell''))^3} \right) \leq 0,$$

which implies that $h(\delta)$ is a concave function with its minimum obtained at the boundary of Δ . Let $\delta^* = \text{argmin}_{\delta \in \Delta} h(\delta)$ and since δ^* is at the boundary of Δ , we have either $\gamma_{\ell'}(\Pi'_{\delta^*}) \in \{\bar{\nu}_{\text{alloc}}/v(\ell'), 1\}$ or $\gamma_{\ell''}(\Pi'_{\delta^*}) \in \{\bar{\nu}_{\text{alloc}}/v(\ell''), 1\}$. Finally, according to Lemma 13, either Π^* is not an optimal prior (when $\min_{\delta \in \Delta} h(\delta) < 0$) or Π'_{δ^*} is also an optimal prior (when $\min_{\delta \in \Delta} h(\delta) = 0$).

As suggested by Lemma 16: in the optimal scheme, the defender should either allocate a unit of treasure to location ℓ and defend it with probability $1 - \bar{\nu}_{\text{alloc}}/v(\ell)$, or allocate $\bar{\nu}_{\text{alloc}}/v(\ell)$ treasure without defense, for at least $|L| - 1$ locations.

Hardness Reduction We reduce from SUBSET SUM.

Definition 5 (Subset Sum). In an instance of SUBSET SUM, we are given a set of n numbers $\{w_1, \dots, w_n\}$ with $0 < w_i < 1$ for all $i \in [n]$ and a target W . The task is to determine whether there exists a subset $C \subseteq [n]$ such that $\sum_{i \in C} w_i = W$. Without loss of generality, we assume both $\sum_{i \in [n]} w_i$ and W are integral.

Given an instance of SUBSET SUM, we construct a security game with $|L| = n$ locations with $L = [n]$, $|M| = n + W - \sum_{i \in [n]} w_i$ treasures, and $|R| = W$ defensive resources. Finally, for location $i \in [n]$, $v(i) = 1/(1-w_i)$. The team max-min game is directly induced from this security game. We finish the reduction by showing that there exists a solution to the instance of SUBSET SUM if and only if the game value of the constructed team max-min game is at least -1 .

Lemma 17. *If there exists a solution to the instance of SUBSET SUM, then the game value of the constructed team max-min game is at least -1 .*

Proof. Let C be the solution of the instance of SUBSET SUM. Consider the following $\tilde{\gamma}(II)$:

- For location $i \in C$, $\gamma_i(II) = 1$;
- For location $i \notin C$, $\gamma_i(II) = 1/v(i)$.

Note that the total units of allocated treasure is

$$\begin{aligned} \sum_{i \in [n]} \gamma_i(II) &= |C| + \sum_{i \notin C} \frac{1}{v(i)} = |C| + \sum_{i \notin C} (1 - w_i) \\ &= n + \sum_{i \in C} w_i - \sum_{i \in [n]} w_i = n + W - \sum_{i \in [n]} w_i = |M|. \end{aligned}$$

Moreover, we have $\text{val}^{II}(i) = v(i)$ for $i \in C$ and $\text{val}^{II}(i) = 1$ for $i \notin C$.

$$\text{Def}(1, \text{val}^{II}) = \sum_{i \in [n]} \left(1 - \frac{1}{\text{val}^{II}(i)}\right) = \sum_{i \in C} \left(1 - \frac{1}{v(i)}\right) = \sum_{i \in C} w_i = W = |R|.$$

According to Lemma 13, we can conclude that the game value is at least -1 .

Lemma 18. *If the game value of the constructed team max-min game is at least -1 , then there exists a solution to the instance of SUBSET SUM.*

Proof. Suppose the game value is $-\kappa$ with $0 \leq \kappa \leq 1$. Then according to Lemma 16, we have that for $n - 1$ locations, we have either $\gamma_i(II^*) = 1$ or $\gamma_i(II^*) = \kappa/v(i) = \kappa \cdot (1 - w_i)$. Denote the remaining location by i^* . For convenience, let $A = \{i \neq i^* \mid \gamma_i(II^*) = 1\}$ and $B = \{i \neq i^* \mid \gamma_i(II^*) = \kappa \cdot (1 - w_i)\}$. Then, for the remaining location i^* , we have

$$\begin{aligned} \gamma_{i^*}(II^*) &= n + W - \sum_{i \in [n]} w_i - |A| - \sum_{i \in B} \kappa \cdot (1 - w_i) \\ &= W - \sum_{i \in A} w_i + (1 - w_{i^*}) + \sum_{i \in B} (1 - \kappa) \cdot (1 - w_i) \end{aligned}$$

and $\kappa/v(i^*) \leq \gamma_{i^*}(\Pi^*) \leq 1$. Moreover, we have that

$$\begin{aligned} \text{Def}(\kappa, \text{val}^{\Pi^*}) &= \sum_{i \in [n]} \left(1 - \frac{\kappa}{\gamma_i(\Pi^*) \cdot v(i)} \right) \\ &= \sum_{i \in A} (1 - \kappa \cdot (1 - w_i)) + \left(1 - \frac{\kappa \cdot (1 - w_{i^*})}{\gamma_{i^*}(\Pi^*)} \right) \\ &\geq \sum_{i \in A} w_i + \left(1 - \frac{\kappa \cdot (1 - w_{i^*})}{\gamma_{i^*}(\Pi^*)} \right), \end{aligned}$$

where the inequality uses the fact that $0 \leq \kappa \leq 1$. Plugging in the expression of $\gamma_{i^*}(\Pi^*)$, we have

$$\begin{aligned} &\text{Def}(\kappa, \text{val}^{\Pi^*}) \\ &\geq \sum_{i \in A} w_i + \left(1 - \frac{\kappa \cdot (1 - w_{i^*})}{W - \sum_{i \in A} w_i + (1 - w_{i^*}) + \sum_{i \in B} (1 - \kappa) \cdot (1 - w_i)} \right) \\ &= \sum_{i \in A} w_i + \frac{W - \sum_{i \in A} w_i + (1 - \kappa) \cdot (1 - w_{i^*}) + \sum_{i \in B} (1 - \kappa) \cdot (1 - w_i)}{W - \sum_{i \in A} w_i + (1 - w_{i^*}) + \sum_{i \in B} (1 - \kappa) \cdot (1 - w_i)} \\ &\geq W + (1 - \kappa) \cdot (1 - w_{i^*}) + \sum_{i \in B} (1 - \kappa) \cdot (1 - w_i) \\ &\geq W, \end{aligned}$$

where the second inequality uses the fact that $\gamma_{i^*}(\Pi^*) \leq 1$ and the last inequality uses the fact that $0 \leq \kappa \leq 1$. However, since there are $|R| = W$ defensive resources, both inequalities must be equalities, which implies that $\gamma_{i^*}(\Pi^*) = 1$ and $\kappa = 1$. Recall the formula of $\gamma_{i^*}(\Pi^*)$, and when $\gamma_{i^*}(\Pi^*) = 1$ and $\kappa = 1$, we have

$$\gamma_{i^*}(\Pi^*) = 1 = W - \sum_{i \in A} w_i + (1 - w_{i^*})$$

which implies that $W = w_{i^*} + \sum_{i \in A} w_i$. Thus, $\{i^*\} \cup A$ is the solution to the SUBSET SUM problem.

Combining Lemma 17, Lemma 18, and Theorem 4, we finish the proof of the hardness result in Theorem 7. We next present our pseudo-polynomial time algorithm.

Pseudo-polynomial Time Algorithm Our algorithm is based on the structural results of $\bar{\gamma}(\Pi^*)$ in Lemma 16. The defender first applies a binary search on the negative of the game value \bar{v}_{alloc} to check whether the game value is achievable. The defender then enumerates a location i^* and without loss of generality, we label the remaining $n - 1$ locations as $\{1, \dots, n - 1\}$.

We consider a dynamic program that solves $\text{OPT}(j, m)$, which indicates the minimum units of defensive resources needed to defend locations $\{1, \dots, j\}$ when

at most m units of treasures have been allocated. Then, the transition function is as follows: $\text{OPT}(j, m)$ is computed as

$$\min \left\{ \text{OPT}(j-1, m-1) + \left(1 - \frac{\bar{v}_{\text{alloc}}}{v(j)}\right), \text{OPT}\left(j-1, m - \frac{\bar{v}_{\text{alloc}}}{v(j)}\right) \right\}$$

since the defender will either allocate 1 unit of treasures at location j and defend it with probability $(1 - \bar{v}_{\text{alloc}}/v(j))$ or allocate $\bar{v}_{\text{alloc}}/v(j)$ units of treasures without defense according to Lemma 16. The boundary conditions are $\text{OPT}(0, 0) = 0$ and $\text{OPT}(0, m) = -\infty$ for $m > 0$. Finally, the defender searches whether there exists $\text{OPT}(n-1, m)$ with $|M| - 1 \leq m \leq |M| - \bar{v}_{\text{alloc}}/v(i^*)$ such that the defender has sufficient defensive resource to defend location i^* allocated with the remaining $|M| - m$ units of treasures, i.e., whether the following inequality holds:

$$\text{OPT}(n-1, m) + \left(1 - \frac{\bar{v}_{\text{alloc}}}{(|M| - m) \cdot v(i^*)}\right) \leq |R|.$$

D.5 Only Defensive Resources are Homogeneous

We will reduce from 3-SET COVER (see Definition 4). For ease of presentation, we assume that both n and Z are even integers. Given an instance of 3-SET COVER, we construct a security game with $|L| = |U| \cdot Z + |U| = 3n \cdot Z + 3n$ locations, $|M| = |U| \cdot Z + n + (Z - n)/2 = 3n \cdot Z + (Z + n)/2$ treasures, and $|R| = n/2$ defensive resources.

In particular, for each pair of $e \in U$ and E_i for $i \in [Z]$, we construct a location denoted by ℓ_{e, E_i} , and for each E_i for $i \in [Z]$, we construct a location denoted by ℓ_{E_i} . Moreover, for each $e \in U$, we construct a treasure denoted by m_e and for each E_i with $i \in [Z]$, we construct a treasure denoted by m_{E_i} . In addition, we add $Q = |M| - |U| - Z = 3n \cdot Z + (Z + n)/2 - 3n - Z$ units of identical treasures, denoted by m_k^* for $k \in [Q]$. The location-treasure importance function v is given as follows:

- $v(m_e, \ell_{e, E_i}) = 1$ if $e \in E_i$; and $v(m_e, \ell_{e', E_i}) = \infty$ if $e \neq e'$ or $e \notin E_i$;
- $v(m_{E_i}, \ell_{e, E_i}) = 6n$ for all $e \in U$; and $v(m_{E_i}, \ell_{e, E_j}) = \infty$ for all $e \in U$ and $j \neq i$;
- $v(m_k^*, \ell_{e, E_j}) = \infty$ for all $k \in [Q]$, $e \in U$, $j \in [Z]$.

Moreover, for the remaining locations:

- $v(m_e, \ell_{E_i}) = \infty$ for all $e \in U$ and $i \in [Z]$;
- $v(m_{E_i}, \ell_{E_i}) = 2$ and $v(m_{E_i}, \ell_{E_j}) = \infty$ for all $j \neq i$ and $j \in [Z]$;
- $v(m_k^*, \ell_{E_i}) = \infty$ for all $k \in [Q]$ and $i \in [Z]$.

Since the defensive resources are homogeneous, the mixed strategy of the defender can still be summarized by a vector $\vec{d} \in [0, 1]^L$ with $\sum_{\ell \in L} d_\ell \leq |R|$. Moreover, recall that we can represent any prior Π with its corresponding fraction matching $\vec{g}(\Pi)$.

The team max-min game is directly induced from this security game. We finish the reduction by showing that there exists a solution to the instance of 3-SET COVER if and only if the game value of the constructed team max-min game is at least -1 .

Lemma 19. *If there exists a solution to the instance of 3-SET COVER, then the game value of the constructed team max-min game is at least -1 .*

Proof. Let E_{k_1}, \dots, E_{k_n} be such a 3-set cover and let $C = \{k_1, \dots, k_n\}$. It suffices to exhibit a scheme which induces an expected utility of the defender at least -1 . Consider the following prior Π and the corresponding strategy of allocating defensive resources for the defender:

- For $i \in C$, $g_{m_{E_i}, \ell_{E_i}} = 1$ and $d_{\ell_{E_i}} = 1/2$;
- For $i \notin C$, $g_{m_{E_i}, \ell_{E_i}} = 1/2$ and $g_{m_{E_i}, \ell_{e, E_i}} = 1/(6n)$ for all $e \in U$;
- For any pair of (e, E_i) with $i \in C$ and $e \in E_i$, $g_{m_e, \ell_{e, E_i}} = 1$.

Note that since E_{k_1}, \dots, E_{k_n} is a 3-set cover, for each e , there exists a unique E_i with $i \in C$ that contains e . Therefore, we have allocated all treasures except m_k^* for $k \in [Q]$ and moreover, we have exhausted the defensive resources. Observe that for locations other than ℓ_{E_i} for $i \in [Z]$, we have allocated totally $(Z - n)/2 + 3n$ units of treasures to $3n \cdot Z$ locations. Hence, we can still allocate $3n \cdot Z - (Z - n)/2 - 3n = Q$ more units to these locations, which means that we can allocate all treasures m_k^* for $k \in [Q]$. Finally, it is straightforward to verify that under such a scheme, the attacker's utility of attacking any location is at most 1.

Lemma 20. *If the game value of the constructed team max-min game is at least -1 , then there exists a solution to the instance of 3-SET COVER.*

Proof. For convenience, let $\bar{L} = \{\ell_{e, E_i}\}_{e \in U, i \in Z}$. We first claim that it is without loss of generality to focus on $\vec{g}(\Pi^*)$ such that $g_{m_k^*, \ell}(\Pi^*) = 0$ for all $k \in [Q]$ and $\ell \notin \bar{L}$. Because if there exists k^* and i^* such that $g_{m_{k^*}^*, \ell^*}(\Pi^*) > 0$ for $\ell^* \notin \bar{L}$, notice that the importance of ℓ^* is ∞ if $m_{k^*}^*$ is allocated, while $v(m_{k^*}^*, \ell) = 0$ for all $\ell \in \bar{L}$. Therefore, moving $m_{k^*}^*$ to \bar{L} if there is still space left or switching $m_{k^*}^*$ with any treasure from $\{m_e\}_{e \in U} \cup \{m_{E_j}\}_{j \in Z}$ that is allocated to \bar{L} would not cause the defender to suffer additional utility loss.

As a result, there are $3n \cdot Z - Q = 3n + (Z - n)/2 = |U| + (Z - n)/2$ space left among \bar{L} . However, notice that we still have $|U| + Z$ treasures from $\{m_e\}_{e \in U} \cup \{m_{E_j}\}_{j \in Z}$ to be allocated. Therefore, there are at least $n + (Z - n)/2$ treasures that must be allocated to locations in $L \setminus \bar{L}$. For convenience, let $\eta(E_i) = g_{m_{E_i}, \ell_{E_i}}(\Pi^*)$.

Proposition 6. *There exists Π^* such that there are exactly n subsets E_i satisfying $\eta(E_i) = 1$ while for any other subset E_j , $\eta(E_j) = 1/2$, which results in*

$$\sum_{i \in [Z]} \eta(E_i) = n + (Z - n)/2.$$

Moreover, for any E_j , we have $g_{m_e, \ell_{E_j}}(\Pi^*) = 0$ for all $e \in U$ and $g_{m_{E_i}, \ell_{E_j}}(\Pi^*) = 0$ for all $i \neq j$.

Proof. First, notice that for a location ℓ_{E_j} with $g_{m_e, \ell_{E_j}}(\Pi^*) > 0$ for some $e \in U$ or $g_{m_{E_i}, \ell_{E_j}}(\Pi^*) > 0$ for some $i \neq j$, its importance is ∞ and the defender must defend it with probability 1. As a result, it is beneficial to allocate a unit of treasure in total to such a location. Let \bar{C} be the set of the indices of these locations.

For location ℓ_{E_i} with $i \notin \bar{C}$, we need to defend it with probability at least $\left(1 - \frac{1}{2 \cdot \eta(E_i)}\right)^+$ to limit the attacker's utility to be at most 1. Hence, it is always optimal to have $\eta(E_i) \geq 1/2$ for $i \notin \bar{C}$. Moreover, since there are $n/2 - |\bar{C}|$ defensive resources available, we have

$$\sum_{i \in [Z] \setminus \bar{C}} \left(1 - \frac{1}{2 \cdot \eta(E_i)}\right)^+ \leq \frac{n}{2} - |\bar{C}|.$$

Similar to the function $h(\delta)$ discussed in the proof of Lemma 16, it can be shown that the left-hand side of the above inequality, subject to a constraint that $\sum_i \eta(E_i) = c$ where c is a constant, is minimized at the boundary where all but one $i \in Z$ satisfy that $\eta(E_i) \in \{1/2, 1\}$. Moreover, notice that it is beneficial to have $\eta(E_i) = 1$. As a result, to maximize $\sum_i \eta(E_i)$, there are exactly $n - 2|\bar{C}|$ subsets E_i satisfying $\eta(E_i) = 1$ while for other subset E_j with $j \notin \bar{C}$, $\eta(E_j) = 1/2$, which results in exactly $\sum_{i \notin \bar{C}} \eta(E_i) = n - 2|\bar{C}| + (Z - n + |\bar{C}|)/2$.

Combining with the $|\bar{C}|$ units of treasures to locations $\{\ell_{E_j}\}_{j \in \bar{C}}$, the total units of treasures is $n + (Z - n)/2 - |\bar{C}|/2 \leq n + (Z - n)/2$. The equality holds only if $|\bar{C}| = 0$. Thus, there are exactly n subsets E_i satisfying $\eta(E_i) = 1$.

Since there are exactly n subsets E_i satisfying $g_{m_{E_i}, \ell_{E_i}} = 1$, the defender must defend each such ℓ_{E_i} with probability $1/2$, and thus, the defensive resources are exhausted. Let $C = \{i \in [Z] \mid \eta(E_i) = 1\}$. We claim that $\cup_{i \in C} E_i = U$. For the sake of contradiction, assume that there exists $e^* \in U$ such that $e^* \notin E_i$ for all $i \in C$. Observe that for $i \notin C$ with $e \in E_i$, the remaining $1/2$ units of treasure m_{E_i} can only be distributed uniformly among the locations $\{\ell_{e, E_i}\}_{e \in U}$. As a result, the treasure m_{e^*} cannot be allocated to any location ℓ_{e, E_i} for all $e \in E_i$ and $i \notin C$. However, once the treasure m_{e^*} is allocated to any other location ℓ , the attacker's utility to attack ℓ would be ∞ since the defender has already exhausted the defensive resources, which produces a contradiction.

Combining Lemma 19, Lemma 20, and Theorem 4, we finish the reduction.

D.6 Only Locations are Homogeneous

We will again reduce from 3-SET COVER (see Definition 4). Given an instance of 3-SET COVER, we construct a security game with $|L| = |U| + n + Z$ locations, $|M| = |U| + n + Z$ treasures, and $|R| = |U| + Z$ defensive resources.

When the locations are homogeneous, the location-treasure importance function such that v can be simplified as $v(m)$ representing the utility loss of the defender if a location with treasure $m \in M$ allocated is attacked without defense. Moreover, the defense-quality function q can be simplified as $q(m, r)$ characterizing the effectiveness of defending treasure m with defensive resource r . Let $\vec{g}(\Pi)$ be the fractional matching between L and M corresponding to the prior Π , and let \vec{d} be the fractional matching between L and R corresponding to a defender's mixed strategy of allocating defensive resources. Then, the attacker's utility of attacking location ℓ is given by

$$\sum_{m \in M} g_{\ell, m}(\Pi) \cdot v(m) \cdot \left(1 - \sum_{r \in R} d_{\ell, r} \cdot q(m, r) \right).$$

For each $e \in U$, we construct a treasure denoted by m_e , and a defensive resource denoted by r_e . For each E_i with $i \in [Z]$, we construct a defensive resource denoted by r_{E_i} . In addition, we construct n treasures, each of which denoted by m_k^{fix} for $k \in [n]$, and Z treasures, each of which denoted by m_k^{bad} for $k \in [Z]$. The treasure importance function v is given by

- $v(m_e) = \infty$ for all $e \in U$;
- $v(m_k^{\text{fix}}) = 1$ for all $k \in [n]$;
- $v(m_k^{\text{bad}}) = Z/(4n)$ for all $k \in [Z]$.

Moreover, the defense-quality importance function is given by

- $q(m_e, r_{e'}) = 1$ if $e = e'$; otherwise, $q(m_e, r_{e'}) = 0$;
- $q(m_e, r_{E_i}) = 1$ if $e \in E_i$; otherwise, $q(m_e, r_{E_i}) = 0$;
- For $k \in [n]$, $q(m_k^{\text{fix}}, r_e) = 0$ for all $e \in U$ and $q(m_k^{\text{fix}}, r_{E_i}) = 0$ for all $i \in [Z]$;
- For $k \in [Z]$, $q(m_k^{\text{bad}}, r_e) = 0$ for all $e \in U$ and $q(m_k^{\text{bad}}, r_{E_i}) = 1$ for all $i \in [Z]$.

The team max-min game is directly induced from this security game. We finish the reduction by showing that there exists a solution to the instance of 3-SET COVER if and only if the game value of the constructed team max-min game is at least $-1/4$.

Lemma 21. *If there exists a solution to the instance of 3-SET COVER, then the game value of the constructed team max-min game is at least $-1/4$.*

Proof. Let E_{k_1}, \dots, E_{k_n} be such a 3-set cover. It suffices to exhibit a scheme which induces an expected utility of the defender at least $-1/4$. For convenience, let $C = \{k_1, \dots, k_n\}$ and the set of locations be $L = \{\ell_e\}_{e \in U} \cup \{\ell_{E_i}\}_{i \in C} \cup \{\ell_k^{\text{bad}}\}_{k \in [n]}$. Consider the following prior Π

- For each $e \in U$, $g_{\ell_e, m_e}(\Pi) = 3/4$ and $g_{\ell_{E_i}, m_e} = 1/4$ for $e \in E_i$ and $i \in C$;
- For each $k \in [Z]$, $g_{\ell_k^{\text{bad}}, m_k^{\text{bad}}}(\Pi) = 1$;

Notice that there are n units of treasures from $\{m_k^{\text{fix}}\}_{k \in [n]}$ that have not been allocated while for each location in $\{\ell_e\}_{e \in U} \cup \{\ell_{E_i}\}_{i \in C}$, there is $1/4$ space left, resulting in $1/4 \cdot (|U| + |C|) = n$ space left. Therefore, we can simply allocate n units of treasures from $\{m_k^{\text{fix}}\}_{k \in [n]}$ to fill the locations in $\{\ell_e\}_{e \in U} \cup \{\ell_{E_i}\}_{i \in C}$. Moreover, the mixed strategy of defending is as follows:

- For each $e \in U$, $d_{\ell_e, r_e} = 1$;
- For each $i \in C$, $d_{\ell_{E_i}, r_{E_i}} = 1$;

Notice that there are $Z - n$ units of defensive resources from $\{r_{E_i}\}_{i \notin C}$ that have not been allocated and for each location ℓ_k^{bad} for $k \in [Z]$, we allocate $(1 - n/Z)$ units of defensive resources from $\{r_{E_i}\}_{i \notin C}$ to defend it. It is now straightforward to verify that for any location, the attacker's utility of attacking is exactly $1/4$.

Lemma 22. *If the game value of the constructed team max-min game is at least $-1/4$, then there exists a solution to the instance of 3-SET COVER.*

Proof. For convenience, let $\bar{d}_\ell = \sum_{e \in U} d_{\ell, r_e} + \sum_{i \in [Z]} d_{\ell, r_{E_i}}$ be the total probability that a location ℓ is defended. Moreover, let $\bar{L} = \{\ell \in L \mid \bar{d}_\ell < 1\}$. We first observe that for a location $\ell \in L$, if $g_{\ell, m_e}(II^*) > 0$, then we must have $\ell \notin \bar{L}$; or otherwise the attacker's utility of attacking ℓ would be ∞ . Hence, for $\ell \in \bar{L}$, the defender can only allocate treasures from $\{m_k^{\text{fix}}\}_{k \in [n]}$ and $\{m_k^{\text{bad}}\}_{k \in [Z]}$.

Since the number of treasures equals to the number of locations, we must allocate a unit of treasure in total to each location. Assume that for $\ell \in \bar{L}$, the defender can allocate κ units of treasures from $\{m_k^{\text{fix}}\}_{k \in [n]}$ and $(1 - \kappa)$ units of treasures $\{m_k^{\text{bad}}\}_{k \in [Z]}$. Recall that the defensive resources are not effective to defend treasures from $\{m_k^{\text{fix}}\}_{k \in [n]}$. Therefore, to achieve a game value $-1/4$, we have

$$\kappa + (1 - \kappa) \cdot \frac{Z}{4n} \cdot (1 - \bar{d}_\ell) \leq \frac{1}{4},$$

which implies that

$$\bar{d}_\ell \geq 1 - \left(4 - \frac{3}{1 - \kappa}\right) \cdot \frac{n}{Z}, \quad (9)$$

which implies that, for any $\ell \in \bar{L}$, we have $\bar{d}_\ell \geq 1 - n/Z$.

Proposition 7. $|\bar{L}| = Z$. *Moreover, locations in \bar{L} only host the treasures from $\{m_k^{\text{bad}}\}_{k \in [Z]}$, and there is no treasure from $\{m_k^{\text{fix}}\}_{k \in [n]}$ that is allocated to locations in \bar{L} . Consequently, the defender allocates $(1 - n/Z)$ units of defensive resources from $\{r_{E_i}\}_{i \in [Z]}$ to defend each location in \bar{L} .*

Proof. First, if $|\bar{L}| = Z - K$ for $K > 0$, then recall that for $\ell \notin \bar{L}$, $\bar{d}_\ell = 1$. Therefore, we have

$$\sum_{\ell \in L} \bar{d}_\ell \geq |L \setminus \bar{L}| + |\bar{L}| \cdot \left(1 - \frac{n}{Z}\right) = 3n + Z + K \cdot \frac{n}{Z} > 3n + Z = |R|,$$

which produces a contradiction.

On the other hand, if $|\bar{L}| = Z + K$ for $K > 0$. Since there are Z units of treasures from $\{m_k^{\text{bad}}\}_{k \in [Z]}$ in total, the defender must allocate at least K units of treasures from $\{m_k^{\text{fix}}\}_{k \in [n]}$ to locations in \bar{L} . Assume that there are κ_ℓ units of treasures from $\{m_k^{\text{fix}}\}_{k \in [n]}$ in total that are allocated to a location $\ell \in \bar{L}$. Note

that we have $\kappa_\ell \in [0, 1/4]$ for all $\ell \in \bar{L}$ and $\sum_{\ell \in \bar{L}} \kappa_\ell \geq K$, which implies that $K \leq 1/4 \cdot |\bar{L}|$. From (9), we have:

$$\begin{aligned} \sum_{\ell \in \bar{L}} \bar{d}_\ell &\geq \sum_{\ell \in \bar{L}} 1 - \left(4 - \frac{3}{1 - \kappa_\ell}\right) \cdot \frac{n}{Z} \\ &\geq |\bar{L}| \cdot \left(1 - \left(4 - \frac{3}{1 - \sum_{\ell \in \bar{L}} \kappa_\ell / |\bar{L}|}\right) \cdot \frac{n}{Z}\right) \\ &\geq (Z + K) \cdot \left(1 + \left(\frac{3K}{Z} - 1\right) \cdot \frac{n}{Z}\right), \end{aligned}$$

where the second inequality follows the Jensen's inequality. However, this results in

$$\sum_{\ell \in L} \bar{d}_\ell = \sum_{\ell \in \bar{L}} \bar{d}_\ell + \sum_{\ell \notin \bar{L}} \bar{d}_\ell \geq (Z + K) \cdot \left(1 + \left(\frac{3K}{Z} - 1\right) \cdot \frac{n}{Z}\right) + (4n - K) > 3n + Z = |R|,$$

because we have

$$\begin{aligned} &(Z + K) \cdot \left(1 + \left(\frac{3K}{Z} - 1\right) \cdot \frac{n}{Z}\right) + (4n - K) > 3n + Z \\ \Leftrightarrow &(Z + K) \cdot \left(1 + \left(\frac{3K}{Z} - 1\right) \cdot \frac{n}{Z}\right) > Z + K - n \\ \Leftrightarrow &(Z + K) \cdot \left(\frac{3K}{Z} - 1\right) \cdot \frac{n}{Z} > -n \\ \Leftrightarrow &(Z + K) \cdot (3K - Z) > -Z^2 \\ \Leftrightarrow &3K^2 + 2KZ > 0. \end{aligned}$$

Therefore, we must have $|\bar{L}| = Z$. When $|\bar{L}| = Z$, from (9), we have

$$\sum_{\ell \in L} \bar{d}_\ell \geq |L \setminus \bar{L}| + |\bar{L}| \cdot \left(1 - \frac{n}{Z}\right) = 3n + Z = |R|,$$

while the equality is obtained only if there is no treasure from $\{m_k^{\text{fix}}\}_{k \in [n]}$ that is allocated to locations in \bar{L} .

According to Proposition 7, all treasures from $\{m_k^{\text{bad}}\}_{k \in [Z]}$ have been allocated. Moreover, notice that the treasures from $\{m_k^{\text{fix}}\}_{k \in [n]}$ can only be allocated to the remaining locations in $L \setminus \bar{L}$ where $|L \setminus \bar{L}| = 4n$. In addition, each location can host at most 1/4 units of treasures from $\{m_k^{\text{fix}}\}_{k \in [n]}$ while there are n units in total, and therefore, we have that each location in $L \setminus \bar{L}$ hosts exactly 1/4 units of treasures from $\{m_k^{\text{fix}}\}_{k \in [n]}$.

For convenience, let $\tilde{L} = \{\ell \in L \setminus \bar{L} \mid \exists e \neq e', g_{\ell, e}(\Pi^*) > 0 \text{ and } g_{\ell, e'}(\Pi^*) > 0\}$ be the set of locations in $L \setminus \bar{L}$ that host at least two different treasures from $\{m_e\}_{e \in U}$.

Proposition 8. $|\tilde{L}| = n$. Moreover, for each $\ell \in \tilde{L}$, the defender must allocate a unit of treasure from $\{r_{E_i}\}_{i \in [Z]}$ to defend it.

Proof. We first argue that $|\tilde{L}| \geq n$. Because if not, then there are at least $3n + 1$ locations $\ell \in L \setminus \tilde{L}$ in which there exists $e \in U$ with $g_{\ell,e}(\Pi^*) = 3/4$. Since there are $|U| = 3n$ elements in total, by the Pigeonhole principle, there exists $e \in U$ such that $\sum_{\ell \in L \setminus \tilde{L}} g_{\ell,e}(\Pi^*) \geq 3/2$. A contradiction.

Next, by Proposition 7, there are n units of defensive resources from $\{r_{E_i}\}_{i \in [Z]}$ left. For a location $\ell \in \tilde{L}$ hosts two different treasures, the defender must allocate a unit of treasure from $\{r_{E_i}\}_{i \in [Z]}$ to defend it; or otherwise, the attacker's utility of attacking ℓ is ∞ .

Proposition 8 implies that for the remaining $3n$ locations $\ell \in L \setminus (\tilde{L} \cup \bar{\tilde{L}})$, each location hosts $3/4$ units of a unique treasure. Since $|U| = 3n$, consequently, for each treasure m_e for $e \in U$, there are $1/4$ units left. Observe that for location $\ell \in \tilde{L}$, it cannot host more than 3 different treasures from $\{m_e\}_{e \in U}$ since there is no defensive resource that is effective for more than 3 different treasures from $\{m_e\}_{e \in U}$. Hence, for each $\ell \in \tilde{L}$, it hosts exactly 3 different treasures $m_{e_1}, m_{e_2}, m_{e_3} \in \{m_e\}_{e \in U}$, each of which with $1/4$ units. Moreover, to achieve a game value $-1/4$, there must exist a subset $E_i = \{e_1, e_2, e_3\}$ so that the defender can defend such a location with a unit of defensive resource r_{E_i} . Thus, there exists a solution to the instance of 3-SET COVER.

Combining Lemma 21, Lemma 22, and Theorem 4, we finish the reduction.