

COMPUTING EQUILIBRIUM BEYOND UNILATERAL DEVIATION

Mingyang Liu, Gabriele Farina & Asuman Ozdaglar

LIDS, EECS

Massachusetts Institute of Technology

Cambridge, MA 02139, USA

{liumy19, gfarina, asuman}@mit.edu

ABSTRACT

Most familiar equilibrium concepts, such as Nash and correlated equilibrium, guarantee only that no single player can improve their utility by deviating unilaterally. They offer no guarantees against profitable coordinated deviations by coalitions. Although the literature proposes notions to address multilateral deviations (*e.g.*, strong Nash and coalition-proof equilibrium), these generally fail to exist. In this paper, we study a solution concept that accommodates multi-player deviations and is guaranteed to exist. We prove a fixed-parameter lower bound on the complexity of computing such an equilibrium and present an algorithm that matches this bound.

1 INTRODUCTION

Most equilibrium concepts studied so far, such as Nash equilibrium (NE) (Nash, 1950), correlated equilibrium (CE) (Aumann, 1974), coarse correlated equilibrium (CCE) (Moulin & Vial, 1978), and Stackelberg equilibrium (Von Stackelberg, 2010), guarantee only that *no individual player* can gain by deviating unilaterally. However, they offer no guarantees when multiple players deviate simultaneously by forming a coalition. In this paper, we address the following question:

What is an appropriate notion for capturing multilateral deviations, and how can it be computed?

Previous notions that address coalition deviations, such as strong NE (Aumann, 1959)¹ and coalition-proof equilibrium (Bernheim et al., 1987), mostly fail to exist in general games (unlike NE). Therefore, instead of searching for a joint strategy immune to all coalition deviations, we focus on computing a joint strategy that *minimizes* the maximum average gain achievable by any coalition, which is the average of improvements over all its members. In other words, we can compute the most stable strategy profile, even if a perfectly stable one does not exist. We refer to this notion as the *Minimum Average-Strong Equilibrium* (MASE).

The difficulty of this optimization problem naturally depends on the complexity of the interactions between players. To formalize this, we introduce the Utility Dependency Graph, $\mathcal{G}(\mathcal{V}, \mathcal{E})$, where each player is a vertex. An edge connects two players, i and j , if and only if there is some player k whose utility is affected by the actions of both i and j . Intuitively, an edge (i, j) signifies that the actions of i and j are linked, as they jointly influence the payoff of some player k . This graph provides a clear map of the game’s interaction structure, and its properties can help us understand the computational complexity of finding the MASE. For games with simple interaction structures (*e.g.*, a sparse Utility Dependency Graph), one might expect to compute the MASE efficiently.

However, computing the MASE is computationally challenging in the general case. We establish two key hardness results that delineate the problem’s complexity.

First, the problem is fundamentally harder than finding equilibria like NE or CE. In those cases, we are solving a feasibility problem: finding a strategy where the maximum gain from deviating is at

¹In this paper, we use the term *strong equilibrium* to broadly refer to any equilibrium concept that considers multilateral deviations.

most zero. For MASE, we must solve an optimization problem: minimizing this maximum gain. This distinction is crucial, and we show that even for the simplest case of single-player deviations (*i.e.*, coalitions of size one), approximating the MASE value to within a factor that is inverse polynomial in the number of players is NP-hard. This indicates that even without considering complex coalitions, the problem is intractable without additional assumptions on the game’s structure.

Second, we show that this complexity is intrinsically tied to the structure of the Utility Dependency Graph. Building on the strong exponential time hypothesis (SETH) (Impagliazzo & Paturi, 2001) (see Theorem 4.3 for details), we prove that solving MASE requires time that is at least exponential in the treewidth² of the Utility Dependency Graph. This holds even when we only consider coalitions of a constant size. This result demonstrates that the treewidth is a fundamental barrier, and an exponential dependence on it is unavoidable.

Finally, we present a positive result that matches the lower-bound time complexity, up to exponential factors. We develop an algorithm that computes the MASE with a time complexity that is exponential in the treewidth of the Utility Dependency Graph. This demonstrates that our hardness result is tight and establishes the treewidth as the definitive parameter characterizing the complexity of computing the MASE. While the problem is hard in general, it becomes tractable for games where the underlying interaction structure is not too complex.

To summarize, our contributions are as follows:

1. Complexity characterization: We establish lower bounds on the computational complexity of computing the minimally deviated equilibrium, showing that the problem is inherently tied to the treewidth of the Utility Dependency Graph.
2. Algorithmic contribution: We design an algorithm that efficiently computes the minimally deviated equilibrium, achieving a running time that matches the established lower bound up to exponential dependence on treewidth.

2 PRELIMINARIES

For any vector $\mathbf{x} \in \mathbb{R}^n$, we use x_i to denote its i^{th} element and $\|\mathbf{x}\|_p$ to denote its p -norm. By default, $\|\mathbf{x}\|$ refers to the 2-norm. For a positive integer N , let $[N] := 1, 2, \dots, N$. We denote the $(n - 1)$ -dimensional probability simplex by $\Delta^n := \{\mathbf{x} \in [0, 1]^n : \sum_{i=1}^n x_i = 1\}$. More generally, for any discrete set S , let $|S|$ denote its cardinality and write Δ^S for the probability simplex over S , whose coordinates are indexed by elements of S (*e.g.*, $\Delta^n = \Delta^{[n]}$). Similarly, \mathbb{R}^S denotes the $|S|$ -dimensional real vector space with coordinates indexed by S . For any set S_1, S_2 , $S_1 \times S_2$ denotes the Cartesian product of sets S_1 and S_2 . Finally, we let $\mathbb{1}(\text{argument})$ denote the indicator function, which equals 1 if the argument is true and 0 otherwise.

2.1 GAMES

A game is represented as a tuple $(N, \{\mathcal{A}_i\}_{i=1}^N, \{\mathcal{U}_i\}_{i=1}^N, \mathcal{S})$, where

- N is the number of players.
- \mathcal{A}_i is the action set of player i . For convenience, let $\mathcal{A} := \times_{i=1}^N \mathcal{A}_i$ denote the joint action set.
- $\mathcal{U}_i : \mathcal{A} \rightarrow [0, 1]$ is the utility function of player $i \in [N]$.
- \mathcal{S} is the set of coalitions, which is a set of subsets of players. For example, if only unilateral deviations are allowed (as in Nash equilibrium or coarse correlated equilibrium), then $\mathcal{S} = \{\{1\}, \{2\}, \dots, \{N\}\}$.

For notational simplicity, for any subset of players $S \subseteq [N]$, we write $\mathcal{A}_S := \times_{i \in S} \mathcal{A}_i$. Throughout the paper, let $A := \max_{i \in [N]} |\mathcal{A}_i|$ denote the size of the largest action set.

For any joint action $\mathbf{a} \in \mathcal{A}$, let a_i denote the action of player i , and let $\mathbf{a}_{-i} = (a_1, a_2, \dots, a_{i-1}, a_{i+1}, \dots, a_N)$ be the joint action of all players except i . More generally, for any subset $S \subseteq [N]$, we write \mathbf{a}_{-S} for the joint action of players outside S .

²Treewidth can be thought of as a formal measure of how sparse and "tree-like" a graph is.

	Confess (C)	Defect (D)
Confess (C)	(0.6, 0.6)	(0, 1)
Defect (D)	(1, 0)	(0.2, 0.2)

Table 1: Utility matrix of the Prisoner’s Dilemma. Each entry (a, b) denotes the payoff of the row player (a) and the column player (b) .

2.2 SUCCINCT REPRESENTATION

This paper focuses on multi-player games with a succinct representation. Specifically, each utility function \mathcal{U}_i can be encoded using a number of bits polynomial in the number of players N , rather than requiring $\mathcal{O}\left(N \prod_{i=1}^N |\mathcal{A}_i|\right)$ bits as in the general case. Examples of succinctly represented games include polymatrix games (Howson Jr, 1972; Eaves, 1973) and congestion games (Rosenthal, 1973). Throughout the paper, we call an algorithm *efficient* if its running time is polynomial in N , as opposed to polynomial in $\prod_{i=1}^N |\mathcal{A}_i|$. We focus on succinct games because MASE can otherwise be solved by a linear program whose size grows exponentially with N (see Section B). Moreover, the study of strong equilibrium is particularly compelling in large games, where exponential dependence on N is computationally prohibitive.

3 MINIMUM AVERAGE-STRONG EQUILIBRIUM (MASE)

Several notions of strong equilibrium have been proposed, including the strong Nash equilibrium (NE) (Aumann, 1959), the sum-strong NE (Hoefer, 2013) (no improvement on the total gain of any coalition), and coalition-proof equilibrium (Bernheim et al., 1987). However, none of these exist in general games. To build intuition, we first illustrate why a strong NE does not exist in the Prisoner’s Dilemma. We further show that the problem persists even when correlated strategies are allowed.

Lemma 3.1. In the Prisoner’s Dilemma, no strong Nash nor strong correlated equilibrium exists when $\mathcal{S} = \{\{1\}, \{2\}, \{1, 2\}\}$ is the set of all non-empty subsets of players.

Since correlated equilibria include all Nash equilibria, it suffices to examine strong correlated equilibria. A strong correlated equilibrium is a correlated joint strategy where no subset of players (a coalition) can jointly deviate in a way that strictly improves the utility of all its members.

As shown in Table 1, any strategy with positive weight on (C, C) , (C, D) , (D, C) yields a profitable deviation for at least one singleton coalition, $\{1\}$ or $\{2\}$. Conversely, placing all weight on (D, D) creates a deviation to (C, C) that benefits the coalition $\{1, 2\}$. Thus, no strong NE exists.

The failure of strong equilibria arises because coalition objectives may conflict, making it impossible to find a strategy that simultaneously satisfies all coalitions. Motivated by this, rather than requiring exact immunity to deviations, we instead seek to minimize the incentive to deviate since a minimizer always exists by Weierstrass theorem. This leads to the following definition of *Minimum Average-Strong Equilibrium* (MASE):

$$\pi^* \in \operatorname{argmin}_{\pi \in \Delta^{\mathcal{A}}} \max_{S \in \mathcal{S}} \max_{\hat{\mathbf{a}}_S \in \mathcal{A}_S} \frac{1}{|S|} \sum_{i \in S} \mathbb{E}_{\mathbf{a} \sim \pi} [\mathcal{U}_i(\hat{\mathbf{a}}_S, \mathbf{a}_{-S}) - \mathcal{U}_i(\mathbf{a})]. \quad (\text{MASE})$$

Intuitively, (MASE) selects the correlated strategy $\pi \in \Delta^{\mathcal{A}}$ that minimizes the maximum average gain attainable by any coalition across all possible coalitions. If this value is less than or equal to zero, then no coalition can simultaneously deviate in a way that yields a strictly positive total gain. Note that the algorithm presented in this paper can also be extended to handle any weighted average over the coalition.

A correlated strategy $\pi \in \Delta^{\mathcal{A}}$ is called an ϵ -MASE if

$$\begin{aligned} & \max_{S \in \mathcal{S}} \max_{\hat{\mathbf{a}}_S \in \mathcal{A}_S} \frac{1}{|S|} \sum_{i \in S} \mathbb{E}_{\mathbf{a} \sim \pi} [\mathcal{U}_i(\hat{\mathbf{a}}_S, \mathbf{a}_{-S}) - \mathcal{U}_i(\mathbf{a})] \\ & \leq \max_{S \in \mathcal{S}} \max_{\hat{\mathbf{a}}_S \in \mathcal{A}_S} \frac{1}{|S|} \sum_{i \in S} \mathbb{E}_{\mathbf{a} \sim \pi^*} [\mathcal{U}_i(\hat{\mathbf{a}}_S, \mathbf{a}_{-S}) - \mathcal{U}_i(\mathbf{a})] + \epsilon. \end{aligned}$$

A Strong Nash equilibrium requires that for any deviating coalition, at least one member does not strictly improve their utility. In contrast, ϵ -MASE aims to minimize the average improvement over all players within any given coalition. From another perspective, ϵ -MASE minimizes the incentive to deviate, even when coalition members can freely reallocate utility within the coalition.

4 HARDNESS OF SOLVING MASE

Recall that we call an algorithm *efficient* if it runs in time polynomial in N . In this section, we first establish the computational hardness of computing ϵ -MASE.

Theorem 4.1. Computing ϵ -MASE is NP-hard, even when \mathcal{S} only contains singletons (coalitions of size one) and $1/\epsilon$ is polynomial in the number of players.

The proof is deferred to Section C. Importantly, Theorem 4.1 highlights a fundamental distinction from the case of CCE, which can be computed efficiently (Papadimitriou & Roughgarden, 2008). The reason is that for CCE it suffices to find a correlated strategy $\pi \in \Delta^{\mathcal{A}}$ such that the deviation gap, $\max_{i \in [N]} \max_{\hat{a}_i \in \mathcal{A}_i} \mathbb{E}_{\mathbf{a} \sim \pi} [\mathcal{U}_i(\hat{a}_i, \mathbf{a}_{-i}) - \mathcal{U}_i(\mathbf{a})]$, is less or equal to zero, whereas here we must find a strategy that minimizes the gap. Together with the linear programming characterization in Section B, this implies that computing ϵ -MASE is actually NP-complete. In fact, Anagnostides et al. (2025) recently showed that even minimizing the *average* deviation gap of CCE across all players (instead of the maximum gap considered here) is also NP-complete.

4.1 FIXED PARAMETER LOWER BOUND

Next, we present a more refined hardness result: a fixed-parameter lower bound for computing MASE. To do so, we first formalize the notion of dependencies among players’ utilities.

For each player $i \in [N]$, define the relevant set $\mathcal{N}(i) \subseteq [N]$ consisting of all players $j \in [N]$ (including $j = i$) such that the action of j can affect the utility of i . Formally, $j \in [N]$ is in $\mathcal{N}(i)$ if and only if there exist $\mathbf{a}_{-j} \in \mathcal{A}_{-j}$ and $a_j, a'_j \in \mathcal{A}_j$ such that $\mathcal{U}_i(a_j, \mathbf{a}_{-j}) \neq \mathcal{U}_i(a'_j, \mathbf{a}_{-j})$. This leads to the following graph representation.

Definition 4.2 (Utility Dependency Graph). The utility dependence graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ is an undirected graph with vertex set $\mathcal{V} = [N]$ representing the players, and edge set $\mathcal{E} = \bigcup_{k \in [N]} \{(i, j) \mid i, j \in \mathcal{N}(k), i \neq j\}$.

Since \mathcal{U}_i depends only on the actions of players in $\mathcal{N}(i)$, we may equivalently write $\mathcal{U}_i(\mathbf{a}_C) = \mathcal{U}_i(\mathbf{a}_C, \mathbf{a}'_{-C})$ for arbitrary $\mathbf{a}'_{-C} \in \mathcal{A}_{-C}$, where $C \supseteq \mathcal{N}(i)$. It is worth noting that this definition differs from the graph of a graphical game (Kakade et al., 2003; Kearns et al., 2001). Here, players i and j are connected if both influence the utility of some other player k , even if i and j do not directly affect each other. Whereas in graphical games, two players i and j are connected if and only if at least one can influence the other’s utility.

With this graph structure in place, we can connect the hardness of computing MASE to the treewidth of \mathcal{G} . Intuitively, treewidth measures how close a graph is to being a tree: the treewidth of \mathcal{G} is one when \mathcal{G} is a tree, and it is $N - 1$ when \mathcal{G} is a complete graph. Throughout this section, let \mathcal{O}^* denote asymptotic complexity with factors polynomial in N suppressed.

Theorem 4.3 (Treewidth). Suppose a tree decomposition of the Utility Dependency Graph is given. Under the Strong Exponential Time Hypothesis (SETH) (Impagliazzo & Paturi, 2001),³ (MASE) cannot be computed in $\mathcal{O}^*((A - \zeta)^{\text{tw}(\mathcal{G})})$ for any $\zeta > 0$. Moreover, under the additional assumption that $\text{BPP} = \text{P}$,⁴ $\frac{1}{9N^2}$ -approximate MASE cannot be computed in $\mathcal{O}^*((A - \zeta)^{\text{tw}(\mathcal{G})})$ for any $\zeta > 0$.

The proof is deferred to Section C.2. For approximate MASEs we assume $\text{BPP} = \text{P}$, which is standard in the literature (Arora & Barak, 2009), because the reduction involves sampling joint actions from

³SETH assumes that SAT cannot be solved in $\mathcal{O}^*((2 - \zeta)^n)$ for any $\zeta > 0$, where n is the number of variables in the SAT instance.

⁴This assumption implies that any problem with a polynomial-time randomized algorithm also has a polynomial-time deterministic algorithm.

the approximate MASE. Since enumerating all $\mathbf{a} \in \mathcal{A}$ is computationally infeasible, we rely on randomized sampling. This yields only a randomized algorithm for the original NP-hard problem, and the assumption $\text{BPP}=\text{P}$ ensures that such a randomized algorithm can be derandomized into a deterministic one, completing the reduction.

Theorem 4.3 shows that the computational complexity of solving MASE is inherently tied to the treewidth of the Utility Dependency Graph. Intuitively, when the treewidth is large, each player’s utility depends on many others, making even the evaluation of coalition deviations computationally demanding (enumerating over all $\hat{\mathbf{a}}_S \in \mathcal{A}_S$). In contrast, when the treewidth is small, such as zero (each player’s utility depends only on their own action), computing MASE becomes trivial, since each player’s utility can be maximized independently. In polymatrix games (Eaves, 1973), the treewidth of the Utility Dependency Graph can be bounded by that of its corresponding graph. Further details are provided in Section I.

5 EFFICIENT COMPUTATION OF MASE

Although an (MASE) lives in an exponentially large space (of size $|\mathcal{A}|$), it can still be computed efficiently. This is because the equilibrium always admits a compact representation.

Theorem 5.1 (Efficient Representation). For any $\epsilon \geq 0$, at least one of the ϵ -MASE can be represented as a linear combination of $\sum_{S \in \mathcal{S}} |S| \cdot A^{\text{tw}(\mathcal{G})} + 1$ pure strategies, where $\text{tw}(\mathcal{G})$ is the treewidth of Utility Dependency Graph.

The proof is deferred to Section D. Intuitively, Theorem 5.1 shows that there must be an ϵ -MASE that always has a sparse representation. Since a pure strategy can be encoded by the index of its unique action with nonzero probability, this compactness makes computation tractable.

5.1 META-GAME BETWEEN THE CORRELATOR AND DEVIATOR

To compute an (MASE), we reformulate the problem as a *meta-game* between two players: the *correlator* and the *deviator* (Hart & Schmeidler, 1989). The correlator chooses the correlated strategy $\pi \in \Delta^{\mathcal{A}}$, while the deviator selects deviations. The game is zero-sum: the correlator aims to minimize the coalition’s gain from deviation, and the deviator aims to maximize it. Formally:

$$\min_{\pi \in \Delta^{\mathcal{A}}} \max_{\mu \in \Delta^{\times_{S \in \mathcal{S}} \mathcal{A}_S}} F(\pi, \mu), \quad (5.1)$$

where

$$F(\pi, \mu) := \sum_{S \in \mathcal{S}} \sum_{\hat{\mathbf{a}}_S \in \mathcal{A}_S} \frac{\mu(S, \hat{\mathbf{a}}_S)}{|S|} \sum_{i \in S} \mathbb{E}_{\mathbf{a} \sim \pi} [\mathcal{U}_i(\hat{\mathbf{a}}_S, \mathbf{a}_{-S}) - \mathcal{U}_i(\mathbf{a})]. \quad (5.2)$$

Here, we extend the deviator’s decision space from a discrete to a continuous set. This relaxation does not strengthen the deviator, since the objective is linear in μ , and the maximum is always attained at an extreme point. Therefore, (5.1) is equivalent to the original definition in (MASE).

A natural idea is to apply no-regret learning algorithms simultaneously for the correlator and deviator. However, directly updating the full distributions π and μ is infeasible, because the underlying spaces are exponentially large.

Fortunately, Theorem 5.1 implies that maintaining the full distributions is unnecessary: it suffices to keep track of a polynomial number of pure strategies, and use their convex combination as the approximate equilibrium. This motivates our use of *Follow the Perturbed Leader* (FTPL) (Hazan et al., 2016), where each decision at a timestep is a pure strategy, which can be represented compactly.

Let $\pi^{(t)} \in \Delta^{\mathcal{A}}$ and $\mu^{(t)} \in \Delta^{\times_{S \in \mathcal{S}} \mathcal{A}_S}$ denote the decision variables at timestep $t \geq 1$ for the correlator and the deviator, respectively. The interaction between these two players can be described

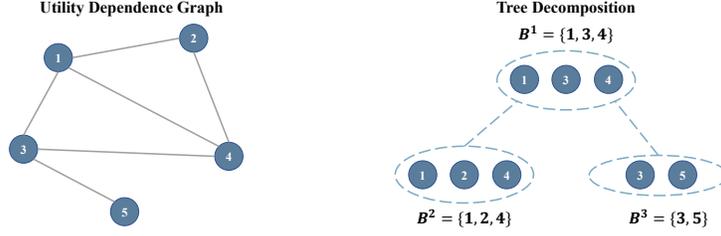


Figure 1: An illustration of a tree decomposition of the Utility Dependency Graph.

by the update rule

$$\begin{aligned} \pi^{(t+1)} &\in \operatorname{argmin}_{\pi \in \Delta^{\mathcal{A}}} \sum_{\tau=1}^t F(\pi, \mu^{(\tau)}) - \langle \tilde{\mathbf{n}}^{(t+1)}, \pi \rangle \\ \mu^{(t+1)} &\in \operatorname{argmax}_{\mu \in \Delta^{\mathcal{S}} \times \mathcal{S}^{\mathcal{A}_S}} \sum_{\tau=1}^t F(\pi^{(\tau)}, \mu) + \langle \tilde{\mathbf{m}}^{(t+1)}, \mu \rangle, \end{aligned} \quad (5.3)$$

where $\tilde{\mathbf{n}}^{(t+1)}$ and $\tilde{\mathbf{m}}^{(t+1)}$ are noise vectors sampled independently at each timestep from some distribution, which we will specify later. These noise terms play the role of regularizers in on-line mirror descent (OMD) (Hazan et al., 2016), ensuring stability in the updates by controlling $\mathbb{E} [\|\pi^{(t+1)} - \pi^{(t)}\|]$ and $\mathbb{E} [\|\mu^{(t+1)} - \mu^{(t)}\|]$.

Since $F(\pi, \mu)$ is bilinear in (π, μ) , both the minimization and maximization problems admit solutions at vertices of their respective decision spaces. In other words, the argmin for the correlator and the argmax for the deviator always contain at least one pure strategy.

In what follows, we will explain in detail how to update π efficiently under this framework. The update of μ is deferred to Section G.

5.2 EFFICIENT UPDATE OF π

The key step in updating $\pi^{(t+1)}$ is to select a pure strategy, *i.e.*, a joint action $\mathbf{a}^{(t+1)} \in \mathcal{A}$ with $\pi^{(t+1)}(\mathbf{a}^{(t+1)}) = 1$, that minimizes the objective. To gain insight into this update rule, we first examine how to compute $\operatorname{argmin}_{\pi \in \Delta^{\mathcal{A}}} F(\pi, \mu)$ for a fixed μ .

Suppose we want to find a joint action $\tilde{\mathbf{a}} \in \mathcal{A}$ such that the pure strategy $\tilde{\pi}$ with $\tilde{\pi}(\tilde{\mathbf{a}}) = 1$ minimizes $F(\tilde{\pi}, \mu)$. Expanding the definition, we obtain

$$\begin{aligned} F(\tilde{\pi}, \mu) &= \sum_{i=1}^N \sum_{S \in \mathcal{S}: i \in S} \sum_{\hat{\mathbf{a}}_S \in \mathcal{A}_S} \frac{\mu(S, \hat{\mathbf{a}}_S)}{|S|} (\mathcal{U}_i(\hat{\mathbf{a}}_S, \tilde{\mathbf{a}}_{-S}) - \mathcal{U}_i(\tilde{\mathbf{a}})) \\ &= \sum_{i=1}^N \sum_{S \in \mathcal{S}: i \in S} \sum_{\hat{\mathbf{a}}_S \in \mathcal{A}_S} \frac{\mu(S, \hat{\mathbf{a}}_S)}{|S|} (\mathcal{U}_i(\hat{\mathbf{a}}_{S \cap \mathcal{N}(i)}, \tilde{\mathbf{a}}_{\mathcal{N}(i) \setminus S}) - \mathcal{U}_i(\tilde{\mathbf{a}}_{\mathcal{N}(i)})). \end{aligned}$$

Therefore, for each candidate $\tilde{\mathbf{a}} \in \mathcal{A}$, only the local actions $\tilde{\mathbf{a}}_{\mathcal{N}(i)}$ matter for the **expression** above. If we can evaluate this **expression** efficiently,⁵ then for each player $i \in [N]$ we may search for $\tilde{\mathbf{a}}_{\mathcal{N}(i)}$ that minimizes it. However, a difficulty arises because $\mathcal{N}(i)$ and $\mathcal{N}(j)$ may overlap across different players. Hence, we must ensure that the local assignments remain globally consistent.

To address this, we now introduce the concept of a tree decomposition and show how it enables us to optimize F efficiently. Throughout the paper, we assume that a tree decomposition is given, and analyze the complexity only with respect to this decomposition.

Tree decomposition. A tree decomposition $\mathcal{T} := B^1, B^2, \dots, B^K$ of the Utility Dependency Graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ is a tree with K nodes (bags), each $B^k \subseteq \mathcal{V}$ where $\mathcal{V} = [N]$, satisfying the following properties (Diestel, 2025):

⁵This is possible since μ is a linear combination of pure strategies when updated according to (5.3).

1. $\bigcup_{k=1}^K B^k = [N]$.
2. For every edge $(i, j) \in \mathcal{E}$, there exists k with $\{i, j\} \subseteq B^k$.
3. For any player $i \in [N]$, if i appears in two bags $B, B' \in \mathcal{T}$, then every bag on the path from B to B' also contains i .

As illustrated in Figure 1, the tree decomposition separates the game into overlapping bags. For example, since B^2 and B^3 only overlap at B^1 , then B^2 and B^3 can be optimized independently, with consistency later enforced at B^1 .

Since any clique in \mathcal{G} is contained in some bag (Diestel, 2025), for every player $i \in [N]$ there exists a bag B with $\mathcal{N}(i) \subseteq B$. We arbitrarily assign each player i to such a bag.

Dynamic programming on the tree. We begin by choosing an arbitrary bag as the root of the tree decomposition and denote it by B^r . For each bag $B \in \mathcal{T}$, let $C(B)$ denote the set of its children. With this setup, we maintain a vector $\mathbf{d}^{(t+1)} \in \mathbb{R}^{\times_{B \in \mathcal{T}} \mathcal{A}_B}$, defined as

$$\begin{aligned} d^{(t+1)}(B, \mathbf{a}_B) &= \sum_{\tau=1}^t \sum_{S \in \mathcal{S}} \frac{1}{|S|} \sum_{\substack{i \in S: \\ i \text{ assigned to } B}} \sum_{\hat{\mathbf{a}}_S \in \mathcal{A}_S} \mu^{(\tau)}(S, \hat{\mathbf{a}}_S) (\mathcal{U}_i((\hat{\mathbf{a}}_{S \cap B}, \mathbf{a}_{B \setminus S})) - \mathcal{U}_i(\mathbf{a}_B)) \\ &+ \sum_{B' \in C(B)} \min_{\substack{\mathbf{a}'_{B'} \in \mathcal{A}_{B'}: \\ \mathbf{a}_{B \cap B'} = \mathbf{a}'_{B \cap B'}}} d^{(t+1)}(B', \mathbf{a}'_{B'}) - n^{(t+1)}(B, \mathbf{a}_B), \end{aligned} \quad (5.4)$$

where $n^{(t+1)}(B, \mathbf{a}_B) \sim \text{Exp}(\eta)^6$ is sampled from an exponential distribution. Therefore, in (5.3), $\tilde{n}^{(t+1)}(\mathbf{a}) = \sum_{B \in \mathcal{T}} n^{(t+1)}(B, \mathbf{a}_B)$. Since each i assigned to B satisfies $\mathcal{N}(i) \subseteq B$, the utility $\mathcal{U}_i(\mathbf{a}_B)$ can be written in terms of \mathbf{a}_B alone. Moreover, the summation $\sum_{\hat{\mathbf{a}}_S \in \mathcal{A}_S}$ can be computed efficiently, since $\mu^{(\tau)}$ is updated via (5.3) and is therefore a pure strategy.

Reconstructing the strategy. The optimal joint action $\mathbf{a}^{(t+1)} \in \mathcal{A}$ is then reconstructed recursively from the root B^r to the leaves:

$$\begin{aligned} \mathbf{a}_{B^r}^{(t+1)} &= \underset{\mathbf{a}_{B^r} \in \mathcal{A}_{B^r}}{\text{argmin}} d^{(t+1)}(B^r, \mathbf{a}_{B^r}) \\ \forall B \in C(B^r), \mathbf{a}_{B \setminus B^r}^{(t+1)} &= \underset{\mathbf{a}_{B \setminus B^r} \in \mathcal{A}_{B \setminus B^r}}{\text{argmin}} d^{(t+1)}\left(B, (\mathbf{a}_{B \setminus B^r}, \mathbf{a}_{B \cap B^r}^{(t+1)})\right). \end{aligned} \quad (5.5)$$

By **Property 1 of Tree Decomposition**, every player's action will be included. Since $\underset{\mathbf{a}_{B \setminus B^r} \in \mathcal{A}_{B \setminus B^r}}{\text{argmin}}$ is taken over $\mathcal{A}_{B \setminus B^r}$, no contradictions arise by **Property 3 of Tree Decomposition**. We then set $\pi^{(t+1)}(\mathbf{a}^{(t+1)}) = 1$.

The regret bound of this procedure is summarized below.

Theorem 5.2. Consider (5.3). For any $\delta > 0$, with probability at least $1 - \delta$, the following holds:

$$\max_{\hat{\pi} \in \Delta^{\mathcal{A}}} \sum_{t=1}^T F\left(\pi^{(t)}, \mu^{(t)}\right) - F\left(\hat{\pi}, \mu^{(t)}\right) \leq 2|\mathcal{T}| \frac{1 + (\text{tw}(\mathcal{G}) + 1) \log A}{\eta} + 2\eta |\mathcal{T}| T + \sqrt{2T \log \frac{1}{\delta}}.$$

The proof is given in Section F. Importantly, Theorem 5.2 shows that by setting $\eta = 1/\sqrt{T}$, we obtain $\mathcal{O}(\sqrt{T})$ regret. Since the update rule for μ mirrors that of π , the detailed analysis is deferred to Section G. We now formally state the regret bound for μ in the following theorem.

Theorem 5.3. Consider the updates in (5.3). For any $\delta > 0$, with probability at least $1 - \delta$, the following holds:

$$\max_{\hat{\mu} \in \Delta^{\times_{S \in \mathcal{S}} \mathcal{A}_S}} \sum_{t=1}^T F\left(\pi^{(t)}, \hat{\mu}\right) - F\left(\pi^{(t)}, \mu^{(t)}\right) \leq 2|\mathcal{T}| \frac{1 + (\text{tw}(\mathcal{G}) + 1) \log A}{\eta} + 2\eta |\mathcal{T}| T + \sqrt{2T \log \frac{1}{\delta}}.$$

The complete proof is provided in Section G.

⁶ $\Pr(x \geq w) = \exp(-\eta w)$ when $x \sim \text{Exp}(\eta)$.

5.3 COMPUTATION OF EQUILIBRIUM

For any $\delta' > 0$, by setting $\delta = \frac{\delta'}{2}$ in Theorem 5.2 and Theorem 5.3, and applying the union bound, we obtain that with probability at least $1 - \delta'$, the following holds:

$$\begin{aligned} & \max_{\hat{\mu} \in \Delta^{\mathcal{S} \times \mathcal{S} \times \mathcal{A}_S}} \sum_{t=1}^T F(\pi^{(t)}, \mu) - \min_{\hat{\pi} \in \Delta^{\mathcal{A}}} \sum_{t=1}^T F(\hat{\pi}, \mu^{(t)}) \\ & \leq 4|\mathcal{T}| \frac{1 + (\text{tw}(\mathcal{G}) + 1) \log A}{\eta} + 4\eta |\mathcal{T}| T + 2\sqrt{2T \log \frac{2}{\delta'}}. \end{aligned} \quad (5.6)$$

We now connect this bound to the convergence of the average strategy profile. Let π^*, μ^* be the solution to (5.1), and define the average strategies $\bar{\pi} := \frac{1}{T} \sum_{t=1}^T \pi^{(t)}$ and $\bar{\mu} := \frac{1}{T} \sum_{t=1}^T \mu^{(t)}$. The left-hand side of (5.6) corresponds to the duality gap: $\max_{\hat{\mu} \in \Delta^{\mathcal{S} \times \mathcal{S} \times \mathcal{A}_S}} F(\bar{\pi}, \hat{\mu}) - \min_{\hat{\pi} \in \Delta^{\mathcal{A}}} F(\hat{\pi}, \bar{\mu})$. Since π^*, μ^* are optimal solutions to (5.6), they satisfy

$$\min_{\hat{\pi} \in \Delta^{\mathcal{A}}} F(\hat{\pi}, \bar{\mu}) \leq F(\pi^*, \mu^*) \leq \max_{\hat{\mu} \in \Delta^{\mathcal{S} \times \mathcal{S} \times \mathcal{A}_S}} F(\bar{\pi}, \hat{\mu}).$$

Combining these pieces, we arrive at the following finite-time convergence guarantee:

Theorem 5.4. Let π^*, μ^* be the solution of (5.1), and define $\bar{\pi} := \frac{1}{T} \sum_{t=1}^T \pi^{(t)}$, $\bar{\mu} := \frac{1}{T} \sum_{t=1}^T \mu^{(t)}$. Then, for any $\delta > 0$, with probability at least $1 - \delta$, we have

$$\max_{\hat{\mu} \in \Delta^{\mathcal{S} \times \mathcal{S} \times \mathcal{A}_S}} F(\bar{\pi}, \hat{\mu}) \leq F(\pi^*, \mu^*) + 4|\mathcal{T}| \frac{1 + (\text{tw}(\mathcal{G}) + 1) \log A}{\eta T} + 4\eta |\mathcal{T}| + 2\sqrt{\frac{2 \log \frac{2}{\delta}}{T}}.$$

With $\eta = \frac{1}{\sqrt{T}}$, the average strategy $\bar{\pi}$ constitutes an $\mathcal{O}\left(\frac{|\mathcal{T}| \cdot \text{tw}(\mathcal{G}) \log A + \sqrt{\log \frac{2}{\delta}}}{\sqrt{T}}\right)$ -MASE. The overall running time is $\mathcal{O}(T \cdot |\mathcal{S}| \cdot |\mathcal{T}| \cdot A^{\text{tw}(\mathcal{G})+1})$. Hence, the exponential dependence aligns with the lower bound in Theorem 4.3.

6 EXPERIMENTS

In this section, we compare our algorithm against several baselines: Follow the Regularized Leader with a Euclidean regularizer (**FTRL**), **Hedge**, Follow the Perturbed Leader with an exponential noise distribution (**FTPL**; all players run FTPL independently), and Online Mirror Descent with a Euclidean regularizer (**OMD**) (Hazan et al., 2016). We also plot the ground-truth MASE computed via linear programming (**LP**) in Section B.

We evaluate the algorithms on three criteria:

- **Exploitability.** ($\max_{i \in [N]} \max_{\hat{a}_i \in \mathcal{A}_i} \mathbb{E}_{\mathbf{a} \sim \pi} [\mathcal{U}_i(\hat{a}_i, \mathbf{a}_{-i}) - \mathcal{U}_i(\mathbf{a})]$): the maximum gain a single player can obtain by deviating unilaterally. Exploitability ≤ 0 indicate a Nash equilibrium (or a correlated equilibrium if π is correlated).
- **Coalition exploitability.** ($\max_{\mu \in \Delta^{\mathcal{S} \times \mathcal{S} \times \mathcal{A}_S}} F(\pi, \mu)$): the maximum average gain when a coalition deviates simultaneously. We take \mathcal{S} to be the set of all non-empty player subsets.
- **Social welfare.** ($\sum_{i=1}^N \mathbb{E}_{\mathbf{a} \sim \pi} [\mathcal{U}_i(\mathbf{a})]$): the sum of all players' utilities.

Utility definitions and additional details are provided in Section H. In the Prisoner's Dilemma (Luce & Raiffa, 1957), the MASE corresponds to players choosing (C, D) and (D, C) with probability 0.5 each, yielding a social welfare of 1.0. In contrast, because the unique NE/CCE in this game is (D, D) , the baselines converge to that outcome, with a lower social welfare of 0.4. Thus, in the Prisoner's Dilemma, MASE promotes cooperation and achieves higher utility.

In the Stag Hunt, there are two Nash equilibria, one of which attains higher utility. As shown in Figure 2, all baselines converge to the worse equilibrium, whereas MASE converges to the better one. Finally, in terms of exploitability (unilateral deviations), MASE remains close to the baselines, while the baselines are substantially more fragile to multilateral deviations.

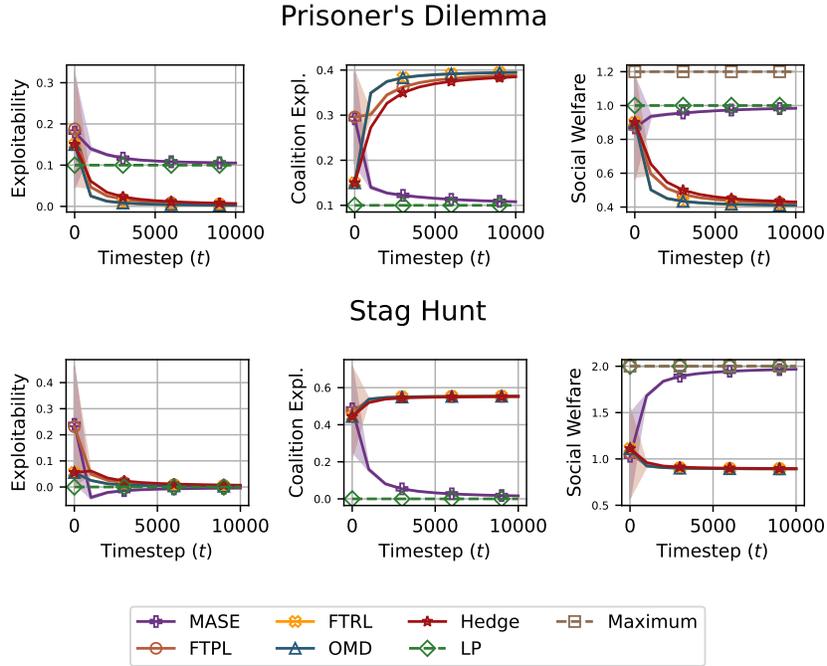


Figure 2: **LP** denotes the linear programming solution from Section B, and **Maximum** denotes the maximum achievable social welfare. The baselines are comparatively fragile to multilateral deviations, while MASE is more robust and achieves higher social welfare. At the same time, MASE’s exploitability is close to that of the baselines.

6.1 TRADE-OFF BETWEEN EXPLOITABILITY AND SOCIAL WELFARE

In Figure 2, we can see that by allowing exploitability to increase from 0.0 to 0.1, the social welfare of MASE increases from 0.4 to 1.0. This raises a natural question:

Given a tolerance $\epsilon \geq 0$, what is the maximum social welfare achievable by an equilibrium with exploitability at most ϵ ?

In other words, if we are willing to sacrifice equilibrium robustness, how much can we improve social welfare? Interestingly, this trade-off can be computed efficiently using a variant of our MASE framework. Specifically, we solve the following weighted objective:

$$\operatorname{argmin}_{\pi \in \Delta^{\mathcal{A}}} \max_{S \in \mathcal{S}} \max_{\hat{\mathbf{a}}_S \in \mathcal{A}_S} \frac{w_S}{|S|} \sum_{i \in S} \mathbb{E}_{\mathbf{a} \sim \pi} [\mathcal{U}_i(\hat{\mathbf{a}}_S, \mathbf{a}_{-S}) - \mathcal{U}_i(\mathbf{a})], \quad (6.1)$$

where $\mathbf{w} \in \mathbb{R}^{\mathcal{S}}$ is a vector of non-negative weights. We have the following lemma.

Lemma 6.1. For any $\epsilon > 0$, computing the CCE with exploitability no more than ϵ that maximizes social welfare is equivalent to (6.1) by setting $\mathcal{S} = \{\{i\}\}_{i \in [N]} \cup \{[N]\}$ and using the weights:

$$w_S = \begin{cases} w & \text{if } |S| = 1 \\ 1 - w & \text{if } S = [N] \end{cases}$$

for some $w \in [0, 1)$. Conversely, solving (6.1) with these parameters corresponds to finding a point on the Pareto frontier of social welfare and exploitability.

The proof is postponed to Section H.4. With Lemma 6.1, we can compute the Pareto frontier by solving (6.1) for different values of w . The results are shown in Figure 3. In the Stag Hunt, since one of the Nash equilibria already maximizes social welfare, the social welfare remains fixed at its optimal value for all $w \in [0, 1)$.

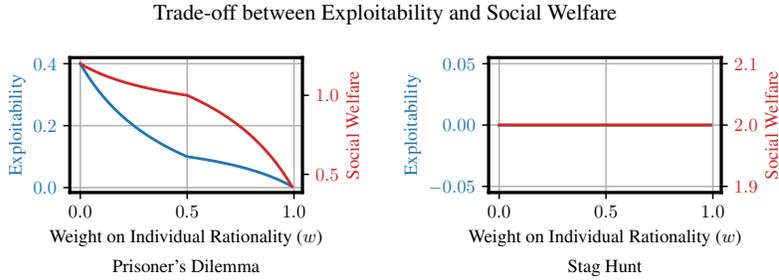


Figure 3: The trade-off between exploitability and social welfare in the Prisoner’s Dilemma and the Stag Hunt.

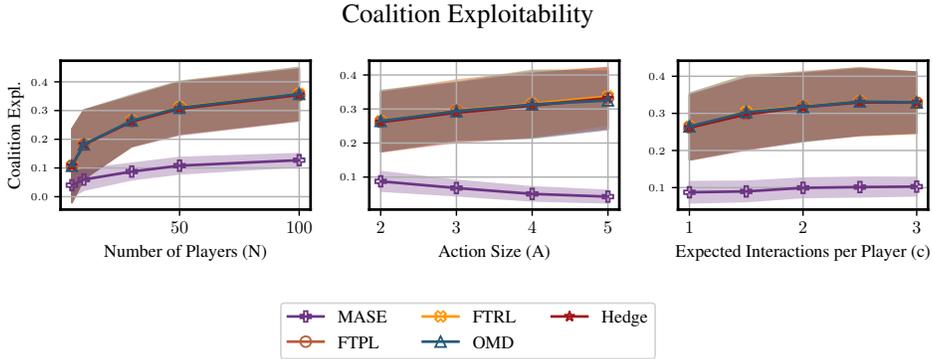


Figure 4: The coalition exploitability of random polymatrix games of different sizes when coalitions with no more than two players are considered. A larger expected number of interactions per player (c) generally corresponds to a larger treewidth of the Utility Dependency Graph.

6.2 COALITION EXPLOITABILITY IN LARGER GAMES

As shown in Figure 4, the coalition exploitability of the average strategy generated by classical no-regret learning algorithms increases as the game size grows. Note that we only consider coalitions of size no more than two. This trend underscores the importance of minimizing coalition exploitability. As games become larger, the equilibria to which these algorithms converge become increasingly fragile to coalition deviations, necessitating approaches that explicitly account for such multilateral deviations. Further details are provided in Section I.

7 CONCLUSION

In this work, we introduced the Minimum Average-Strong Equilibrium (MASE), a tractable solution concept that accounts for multilateral deviations by minimizing each coalition’s average incentive to deviate. We established that computing an approximate MASE is NP-hard even with singleton coalitions and proved a fixed-parameter lower bound showing unavoidable exponential dependence on the treewidth of the Utility Dependency Graph. We then designed an algorithm—combining a correlator–deviator meta-game with FTPL updates and dynamic programming over a tree decomposition, whose running time matches this lower bound up to the treewidth factor. Empirically, MASE is substantially more robust to coalition deviations than standard baselines while improving social welfare in canonical games, all without materially worsening unilateral exploitability.

In the future, it is natural to move beyond uniform averaging within coalitions. A compelling open direction is to characterize lower and upper bounds for objectives that minimize the minimal incentive within each coalition. More broadly, extending these ideas to richer coalition objectives would mature the strong-equilibrium framework and yield more solution concepts that go beyond unilateral deviations.

ACKNOWLEDGEMENT

The authors are grateful to the anonymous reviewers and the area chair for their constructive feedback. M.L. was supported by the MathWorks Fellowship. A.O. and G.F. were supported in part by the ONR grant N000142512296. G.F. was additionally supported by CCF-2443068 and an AI2050 Early Career Fellowship.

8 ETHICS STATEMENT

This paper presents work that aims to advance the field of game theory. There are many potential societal consequences of our work, none of which we feel must be specifically highlighted here.

9 REPRODUCIBILITY STATEMENT

The proof and assumptions are stated in Sections C, D, F and G.2.

10 USE OF LARGE LANGUAGE MODELS

In this paper, we use large language models (LLMs). For example, by correcting grammatical errors, identifying related work to reduce the risk of overlooking relevant papers, helping create illustrative figures (Figure 1), and assisting with coding.

REFERENCES

- Naman Agarwal, Alon Gonen, and Elad Hazan. Learning in non-convex games with an optimization oracle. In *Conference on Learning Theory (COLT)*, 2019.
- Ioannis Anagnostides, Gabriele Farina, Tuomas Sandholm, and Brian Hu Zhang. A polynomial-time algorithm for variational inequalities under the Minty condition. *arXiv preprint arXiv:2504.03432*, 2025.
- Sanjeev Arora and Boaz Barak. *Computational complexity: a modern approach*. Cambridge University Press, 2009.
- Robert Aumann. *Acceptable points in General Cooperative n-person Games*, pp. 287–324. Princeton University Press, 01 1959. doi: 10.1090/fic/023/01.
- Robert J Aumann. Subjectivity and correlation in randomized strategies. *Journal of mathematical Economics*, 1(1):67–96, 1974.
- Carl T Bergstrom and Peter Godfrey-Smith. On the evolution of behavioral heterogeneity in individuals and populations. *Biology and Philosophy*, 13(2):205–231, 1998.
- B Douglas Bernheim, Bezalel Peleg, and Michael D Whinston. Coalition-proof Nash equilibria I. concepts. *Journal of economic theory*, 42(1):1–12, 1987.
- Marie Louisa Tølbøll Berthelsen and Kristoffer Arnsfelt Hansen. On the computational complexity of decision problems about multi-player Nash equilibria. *Theory of Computing Systems*, 66(3): 519–545, 2022.
- Nicolo Cesa-Bianchi and Gábor Lugosi. *Prediction, learning, and games*. Cambridge university press, 2006.
- Xi Chen and Xiaotie Deng. Settling the complexity of two-player Nash equilibrium. In *Symposium on Foundations of Computer Science (FOCS)*, 2006.
- Vincent Conitzer and Tuomas Sandholm. New complexity results about Nash equilibria. *Games and Economic Behavior*, 63(2):621–641, 2008.

- Constantinos Daskalakis, Paul W Goldberg, and Christos H Papadimitriou. The complexity of computing a Nash equilibrium. *Communications of the ACM*, 52(2):89–97, 2009.
- Reinhard Diestel. *Graph theory*, volume 173. Springer Nature, 2025.
- B Curtis Eaves. Polymatrix games with joint constraints. *SIAM Journal on Applied Mathematics*, 24(3):418–423, 1973.
- Barış Can Esmer, Jacob Focke, Dániel Marx, and Paweł Rzażewski. Fundamental problems on bounded-treewidth graphs: The real source of hardness. In *International Colloquium on Automata, Languages and Programming (ICALP)*, 2024.
- Nicola Gatti, Marco Rocco, and Tuomas Sandholm. On the verification and computation of strong Nash equilibrium. In *International Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*, 2013.
- Sergiu Hart and David Schmeidler. Existence of correlated equilibria. *Mathematics of Operations Research*, 14(1):18–25, 1989.
- Elad Hazan et al. Introduction to online convex optimization. *Foundations and Trends® in Optimization*, 2(3-4):157–325, 2016.
- Martin Hoefer. Strategic cooperation in cost sharing games. *International Journal of Game Theory*, 42(1):29–53, 2013.
- Ron Holzman and Nissan Law-Yone. Strong equilibrium in congestion games. *Games and economic behavior*, 21(1-2):85–101, 1997.
- Joseph T Howson Jr. Equilibria of polymatrix games. *Management Science*, 18(5-part-1):312–318, 1972.
- Russell Impagliazzo and Ramamohan Paturi. On the complexity of k-SAT. *Journal of Computer and System Sciences*, 62(2):367–375, 2001.
- Sham Kakade, Michael Kearns, John Langford, and Luis Ortiz. Correlated equilibria in graphical games. In *ACM Conference on Electronic Commerce*, 2003.
- Yasushi Kawase and Hanna Sumita. On the max-min fair stochastic allocation of indivisible goods. In *AAAI Conference on Artificial Intelligence (AAAI)*, 2020.
- Michael J. Kearns, Michael L. Littman, and Satinder Singh. Graphical models for game theory. In *Conference on Uncertainty in Artificial Intelligence (UAI)*, 2001.
- Daniel Lokshtanov, Dániel Marx, and Saket Saurabh. Known algorithms on graphs of bounded treewidth are probably optimal. In *ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2011.
- Duncan R.. Luce and Howard Raiffa. *Games and Decisions*. wiley, 1957.
- Hervé Moulin and J-P Vial. Strategically zero-sum games: the class of games whose completely mixed equilibria cannot be improved upon. *International Journal of Game Theory*, 7(3-4):201–221, 1978.
- John F Nash. Equilibrium points in n -person games. *Proceedings of the national academy of sciences*, 36(1):48–49, 1950.
- Rabia Nessah and Guoqiang Tian. On the existence of strong Nash equilibria. *Journal of Mathematical Analysis and Applications*, 414(2):871–885, 2014.
- Christos H Papadimitriou and Tim Roughgarden. Computing correlated equilibria in multi-player games. *Journal of the ACM (JACM)*, 55(3):1–29, 2008.
- Arthur Pigou. *The economics of welfare*. Routledge, 2017.
- Mona Rahn and Guido Schäfer. Efficient equilibria in polymatrix coordination games. In *International Symposium on Mathematical Foundations of Computer Science (MFCS)*, 2015.

Robert W Rosenthal. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory*, 2(1):65–67, 1973.

Ola Rozenfeld and Moshe Tennenholtz. Strong and correlated strong equilibria in monotone congestion games. In *International Workshop on Internet and Network Economics*, pp. 74–86. Springer, 2006.

Arun Sai Suggala and Praneeth Netrapalli. Online non-convex learning: Following the perturbed leader is optimal. In *Algorithmic Learning Theory (ALT)*, 2020.

Heinrich Von Stackelberg. *Market structure and equilibrium*. Springer Science & Business Media, 2010.

A RELATED WORK

In this section, we review the literature on strong equilibrium from three perspectives: existence, time complexity, and computation.

Existence of Strong Equilibrium. Aumann (1959) introduced the strong NE, where no coalition (a nonempty subset of players) can deviate in a way that strictly improves the utility of all its members. However, even in simple two-player games such as the Prisoner’s Dilemma (Luce & Raiffa, 1957), a strong NE does not exist when players can coordinate. To address this, Bernheim et al. (1987) proposed the coalition-proof equilibrium, which restricts the set of deviations. Yet, this concept also fails to guarantee existence, already in three-player games (Bernheim et al., 1987). More recently, Rahn & Schäfer (2015) studied the notion of α -approximate k -equilibrium, where no coalition of size at most k can deviate to improve each member’s utility by at least a factor of $\alpha \geq 1$. They further showed that such equilibria exist in graph coordination games only under specific conditions, for instance, when $\alpha \geq 2$. Motivated by these non-existence results, we instead focus on minimizing the maximum average gain from coalition deviations (MASE), a quantity that is always well defined.

Complexity of Strong Equilibrium. Since a strong NE degenerates to an NE when only singleton coalitions are considered, computing a strong NE is PPAD-hard in general (Daskalakis et al., 2009; Chen & Deng, 2006). Beyond computation, Conitzer & Sandholm (2008) showed that even deciding whether a strong NE exists is NP-complete in two-player symmetric games, and Berthelsen & Hansen (2022) further established that the problem is $\exists\mathbb{R}$ -complete for three-player games. Similarly, Rahn & Schäfer (2015) proved that determining the existence of a strong NE is NP-complete in graph coordination games, even when restricting attention to coalitions of constant size. To the best of our knowledge, however, no hardness results are known for computing strong equilibria when correlation on the joint strategy is allowed, *i.e.*, a correlated strategy immune to coalition deviations. In this paper, we show that computing a correlated strategy that minimizes the average gain from coalition deviations is NP-hard. Moreover, we establish a fixed-parameter lower bound based on the treewidth of the Utility Dependency Graph, demonstrating an inherent computational barrier in solving the MASE considered here.

Computation of Strong Equilibrium. Holzman & Law-Yone (1997) and Rozenfeld & Tennenholtz (2006) developed algorithms to compute strong NE and correlated strong equilibria in congestion games under certain conditions in polynomial time. Rahn & Schäfer (2015) showed that a strong NE can also be computed in polynomial time when the graph coordination game is defined on a tree. In contrast, Gatti et al. (2013) proposed a spatial branch-and-bound algorithm for computing strong NE more generally, but its runtime is exponential. Along the same lines, Nessah & Tian (2014) also provided a computationally intractable algorithm. Of independent interest, Papadimitriou & Roughgarden (2008) introduced an efficient algorithm for computing optimal CE—*e.g.*, a CE that maximizes the social welfare—in graphical games (Kearns et al., 2001; Kakade et al., 2003) with bounded treewidth, using linear programming. In this paper, we develop a new algorithm for computing MASE based on no-regret learning, with a time complexity that matches the lower bound dictated by the treewidth of the Utility Dependency Graph.

B LINEAR PROGRAMMING FOR SOLVING (MASE)

(MASE) can be solved by the following linear programming

$$\begin{aligned} & \min_{\pi, w} w \\ & w \geq \frac{1}{|S|} \sum_{i \in S} \mathbb{E}_{\mathbf{a} \sim \pi} [\mathcal{U}_i(\hat{\mathbf{a}}_S, \mathbf{a}_{-S}) - \mathcal{U}_i(\mathbf{a})] \quad \forall S \in \mathcal{S}, \hat{\mathbf{a}}_S \in \mathcal{A}_S \\ & \pi(\mathbf{a}) \geq 0 \quad \forall \mathbf{a} \in \mathcal{A} \\ & \sum_{\mathbf{a} \in \mathcal{A}} \pi(\mathbf{a}) = 1. \end{aligned}$$

Since

$$\begin{aligned}\mathbb{E}_{\mathbf{a} \sim \pi} [\mathcal{U}_i(\hat{\mathbf{a}}_S, \mathbf{a}_{-S})] &= \sum_{\mathbf{a} \in \mathcal{A}} \pi(\mathbf{a}) \mathcal{U}_i(\hat{\mathbf{a}}_S, \mathbf{a}_{-S}) \\ \mathbb{E}_{\mathbf{a} \sim \pi} [\mathcal{U}_i(\mathbf{a})] &= \sum_{\mathbf{a} \in \mathcal{A}} \pi(\mathbf{a}) \mathcal{U}_i(\mathbf{a})\end{aligned}$$

are linear in π , the linear programming above is valid. Note that the linear program contains exponentially many variables ($\pi \in \Delta^{\mathcal{A}}$), so its complexity is necessarily exponential in N .

C OMITTED PROOFS IN SECTION 4

This section presents the omitted proofs in Section 4.

C.1 PROOF OF THEOREM 4.1

Theorem 4.1. Computing ϵ -MASE is NP-hard, even when \mathcal{S} only contains singletons (coalitions of size one) and $1/\epsilon$ is polynomial in the number of players.

Proof. We will introduce the allocation problem (NP-hard) and show that it can be reduced to computing the correlated strong equilibrium.

Definition C.1 (Allocation Problem). There are n agents and m goods. An assignment $X: [m] \rightarrow [n]$ is a mapping from each good to an agent. Agent i 's utility is $u_i(X)$ for an assignment X . A stochastic allocation $\mathbf{p} \in \Delta^{[m]^{[n]}}$ is a distribution over all possible assignments. The egalitarian social welfare (ESW) maximization is defined as

$$\max_{\mathbf{p} \in \Delta^{[m]^{[n]}}} \min_{i \in [n]} \sum_{X \in [m]^{[n]}} p(X) u_i(X). \quad (\text{C.1})$$

For any allocation problem, we can create a game with $N = n + m$ players. The action set of player $i \leq n$ is $\mathcal{A}_i = \{0, 1\}$, while the action set of player $j > n$ is $\mathcal{A}_j = [n]$. For any joint action $\mathbf{a} \in \mathcal{A}$, $\mathcal{U}_i(\mathbf{a}) = u_i(\mathbf{a}_{-[n]})$ for $i \leq n$ when $a_1 = a_2 = \dots = a_n$, otherwise $\mathcal{U}_i(\mathbf{a}) = -u_i(\mathbf{a}_{-[n]})$. For $j > n$, $\mathcal{U}_j(\mathbf{a}) = 0$. Moreover, let $\mathcal{S} = \{\{1\}, \{2\}, \dots, \{n\}\}$. We further define

$$\begin{aligned}a &:= \sum_{\substack{\mathbf{a} \in \mathcal{A}: \\ a_2 = a_3 = \dots = a_n = 0}} \pi(\mathbf{a}) u_i(\mathbf{a}_{-[n]}) \\ b &:= \sum_{\substack{\mathbf{a} \in \mathcal{A}: \\ a_2 = a_3 = \dots = a_n = 1}} \pi(\mathbf{a}) u_i(\mathbf{a}_{-[n]}) \\ c &:= - \sum_{\substack{\mathbf{a} \in \mathcal{A}: \\ \exists i, j \in [n] \setminus \{1\}, a_i \neq a_j}} \pi(\mathbf{a}) u_i(\mathbf{a}_{-[n]}).\end{aligned}$$

Then, the gap of player 1 is lower bounded by

$$\max_{\hat{a}_1 \in \{0, 1\}} \mathbb{E}_{\mathbf{a} \sim \pi} [\mathcal{U}_i(\hat{\mathbf{a}}_1, \mathbf{a}_{-1}) - \mathcal{U}_i(\mathbf{a})] \geq \max(a, b) + c - (a + b + c) = -\min(a, b).$$

The equation holds when a_1 is always equal to a_2 when $a_2 = a_3 = \dots = a_n$. Therefore, the optimal strategy π should satisfy $\pi_{[n]}((0, \dots, 0)) = \pi_{[n]}((1, \dots, 1)) = \frac{1}{2}$.

Then, (MASE) is equivalent to

$$\min_{\pi_{[m]} \in \Delta^{[m]^{[n]}}} \max_{i \in [n]} \sum_{\mathbf{a} \in \mathcal{A}} -\pi(\mathbf{a}) u_i(\mathbf{a}_{-[n]}) = \min_{\pi_{[m]} \in \Delta^{[m]^{[n]}}} \max_{i \in [n]} \sum_{\mathbf{a}_{-[n]} \in [m]^{[n]}} -\pi_{-[n]}(\mathbf{a}_{-[n]}) u_i(\mathbf{a}_{-[n]}),$$

which is equivalent to (C.1). Finally, according to Kawase & Sumita (2020, Corollary 1), it is NP-hard to approximate (C.1) up to $1 - \frac{1}{e}$.

The hard instance constructed in Kawase & Sumita (2020) satisfies that $\max_i \max_{X \in [m]^{[n]}} u_i(X) = \text{Poly}(n, m)$ and the solution to (C.1) is 1. Therefore, $\text{Poly}(N, \frac{1}{\epsilon})$ algorithm does not exist unless P=NP for solving an ϵ -MASE. \square

C.2 PROOF OF THEOREM 4.3

Proof. According to Lokshtanov et al. (2011), under SETH, q -coloring cannot be solved in $\mathcal{O}((q - \zeta)^{tw} \cdot \text{Poly}(|I|))$ for arbitrary graph G , when a tree decomposition of width tw is given.⁷ In the sequel, we construct a game such that computing $\frac{1}{N}$ -approximate (MASE) is equivalent to determining the q -coloring.

For any q -coloring problem on $G = (V, E)$, we will construct a game with $N = |V| + |E|$ players. For each player $i \leq |V|$, the action set $\mathcal{A}_i = \{1, 2, \dots, q\}$ and the utility function $\mathcal{U}_i \equiv 0$ is a constant function equal to zero. For each player $j > |V|$, the action set is $\{1\}$ and the utility function is $\mathcal{U}_j(\mathbf{a}) = \mathbb{1}(a_{e_{j-|V|,1}} \neq a_{e_{j-|V|,2}})$, where $(e_{j-|V|,1}, e_{j-|V|,2})$ is the $(j - |V|)^{th}$ edge in E and $\mathbb{1}$ is the indicator function (equals one when the argument is true and otherwise zero). In this game, $\mathcal{S} = \{\{i_1, i_2, j + |V|\} \mid e_j = (i_1, i_2)\}$.

Firstly, for any proper coloring $\mathbf{c} \in [q]^{|V|}$, the associated pure strategy is π^c , where $\pi_i^c(a_i) = 1$ if and only if $a_i = c_i$ and 0 otherwise. It satisfies (MASE). Because for any $S \ni S = (i_1, i_2, j + |V|)$, the maximum of $\frac{1}{|S|} \sum_{i \in S} \mathcal{U}_i(\pi)$ is $\frac{1}{3}$, which is attained when the colors of node i_1, i_2 are different. Therefore, π^c obtains the maximum for every coalition $S \in \mathcal{S}$, which implies the satisfaction of (MASE).

Secondly, for any joint strategy $\pi \in \Delta^{\mathcal{A}}$ satisfying (MASE), we have

$$\begin{aligned} \min_{S \in \mathcal{S}} \frac{\sum_{i \in S} \mathcal{U}_i(\pi)}{|S|} &\leq \frac{1}{|\mathcal{S}|} \sum_{S \in \mathcal{S}} \frac{\sum_{i \in S} \mathcal{U}_i(\pi)}{|S|} = \sum_{\mathbf{a} \in \mathcal{A}} \pi(\mathbf{a}) \frac{1}{|\mathcal{S}|} \sum_{S \in \mathcal{S}} \frac{\sum_{i \in S} \mathcal{U}_i(\mathbf{a})}{|S|} \\ &\leq \frac{1}{3} - \frac{1}{3|E|} \sum_{\mathbf{a} \in \mathcal{A}} \pi(\mathbf{a}) \mathbb{1}(\mathbf{a}_{[|V|]} \text{ is not a proper coloring}). \end{aligned}$$

Because there must exist at least one edge with both of its nodes in the same color for any improper coloring. On the other hand, let $\hat{\pi} = \pi^c$ for some proper coloring \mathbf{c} . Then, for any $S \in \mathcal{S}$, we have

$$\frac{\sum_{i \in S} \mathcal{U}_i(\hat{\pi})}{|S|} = \frac{1}{3}.$$

Therefore, the approximation error of (MASE) is at least

$$\frac{1}{3|E|} \sum_{\mathbf{a} \in \mathcal{A}} \pi(\mathbf{a}) \mathbb{1}(\mathbf{a}_{[|V|]} \text{ is not a proper coloring})$$

for any joint strategy π . When we get $\frac{1}{9N^2}$ approximation, since $|E| \leq N^2$, we have

$$\sum_{\mathbf{a} \in \mathcal{A}} \pi(\mathbf{a}) \mathbb{1}(\mathbf{a}_{[|V|]} \text{ is not a proper coloring}) \leq \frac{1}{3}.$$

Therefore, when sampling $\mathbf{a} \sim \pi$, we will get a proper coloring with probability at least $\frac{2}{3}$, which is in complexity class RP. As a result, when $\mathbb{P}=\text{RP}$, the time complexity of computing (MASE) is at least $O^*(A^{tw(G)})$, where G is the trust graph and A is the size of the maximal action set. \square

D PROOF OF THEOREM 5.1

Theorem 5.1 (Efficient Representation). For any $\epsilon \geq 0$, at least one of the ϵ -MASE can be represented as a linear combination of $\sum_{S \in \mathcal{S}} |S| \cdot A^{tw(G)} + 1$ pure strategies, where $tw(\mathcal{G})$ is the treewidth of Utility Dependency Graph.

⁷As summarized in Esmer et al. (2024), the proof of the q -coloring complexity implicitly implies that the complexity is lower bounded by $\mathcal{O}(q^{tw} \cdot \text{Poly}(|I|))$, even though a tree decomposition of width tw is given. In other words, aside from computing a tree decomposition, the q -coloring itself has an intrinsic computational barrier.

Proof. Let $D := \sum_{S \in \mathcal{S}} \sum_{i \in S} |\mathcal{A}_{S \cap \mathcal{N}(i)}|$. For any joint strategy $\pi \in \Delta^{\mathcal{A}}$, consider the vector $\mathbf{v}^\pi \in \mathbb{R}^D$, where

$$v^\pi(S, i, \hat{\mathbf{a}}_{S \cap \mathcal{N}(i)}) = \mathbb{E}_{\mathbf{a} \sim \pi} [\mathcal{U}_i(\hat{\mathbf{a}}_{S \cap \mathcal{N}(i)}, \mathbf{a}_{-(S \cap \mathcal{N}(i))}) - \mathcal{U}_i(\mathbf{a})],$$

for any $S \in \mathcal{S}$ and $\hat{\mathbf{a}}_{S \cap \mathcal{N}(i)} \in \mathcal{A}_{S \cap \mathcal{N}(i)}$. By definition of \mathbf{v}^π , it is linear in π . Therefore, the vertex set of $\{\mathbf{v}^\pi \mid \pi \in \Delta^{\mathcal{A}}\}$ should correspond to a subset of $\Delta^{\mathcal{A}}$'s vertex set, which is the set of all pure strategies.

By Carathéodory's theorem, for any $\pi^* \in \Delta^{\mathcal{A}}$, \mathbf{v}^{π^*} can be represented as the linear combination of $D + 1$ vertices, which further implies it can be written as a linear combination of $D + 1$ vectors in $\{\mathbf{v}^\pi \mid \pi \text{ is a pure strategy}\}$. Then, when $\mathbf{v}^{\pi^*} = \sum_{k=1}^{D+1} \lambda^k \mathbf{v}^{\pi^k}$ with $\boldsymbol{\lambda} \in \Delta^{D+1}$ and π^1, π^2, \dots are pure strategies, due to the linearity of \mathbf{v}^π , we have

$$\mathbf{v}^{\pi^*} = \sum_{k=1}^{D+1} \lambda^k \mathbf{v}^{\pi^k} = \mathbf{v}^{\sum_{k=1}^{D+1} \lambda^k \pi^k}.$$

Finally,

$$\begin{aligned} & \max_{S \in \mathcal{S}} \max_{\hat{\mathbf{a}}_S \in \mathcal{A}_S} \frac{1}{|S|} \sum_{i \in S} \mathbb{E}_{\mathbf{a} \sim \pi} [\mathcal{U}_i(\hat{\mathbf{a}}_S, \mathbf{a}_{-S}) - \mathcal{U}_i(\mathbf{a})] \\ &= \max_{S \in \mathcal{S}} \max_{\hat{\mathbf{a}}_S \in \mathcal{A}_S} \frac{1}{|S|} \sum_{i \in S} \mathbb{E}_{\mathbf{a} \sim \pi} [\mathcal{U}_i(\hat{\mathbf{a}}_{S \cap \mathcal{N}(i)}, \mathbf{a}_{-(S \cap \mathcal{N}(i))}) - \mathcal{U}_i(\mathbf{a})] \\ &= \max_{S \in \mathcal{S}} \max_{\hat{\mathbf{a}}_S \in \mathcal{A}_S} \frac{1}{|S|} \sum_{i \in S} v^{\pi^*}(S, i, \hat{\mathbf{a}}_{S \cap \mathcal{N}(i)}) \\ &= \max_{S \in \mathcal{S}} \max_{\hat{\mathbf{a}}_S \in \mathcal{A}_S} \frac{1}{|S|} \sum_{i \in S} v^{\sum_{k=1}^{D+1} \lambda^k \pi^k}(S, i, \hat{\mathbf{a}}_{S \cap \mathcal{N}(i)}). \end{aligned}$$

Hence, once π^* satisfies (MASE), there exists a linear combination of $D + 1$ pure strategies also satisfying (MASE). Given $\mathcal{N}(i) \leq \text{tw}(\mathcal{G})$, we have $D \leq \sum_{S \in \mathcal{S}} |S| \cdot A^{\text{tw}(\mathcal{G})}$. \square

E THE OPTIMALITY OF DYNAMIC PROGRAMMING ON TREE DECOMPOSITION

This section shows that the dynamic programming (e.g., (5.4)) will compute the optimality. Let $u_i(\mathbf{a}_{\mathcal{N}(i)})$ be the contribution of player i 's utility to the final objective. Recall that $\mathcal{N}_i^S := \mathcal{N}(i) \cap S$. Then, in (5.4),

$$u_i(\mathbf{a}_{\mathcal{N}(i)}) = - \sum_{\tau=1}^t \sum_{S \in \mathcal{S}} \frac{1}{|S|} \sum_{i \in S} \sum_{\hat{\mathbf{a}}_{\mathcal{N}_i^S} \in \mathcal{A}_{\mathcal{N}_i^S}} \mu^{(\tau)}(S, \hat{\mathbf{a}}_{\mathcal{N}_i^S}) \left(\mathcal{U}_i(\hat{\mathbf{a}}_{\mathcal{N}_i^S}, \mathbf{a}_{\mathcal{N}(i) \setminus S}) - \mathcal{U}_i(\mathbf{a}_{\mathcal{N}(i)}) \right)$$

at timestep t . We consider the following update rule in this section, which generalizes (5.4) and (G.1) (see the proof of Lemma F.3 and Lemma G.1),

$$h(B, \mathbf{a}_B) = \sum_{\substack{i \in [N]: \\ i \text{ assigned to } B}} u_i(\mathbf{a}_{\mathcal{N}(i)}) + \sum_{B' \in \mathcal{C}(B)} \max_{\substack{\mathbf{a}'_{B'} \in \mathcal{A}_{B'}: \\ \mathbf{a}_{B \cap B'} = \mathbf{a}'_{B \cap B'}}} h(B', \mathbf{a}'_{B'}). \quad (\text{E.1})$$

In the following, we will show that (E.1) is optimal.

Lemma E.1. For any bag $B \in \mathcal{T}$, let $\text{st}(B) := \{i\}_{i \text{ assigned to } B} \cup \bigcup_{B' \in \mathcal{C}(B)} \text{st}(B')$ be the set of players assigned to B and bags in its subtree. Then, for any bag $B \in \mathcal{T}$ and $\mathbf{a}_B \in \mathcal{A}_B$, we have

$$h(B, \mathbf{a}_B) = \max_{\mathbf{a}_{-B} \in \mathcal{A}_{-B}} \sum_{i \in \text{st}(B)} u_i(\mathbf{a}_{\mathcal{N}(i)}). \quad (\text{E.2})$$

The proof is postponed to the end of this section. Note that for the root bag B^r , since $\text{st}(B^r) = [N]$, Lemma E.1 implies that $\max_{\mathbf{a}_{B^r} \in \mathcal{A}_{B^r}} h(B^r, \mathbf{a}_{B^r}) = \max_{\mathbf{a} \in \mathcal{A}} \sum_{i=1}^N u_i(\mathbf{a}_{\mathcal{N}(i)})$. Therefore, we find the maximum of $\sum_{i=1}^N u_i(\mathbf{a}_{\mathcal{N}(i)})$, and the optimal joint action $\mathbf{a} \in \mathcal{A}$ can be extracted recursively.

Specifically, let

$$\begin{aligned} \mathbf{a}_{B^r}^* &= \operatorname{argmax}_{\mathbf{a}_{B^r} \in \mathcal{A}_{B^r}} h(B^r, \mathbf{a}_{B^r}) \\ \forall B \in C(B^r), \mathbf{a}_{B \setminus B^r}^* &= \operatorname{argmax}_{\mathbf{a}_{B \setminus B^r} \in \mathcal{A}_{B \setminus B^r}} d\left(B, (\mathbf{a}_{B \setminus B^r}, \mathbf{a}_{B \cap B^r}^{(t+1)})\right). \end{aligned} \quad (\text{E.3})$$

We will do this recursively until we find the whole $\mathbf{a}^* \in \mathcal{A}$. The tie-breaking rule can be arbitrary, and we use the lexicographic order of joint actions for simplicity. Hence, we prove the optimality of the update rule (E.1). \square

Lemma E.1. For any bag $B \in \mathcal{T}$, let $\text{st}(B) := \{i\}_{i \text{ assigned to } B} \cup \bigcup_{B' \in C(B)} \text{st}(B')$ be the set of players assigned to B and bags in its subtree. Then, for any bag $B \in \mathcal{T}$ and $\mathbf{a}_B \in \mathcal{A}_B$, we have

$$h(B, \mathbf{a}_B) = \max_{\mathbf{a}_{-B} \in \mathcal{A}_{-B}} \sum_{i \in \text{st}(B)} u_i(\mathbf{a}_{\mathcal{N}(i)}). \quad (\text{E.2})$$

Proof. For leaf bags B ($C(B) = \emptyset$), for any joint action $\mathbf{a}_B \in \mathcal{A}_B$, we have

$$h(B, \mathbf{a}_B) = \sum_{\substack{i \in [N]: \\ i \text{ assigned to } B}} u_i(\mathbf{a}_{\mathcal{N}(i)}) = \sum_{i \in \text{st}(B)} u_i(\mathbf{a}_{\mathcal{N}(i)}) \stackrel{(i)}{=} \max_{\mathbf{a}_{-B} \in \mathcal{A}_{-B}} \sum_{i \in \text{st}(B)} u_i(\mathbf{a}_{\mathcal{N}(i)}).$$

(i) is because $\mathcal{N}(i) \subseteq B$ for any i assigned to B by definition. Additionally, since B is a leaf bag, $\text{st}(B) = \{i \in B: i \text{ assigned to } B\}$.

Then, for any bag B with all of its children $B' \in C(B)$ satisfying (E.2), we have

$$\begin{aligned} h(B, \mathbf{a}_B) &= \sum_{\substack{i \in [N]: \\ i \text{ assigned to } B}} u_i(\mathbf{a}_{\mathcal{N}(i)}) + \sum_{B' \in C(B)} \max_{\substack{\mathbf{a}'_{B'} \in \mathcal{A}_{B'}: \\ \mathbf{a}_{B \cap B'} = \mathbf{a}'_{B \cap B'}}} h(B', \mathbf{a}'_{B'}) \\ &\stackrel{(i)}{=} \sum_{\substack{i \in [N]: \\ i \text{ assigned to } B}} u_i(\mathbf{a}_{\mathcal{N}(i)}) + \sum_{B' \in C(B)} \max_{\substack{\mathbf{a}'_{B'} \in \mathcal{A}_{B'}: \\ \mathbf{a}_{B \cap B'} = \mathbf{a}'_{B \cap B'}}} \max_{\mathbf{a}'_{-B'} \in \mathcal{A}_{-B'}} \sum_{i \in \text{st}(B')} u_i(\mathbf{a}'_{\mathcal{N}(i)}) \\ &= \sum_{\substack{i \in [N]: \\ i \text{ assigned to } B}} u_i(\mathbf{a}_{\mathcal{N}(i)}) + \sum_{B' \in C(B)} \max_{\substack{\mathbf{a}' \in \mathcal{A}: \\ \mathbf{a}_{B \cap B'} = \mathbf{a}'_{B \cap B'}}} \sum_{i \in \text{st}(B')} u_i(\mathbf{a}'_{\mathcal{N}(i)}). \end{aligned}$$

(i) uses the induction hypothesis. By **Property 3 of Tree Decomposition**, for any $B' \in C(B)$ and $i \in \text{st}(B')$, $\mathcal{N}(i) \cap (B \setminus B') = \emptyset$. Because for any $i \in \text{st}(B')$, there must be a bag B'' in the subtree of B' such that $\mathcal{N}(i) \subseteq B''$, and **Property 3 of Tree Decomposition** will be violated if $\mathcal{N}(i) \cap B \setminus B' \neq \emptyset$. Then, modifying the constraint $\mathbf{a}_{B \cap B'} = \mathbf{a}'_{B \cap B'}$ to $\mathbf{a}_B = \mathbf{a}'_B$ will not change the value of $u_i(\mathbf{a}'_{\mathcal{N}(i)})$ for any $i \in \text{st}(B')$. Hence,

$$h(B, \mathbf{a}_B) = \sum_{\substack{i \in [N]: \\ i \text{ assigned to } B}} u_i(\mathbf{a}_{\mathcal{N}(i)}) + \sum_{B' \in C(B)} \max_{\substack{\mathbf{a}' \in \mathcal{A}: \\ \mathbf{a}_B = \mathbf{a}'_B}} \sum_{i \in \text{st}(B')} u_i(\mathbf{a}'_{\mathcal{N}(i)}).$$

Furthermore, by **Property 3 of Tree Decomposition**, for any $B', B'' \in C(B)$ and $i' \in \text{st}(B'), i'' \in \text{st}(B'')$, we have $\mathcal{N}(i') \cap \mathcal{N}(i'') \subseteq B$. Finally,

$$\begin{aligned} h(B, \mathbf{a}_B) &= \sum_{\substack{i \in [N]: \\ i \text{ assigned to } B}} u_i(\mathbf{a}_{\mathcal{N}(i)}) + \max_{\substack{\mathbf{a}' \in \mathcal{A}: \\ \mathbf{a}_B = \mathbf{a}'_B}} \sum_{B' \in C(B)} \sum_{i \in \text{st}(B')} u_i(\mathbf{a}'_{\mathcal{N}(i)}) \\ &= \max_{\mathbf{a}_{-B} \in \mathcal{A}_{-B}} \left(\sum_{\substack{i \in [N]: \\ i \text{ assigned to } B}} u_i(\mathbf{a}_{\mathcal{N}(i)}) + \sum_{B' \in C(B)} \sum_{i \in \text{st}(B')} u_i(\mathbf{a}'_{\mathcal{N}(i)}) \right) \\ &= \max_{\mathbf{a}_{-B} \in \mathcal{A}_{-B}} \sum_{i \in \text{st}(B)} u_i(\mathbf{a}_{\mathcal{N}(i)}). \end{aligned}$$

This completes the induction. \square

F PROOF OF THEOREM 5.2

Theorem 5.2. Consider (5.3). For any $\delta > 0$, with probability at least $1 - \delta$, the following holds:

$$\max_{\hat{\pi} \in \Delta^{\mathcal{A}}} \sum_{t=1}^T F(\pi^{(t)}, \mu^{(t)}) - F(\hat{\pi}, \mu^{(t)}) \leq 2|\mathcal{T}| \frac{1 + (\text{tw}(\mathcal{G}) + 1) \log A}{\eta} + 2\eta |\mathcal{T}| T + \sqrt{2T \log \frac{1}{\delta}}.$$

Proof. The proof of Theorem 5.2 can be decomposed into three steps.

Firstly, we show that without loss of generality, if FTPL with a fixed noise $\tilde{\mathbf{n}}$ for all timesteps $t = 1, 2, \dots$ attains sublinear regret when the adversary is oblivious,⁸ then FTPL with independent noise vectors $\tilde{\mathbf{n}}^{(t)}$ also attains the same regret confronting an adaptive adversary. The reduction to the oblivious setting is common in the literature (Agarwal et al., 2019; Suggala & Netrapalli, 2020), and we include it here for completeness.

Secondly, we will show that the regret of a fictitious algorithm $\pi^{(t+1)} \in \arg\min_{\pi \in \Delta^{\mathcal{A}}} \sum_{\tau=1}^{(t+1)} F(\pi, \mu^{(\tau)}) + \langle \tilde{\mathbf{n}}^{(t+1)}, \pi \rangle$ is sublinear.

Finally, we will show that the regret of (5.3) and that of the fictitious algorithm are close.

F.1 FIXED NOISE VECTOR

In this section, for completeness, we will show a reduction from an adaptive adversary to an oblivious adversary. For ease of representation, we will take correlator (π) as the *no-regret learner*, and the deviator (μ) as the *adversary*.

An adaptive adversary determines the utility function at timestep t , which is $\mu^{(t)}$ in this section, according to our past strategies, $\pi^{(1)}, \dots, \pi^{(t-1)}$. In contrast, an oblivious adversary determines all utility functions, *i.e.*, $\mu^{(1)}, \dots, \mu^{(T)}$, at the beginning (timestep 0), such that $\mu^{(t)}$ is irrelevant to $\pi^{(1)}, \dots, \pi^{(t-1)}$. In the following, we will show that a sublinear regret against an oblivious adversary implies a sublinear regret against an adaptive adversary.

Intuitively, when the random noise $\tilde{\mathbf{n}}^{(1)}, \dots, \tilde{\mathbf{n}}^{(T)}$ are independent, $\pi^{(t)}$ only depends on $\mu^{(1)}, \dots, \mu^{(t-1)}$, which is known to both the oblivious and adaptive adversary, due to the update rule (5.3). Hence, an additional observation on $\pi^{(1)}, \dots, \pi^{(t-1)}$ does not make adversary more powerful. Formally, we have the following lemma (Cesa-Bianchi & Lugosi, 2006, Lemma 4.1).

Lemma F.1 (Reformulation of Lemma 4.1 in Cesa-Bianchi & Lugosi (2006)). Consider any randomized no-regret learner and the *distribution* of the decision variable $\pi^{(t)}$ is fully determined by $\mu^{(1)}, \dots, \mu^{(t-1)}$. Assume the no-regret learner’s regret against any sequence of $\mu^{(1)}, \dots, \mu^{(T)}$ generated by an oblivious adversary satisfies that

$$\underbrace{\mathbb{E}_{\tilde{\mathbf{n}}^{(1)}, \dots, \tilde{\mathbf{n}}^{(T)}} \left[\max_{\hat{\pi} \in \Delta^{\mathcal{A}}} \sum_{t=1}^T F(\pi^{(t)}, \mu^{(t)}) - F(\hat{\pi}, \mu^{(t)}) \right]}_{\text{Expected Regret}} \leq R.$$

Then, for any sequence of $\mu^{(1)}, \dots, \mu^{(T)}$ generated by an adaptive adversary and $\delta > 0$, with probability at least $1 - \delta$, we have

$$\max_{\hat{\pi} \in \Delta^{\mathcal{A}}} \sum_{t=1}^T F(\pi^{(t)}, \mu^{(t)}) - F(\hat{\pi}, \mu^{(t)}) \leq R + \sqrt{2T \log \frac{1}{\delta}}.$$

⁸An oblivious adversary will choose all the utility functions at timestep 0, while an adaptive adversary will choose the utility functions at timestep t according to $\pi^{(1)}, \pi^{(2)}, \dots, \pi^{(t-1)}$.

It is easy to see that the distribution of $\pi^{(t)}$ generated by (5.3), whose randomness is induced by $\tilde{\mathbf{n}}^{(t)}$, is fully determined by $\mu^{(1)}, \dots, \mu^{(t-1)}$, given $\tilde{\mathbf{n}}^{(t)}$ is generated by a fixed distribution independently at each timestep.

Then, we will show that the expected regret of FTPL using independent noise vectors and FTPL with a fixed noise vector is the same, while facing an oblivious adversary.

$$\begin{aligned} \mathbb{E}_{\tilde{\mathbf{n}}^{(1)}, \dots, \tilde{\mathbf{n}}^{(T)}} \left[\sum_{t=1}^T F \left(\pi^{(t)}, \mu^{(t)} \right) \right] &= \sum_{t=1}^T \mathbb{E}_{\tilde{\mathbf{n}}^{(1)}, \dots, \tilde{\mathbf{n}}^{(T)}} \left[F \left(\pi^{(t)}, \mu^{(t)} \right) \right] \\ &\stackrel{(i)}{=} \sum_{t=1}^T \mathbb{E}_{\tilde{\mathbf{n}}^{(t)}} \left[F \left(\pi^{(t)}, \mu^{(t)} \right) \right]. \end{aligned}$$

(i) uses the fact that both $\pi^{(t)}$ and $\mu^{(t)}$ are independent of $\tilde{\mathbf{n}}^{(1)}, \dots, \tilde{\mathbf{n}}^{(t-1)}$, when the adversary controlling μ is oblivious. Finally, the expectation $\mathbb{E}_{\tilde{\mathbf{n}}^{(1)}} \left[F \left(\pi^{(t)}, \mu^{(t)} \right) \right]$ of

$$\pi^{(t+1)} \in \operatorname{argmin}_{\pi \in \Delta^{\mathcal{A}}} \sum_{\tau=1}^t F \left(\pi, \mu^{(\tau)} \right) - \langle \tilde{\mathbf{n}}^{(1)}, \pi \rangle$$

is equal to $\mathbb{E}_{\tilde{\mathbf{n}}^{(t)}} \left[F \left(\pi^{(t)}, \mu^{(t)} \right) \right]$, when $\mathbf{n}^{(t)}$ and $\mathbf{n}^{(1)}$ are sampled from an identical distribution. In summary,

- fixed noise and an oblivious adversary (Expectation)
- \Rightarrow independent noise and an oblivious adversary (Expectation)
- \Rightarrow independent noise and an adaptive adversary (High Probability Bound).

Hence, the problem reduces to proving sublinear regret against an oblivious adversary with a fixed noise vector for all timesteps.

F.2 LOW REGRET WITH ACCURATE PREDICTION

The discussion above suggests that we only need to show the sublinear regret against an oblivious adversary, when all timesteps share the same noise vector. In other words, we consider the following update rule,

$$\pi^{(t+1)} \in \operatorname{argmin}_{\pi \in \Delta^{\mathcal{A}}} \sum_{\tau=1}^t F \left(\pi, \mu^{(\tau)} \right) - \langle \tilde{\mathbf{n}}, \pi \rangle, \quad (\text{F.1})$$

where $\tilde{\mathbf{n}}, \tilde{\mathbf{n}}^{(1)}, \dots, \tilde{\mathbf{n}}^{(T)}$ are identically distributed.

Next, we will show that if the regret minimizer can make the decision $\pi^{(t+1)}$ with an accurate prediction of $\mu^{(t+1)}$, then we can achieve sublinear regret. In particular, the decision variable at timestep $t+1$ is chosen according to (F.2).

$$\pi^{(t+1)} \in \operatorname{argmin}_{\pi \in \Delta^{\mathcal{A}}} \sum_{\tau=1}^{t+1} F \left(\pi, \mu^{(\tau)} \right) - \langle \tilde{\mathbf{n}}, \pi \rangle. \quad (\text{F.2})$$

Actually, we can see that (F.2) is exactly the original update rule of $\pi^{(t+2)}$. Therefore, we will prove the following lemma in the sequel.

Lemma F.2. Consider (F.2). For any timestep $t = 1, 2, \dots$, $\mu^{(1)}, \mu^{(2)}, \dots$, and any $\hat{\pi} \in \Delta^{\mathcal{A}}$, we have

$$\sum_{\tau=1}^t \left(F \left(\pi^{(\tau+1)}, \mu^{(\tau)} \right) - F \left(\hat{\pi}, \mu^{(\tau)} \right) \right) \leq \langle \tilde{\mathbf{n}}, \pi^{(2)} - \hat{\pi} \rangle. \quad (\text{F.3})$$

Proof. We will prove the lemma by induction. When $t = 1$, we have

$$\begin{aligned} &F \left(\pi^{(2)}, \mu^{(1)} \right) - F \left(\hat{\pi}, \mu^{(1)} \right) \\ &= \left(F \left(\pi^{(2)}, \mu^{(1)} \right) - \langle \tilde{\mathbf{n}}, \pi^{(2)} \rangle \right) - \left(F \left(\hat{\pi}, \mu^{(1)} \right) - \langle \tilde{\mathbf{n}}, \hat{\pi} \rangle \right) + \langle \tilde{\mathbf{n}}, \pi^{(2)} - \hat{\pi} \rangle \\ &\stackrel{(i)}{\leq} \langle \tilde{\mathbf{n}}, \pi^{(2)} - \hat{\pi} \rangle. \end{aligned}$$

(i) is because $\pi^{(2)} \in \operatorname{argmin}_{\pi \in \Delta^{\mathcal{A}}} F(\pi, \mu^{(1)}) - \langle \tilde{\mathbf{n}}, \pi \rangle$.

Next, we will show that when (F.3) holds for $t = t_0$, then it also holds for $t = t_0 + 1$. For any $\hat{\pi} \in \Delta^{\mathcal{A}}$, we have

$$\begin{aligned} & \sum_{\tau=1}^{t_0+1} \left(F(\pi^{(\tau+1)}, \mu^{(\tau)}) - F(\hat{\pi}, \mu^{(\tau)}) \right) \\ &= \sum_{\tau=1}^{t_0+1} F(\pi^{(\tau+1)}, \mu^{(\tau)}) - \left(\sum_{\tau=1}^{t_0+1} F(\hat{\pi}, \mu^{(\tau)}) - \langle \tilde{\mathbf{n}}, \hat{\pi} \rangle \right) - \langle \tilde{\mathbf{n}}, \hat{\pi} \rangle \\ &\stackrel{(i)}{\leq} \sum_{\tau=1}^{t_0+1} F(\pi^{(\tau+1)}, \mu^{(\tau)}) - \left(\sum_{\tau=1}^{t_0+1} F(\pi^{(t_0+2)}, \mu^{(\tau)}) - \langle \tilde{\mathbf{n}}, \pi^{(t_0+2)} \rangle \right) - \langle \tilde{\mathbf{n}}, \hat{\pi} \rangle \\ &= \sum_{\tau=1}^{t_0} \left(F(\pi^{(\tau+1)}, \mu^{(\tau)}) - F(\pi^{(t_0+2)}, \mu^{(\tau)}) \right) + \langle \tilde{\mathbf{n}}, \pi^{(t_0+2)} - \hat{\pi} \rangle. \end{aligned}$$

(i) is because $\pi^{(t_0+2)} \in \operatorname{argmin}_{\pi \in \Delta^{\mathcal{A}}} \sum_{\tau=1}^{t_0+1} F(\pi, \mu^{(\tau)}) - \langle \tilde{\mathbf{n}}, \pi \rangle$. Next, by setting $\hat{\pi} = \pi^{(t_0+2)}$ in the induction hypothesis, we have

$$\begin{aligned} & \sum_{\tau=1}^{t_0} \left(F(\pi^{(\tau+1)}, \mu^{(\tau)}) - F(\pi^{(t_0+2)}, \mu^{(\tau)}) \right) + \langle \tilde{\mathbf{n}}, \pi^{(t_0+2)} - \hat{\pi} \rangle \\ &\leq \langle \tilde{\mathbf{n}}, \pi^{(2)} - \pi^{(t_0+2)} \rangle + \langle \tilde{\mathbf{n}}, \pi^{(t_0+2)} - \hat{\pi} \rangle \\ &= \langle \tilde{\mathbf{n}}, \pi^{(2)} - \hat{\pi} \rangle. \quad \square \end{aligned}$$

F.3 SUBLINEAR VARIATION

In this section, we will show that the regret of FTPL with/without a prediction of $\mu^{(t+1)}$ is close. Formally, for any $\hat{\pi} \in \Delta^{\mathcal{A}}$, we have

$$\begin{aligned} & \sum_{t=1}^T \left(F(\pi^{(t)}, \mu^{(t)}) - F(\hat{\pi}, \mu^{(t)}) \right) \\ &= \sum_{t=1}^T \left(F(\pi^{(t+1)}, \mu^{(t)}) - F(\hat{\pi}, \mu^{(t)}) \right) + \sum_{t=1}^T \left(F(\pi^{(t)}, \mu^{(t)}) - F(\pi^{(t+1)}, \mu^{(t)}) \right) \\ &\stackrel{(i)}{\leq} \langle \tilde{\mathbf{n}}, \pi^{(2)} - \hat{\pi} \rangle + \sum_{t=1}^T \left(F(\pi^{(t)}, \mu^{(t)}) - F(\pi^{(t+1)}, \mu^{(t)}) \right). \end{aligned}$$

(i) uses Lemma F.2.

Moreover, since $\mathcal{U}_i(\mathbf{a}) \in [0, 1]$ for any $i \in [N]$ and $\mathbf{a} \in \mathcal{A}$, we have

$$\max_{\substack{\pi \in \Delta^{\mathcal{A}}, \\ \mu \in \Delta^{\times_{S \in \mathcal{S}} \mathcal{A}_S}}} |F(\pi, \mu)| \leq 1.$$

Then,

$$\mathbb{E} \left[F(\pi^{(t)}, \mu^{(t)}) - F(\pi^{(t+1)}, \mu^{(t)}) \right] \leq \Pr_{\tilde{\mathbf{n}}} \left(\pi^{(t)} \neq \pi^{(t+1)} \right).$$

Hence, we only need to lower bound $\Pr_{\tilde{\mathbf{n}}}(\pi^{(t)} = \pi^{(t+1)})$ in the sequel. Recall that $\pi^{(t)}$ is a pure strategy with $\pi^{(t)}(\mathbf{a}^{(t)}) = 1$ for some joint action $\mathbf{a}^{(t)} \in \mathcal{A}$. For any bag $B \in \mathcal{T}$, let $fa(B)$ denote its father ($fa(B) = \emptyset$ if B is the root). Then,

$$\begin{aligned} \Pr_{\tilde{\mathbf{n}}}(\pi^{(t)} = \pi^{(t+1)}) &= \Pr_{\tilde{\mathbf{n}}}(\mathbf{a}^{(t)} = \mathbf{a}^{(t+1)}) \\ &\stackrel{(i)}{=} \prod_{B \in \mathcal{T}} \Pr_{\tilde{\mathbf{n}}}(\mathbf{a}_{B \setminus fa(B)}^{(t)} = \mathbf{a}_{B \setminus fa(B)}^{(t+1)} \mid \mathbf{a}_{fa(B)}^{(t)} = \mathbf{a}_{fa(B)}^{(t+1)}). \end{aligned}$$

According to (5.5), each bag $B \in \mathcal{T}$ only determines $\mathbf{a}_{B \setminus fa(B)}^{(t)}$. Since $\mathbf{n}^{(t)}(B, \cdot)$ is sampled independently for every bag B , it follows that $\mathbf{a}_{B' \setminus B}^{(t)}, \mathbf{a}_{B'' \setminus B}^{(t)}$ are independent for $B', B'' \in C(B)$ with $B' \neq B''$, by **Property 3 of Tree Decomposition**. Hence, (i) holds.

For any $B \in \mathcal{T}$, to lower bound $\Pr_{\tilde{\mathbf{n}}} \left(\mathbf{a}_{B \setminus fa(B)}^{(t)} = \mathbf{a}_{B \setminus fa(B)}^{(t+1)} \mid \mathbf{a}_{fa(B)}^{(t)} = \mathbf{a}_{fa(B)}^{(t+1)} \right)$, we will first get its lower bound while further conditioning on $\mathbf{a}_{B \cup fa(B)}^{(t)}$'s value and $\mathbf{n}(B, \cdot)$'s value.

Then, $\Pr_{\tilde{\mathbf{n}}} \left(\mathbf{a}_{B \setminus fa(B)}^{(t)} = \mathbf{a}_{B \setminus fa(B)}^{(t+1)} \mid \mathbf{a}_{fa(B)}^{(t)} = \mathbf{a}_{fa(B)}^{(t+1)} \right)$ is equal to this conditioned probability integrating over all possible values of $\mathbf{a}_{B \cup fa(B)}^{(t)}$ and $\mathbf{n}(B, \cdot)$. Formally, we want to lower bound

$$\begin{aligned} p_B^{(t)}(\mathbf{a}'_{B \cup fa(B)}, \mathbf{x}) &:= \Pr \left(\mathbf{a}_{B \setminus fa(B)}^{(t+1)} = \mathbf{a}'_{B \setminus fa(B)} \mid \mathbf{a}_{fa(B)}^{(t)} = \mathbf{a}_{fa(B)}^{(t+1)}, \mathbf{a}_{B \cup fa(B)}^{(t)} = \mathbf{a}'_{B \cup fa(B)}, \right. \\ &\quad \text{and } \forall \mathbf{a}_{B \setminus fa(B)} \in \mathcal{A}_{B \setminus fa(B)} \setminus \left\{ \mathbf{a}'_{B \setminus fa(B)} \right\}, \\ &\quad \left. n(B, (\mathbf{a}'_{B \cap fa(B)}, \mathbf{a}_{B \setminus fa(B)})) = x((\mathbf{a}'_{B \cap fa(B)}, \mathbf{a}_{B \setminus fa(B)})) \right) \end{aligned}$$

for any $\mathbf{a}'_{B \cup fa(B)} \in \mathcal{A}_{B \cup fa(B)}$ and $\mathbf{x} \in \mathbb{R}^{\mathcal{A}_B}$. Then,

$$\Pr_{\tilde{\mathbf{n}}} \left(\mathbf{a}_{B \setminus fa(B)}^{(t)} = \mathbf{a}_{B \setminus fa(B)}^{(t+1)} \mid \mathbf{a}_{fa(B)}^{(t)} = \mathbf{a}_{fa(B)}^{(t+1)} \right) \geq \inf_{\substack{\mathbf{a}'_{B \cup fa(B)} \in \mathcal{A}_{B \cup fa(B)}, \\ \mathbf{x} \in \mathbb{R}^{\mathcal{A}_B}}} p_B^{(t)}(\mathbf{a}'_{B \cup fa(B)}, \mathbf{x}),$$

since $\Pr_{\tilde{\mathbf{n}}} \left(\mathbf{a}_{B \setminus fa(B)}^{(t)} = \mathbf{a}_{B \setminus fa(B)}^{(t+1)} \mid \mathbf{a}_{fa(B)}^{(t)} = \mathbf{a}_{fa(B)}^{(t+1)} \right)$ is equal to $p_B^{(t)}(\mathbf{a}'_{B \cup fa(B)}, \mathbf{x})$ integrating over $\mathbf{a}'_{B \cup fa(B)}$ and \mathbf{x} .

Since $\mathbf{a}_{B \setminus fa(B)}^{(t)} = \mathbf{a}'_{B \setminus fa(B)}$, for any $\mathbf{a}_{B \setminus fa(B)} \in \mathcal{A}_{B \setminus fa(B)} \setminus \left\{ \mathbf{a}'_{B \setminus fa(B)} \right\}$, we have

$$d^{(t)}(B, \mathbf{a}'_B) \leq d^{(t)}\left(B, (\mathbf{a}'_{B \cap fa(B)}, \mathbf{a}_{B \setminus fa(B)})\right).$$

This can be equivalently written as

$$n(B, \mathbf{a}'_B) \geq \left(d^{(t)}(B, \mathbf{a}'_B) + n(B, \mathbf{a}'_B) \right) - d^{(t)}\left(B, (\mathbf{a}'_{B \cap fa(B)}, \mathbf{a}_{B \setminus fa(B)})\right).$$

Then, $\mathbf{a}_{B \setminus fa(B)}^{(t+1)} = \mathbf{a}'_{B \setminus fa(B)}$ is equivalent to

$$\begin{aligned} n(B, \mathbf{a}'_B) &\geq \left(d^{(t+1)}(B, \mathbf{a}'_B) + n(B, \mathbf{a}'_B) \right) - d^{(t+1)}\left(B, (\mathbf{a}'_{B \cap fa(B)}, \mathbf{a}_{B \setminus fa(B)})\right) \quad (\text{F.4}) \\ &= \left(d^{(t)}(B, \mathbf{a}'_B) + n(B, \mathbf{a}'_B) \right) - d^{(t)}\left(B, (\mathbf{a}'_{B \cap fa(B)}, \mathbf{a}_{B \setminus fa(B)})\right) \\ &\quad + \left(d^{(t)}\left(B, (\mathbf{a}'_{B \cap fa(B)}, \mathbf{a}_{B \setminus fa(B)})\right) - d^{(t+1)}\left(B, (\mathbf{a}'_{B \cap fa(B)}, \mathbf{a}_{B \setminus fa(B)})\right) \right) \\ &\quad + \left(d^{(t+1)}(B, \mathbf{a}'_B) - d^{(t)}(B, \mathbf{a}'_B) \right). \end{aligned}$$

for any $\mathbf{a}_{B \setminus fa(B)} \in \mathcal{A}_{B \setminus fa(B)}$. In Lemma F.3, we show that the variation of d is bounded by 1. Therefore,

$$n(B, \mathbf{a}'_B) \geq \left(d^{(t)}(B, \mathbf{a}'_B) + n(B, \mathbf{a}'_B) \right) - d^{(t)}\left(B, (\mathbf{a}'_{B \cap fa(B)}, \mathbf{a}_{B \setminus fa(B)})\right) + 2$$

implies (F.4).

Lemma F.3. Consider the update rule (5.3). For any timestep $t = 1, 2, \dots, T$, bag $B \in \mathcal{T}$, joint action $\mathbf{a}_B \in \mathcal{A}_B$, and noise $\tilde{\mathbf{n}} \in \mathbb{R}^{\times_{B \in \mathcal{T}} \mathcal{A}_B}$, we have

$$\left| d^{(t+1)}(B, \mathbf{a}_B) - d^{(t)}(B, \mathbf{a}_B) \right| \leq 1.$$

The proof is postponed to the end of this section. Let

$$w = \max_{\mathbf{a}_{B \setminus f_a(B)} \in \mathcal{A}_{B \setminus f_a(B)} \setminus \{\mathbf{a}'_{B \setminus f_a(B)}\}} \left(d^{(t)}(B, \mathbf{a}'_B) + n(B, \mathbf{a}'_B) \right) - d^{(t)}\left(B, (\mathbf{a}'_{B \cap f_a(B)}, \mathbf{a}_{B \setminus f_a(B)})\right).$$

Note that w only depends on $\mu^{(1)}, \dots, \mu^{(t-1)}$ and \mathbf{x} . Then,

$$\begin{aligned} p_B^{(t)}(\mathbf{a}'_{B \cup f_a(B)}, \mathbf{x}) &\geq \Pr(n(B, \mathbf{a}'_B) \geq w + 2 \mid n(B, \mathbf{a}'_B) \geq w) \\ &= \frac{\Pr(n(B, \mathbf{a}'_B) \geq w + 2)}{\Pr(n(B, \mathbf{a}'_B) \geq w)} \\ &\stackrel{(i)}{=} \frac{\exp(-\eta(w + 2))}{\exp(-\eta w)} \\ &= \exp(-2\eta). \end{aligned}$$

(i) is because $n(B, \mathbf{a}'_B) \sim \text{Exp}(\eta)$. Finally, by union bound,

$$\begin{aligned} \Pr_{\tilde{\mathbf{n}}}(\pi^{(t)} = \pi^{(t+1)}) &\geq 1 - \sum_{B \in \mathcal{T}} \left(1 - p_B^{(t)}(\mathbf{a}'_{B \cup f_a(B)}, \mathbf{x})\right) \\ &\geq 1 - \sum_{B \in \mathcal{T}} (1 - \exp(-2\eta)) \\ &\geq 1 - \sum_{B \in \mathcal{T}} 2\eta \\ &= 1 - 2\eta |\mathcal{T}|. \end{aligned}$$

Therefore,

$$\begin{aligned} &\mathbb{E} \left[\sum_{t=1}^T F(\pi^{(t)}, \mu^{(t)}) - F(\hat{\pi}, \mu^{(t)}) \right] \\ &\leq \mathbb{E} \left[\langle \tilde{\mathbf{n}}, \pi^{(2)} - \hat{\pi} \rangle \right] + \sum_{t=1}^T \mathbb{E} \left[F(\pi^{(t)}, \mu^{(t)}) - F(\pi^{(t+1)}, \mu^{(t)}) \right] \\ &\leq \mathbb{E} \left[\langle \tilde{\mathbf{n}}, \pi^{(2)} - \hat{\pi} \rangle \right] + 2\eta |\mathcal{T}| T. \end{aligned}$$

Since $\exp(x) \geq 1 + x$ for any $x \in \mathbb{R}$, $(1 - \exp(-2\eta |\mathcal{T}|)) \leq 2\eta |\mathcal{T}|$. Additionally,

$$\begin{aligned} \mathbb{E} \left[\langle \tilde{\mathbf{n}}, \pi^{(2)} - \hat{\pi} \rangle \right] &\stackrel{(i)}{\leq} \mathbb{E} \left[\|\tilde{\mathbf{n}}\|_{\infty} \cdot \|\pi^{(2)} - \hat{\pi}\|_1 \right] \leq 2\mathbb{E}[\|\tilde{\mathbf{n}}\|_{\infty}] \\ &\leq 2 \sum_{B \in \mathcal{T}} \max_{\mathbf{a}_B \in \mathcal{A}_B} n(B, \mathbf{a}_B) \\ &\stackrel{(i)}{\leq} 2 \sum_{B \in \mathcal{T}} \frac{1 + \log |\mathcal{A}_B|}{\eta}. \end{aligned}$$

(i) is by Hölder's Inequality. (ii) is because the expectation of the maximum of n i.i.d. random variable sampled from $\text{Exp}(\eta)$ is upper bounded by $\frac{1 + \log n}{\eta}$ (Agarwal et al., 2019). Furthermore, $\log |\mathcal{A}_B| \leq |B| \cdot \log A \leq (\text{tw}(\mathcal{G}) + 1) \log A$. Hence,

$$\mathbb{E} \left[\sum_{t=1}^T F(\pi^{(t)}, \mu^{(t)}) - F(\hat{\pi}, \mu^{(t)}) \right] \leq 2|\mathcal{T}| \frac{1 + (\text{tw}(\mathcal{G}) + 1) \log A}{\eta} + 2\eta |\mathcal{T}| T.$$

□

F.4 PROOF OF AUXILIARY LEMMAS

Lemma F.3. Consider the update rule (5.3). For any timestep $t = 1, 2, \dots, T$, bag $B \in \mathcal{T}$, joint action $\mathbf{a}_B \in \mathcal{A}_B$, and noise $\tilde{\mathbf{n}} \in \mathbb{R}^{\times_{B \in \mathcal{T}} \mathcal{A}_B}$, we have

$$\left| d^{(t+1)}(B, \mathbf{a}_B) - d^{(t)}(B, \mathbf{a}_B) \right| \leq 1.$$

Proof. Recall Lemma E.1. We can add $|\mathcal{T}|$ players as the noise player, each assigned to a bag in \mathcal{T} , with $u_i(\mathbf{a}_B) = n(B, \mathbf{a}_B)$ so that $\mathcal{N}(i) = B$. Recall that u_i is the contribution of player i to the objective function F . Then, by Lemma E.1, for any $B \in \mathcal{T}$ and $\mathbf{a}_B \in \mathcal{A}_B$, we have

$$d^{(t)}(B, \mathbf{a}_B) = \min_{\mathbf{a}_{-B} \in \mathcal{A}_{-B}} \sum_{i \in \text{st}(B)} u_i^{(t)}(\mathbf{a}_{\mathcal{N}(i)}),$$

where

$$u_i^{(t)}(\mathbf{a}_{\mathcal{N}(i)}) = - \sum_{\tau=1}^t \sum_{\substack{S \in \mathcal{S}: \\ i \in S}} \frac{1}{|S|} \sum_{\hat{\mathbf{a}}_{\mathcal{N}_i^S} \in \mathcal{A}_{\mathcal{N}_i^S}} \mu^{(\tau)}(S, \hat{\mathbf{a}}_{\mathcal{N}_i^S}) \left(\mathcal{U}_i(\hat{\mathbf{a}}_{\mathcal{N}_i^S}, \mathbf{a}_{\mathcal{N}(i) \setminus S}) - \mathcal{U}_i(\mathbf{a}_{\mathcal{N}(i)}) \right) \text{ for } i \in [N]$$

$u_i^{(t)}(\mathbf{a}_{\mathcal{N}(i)}) = n(B, \mathbf{a}_{\mathcal{N}(i)})$ for i as the noise player assigned to bag B .

For any noise player, we can see that $u_i^{(t)}(\mathbf{a}_{\mathcal{N}(i)}) - u_i^{(t+1)}(\mathbf{a}_{\mathcal{N}(i)}) = 0$. For any $i \in [N]$,

$$\begin{aligned} & \left| u_i^{(t)}(\mathbf{a}_{\mathcal{N}(i)}) - u_i^{(t+1)}(\mathbf{a}_{\mathcal{N}(i)}) \right| \\ &= \left| \sum_{\substack{S \in \mathcal{S}: \\ i \in S}} \frac{1}{|S|} \sum_{\hat{\mathbf{a}}_{\mathcal{N}_i^S} \in \mathcal{A}_{\mathcal{N}_i^S}} \mu^{(t+1)}(S, \hat{\mathbf{a}}_{\mathcal{N}_i^S}) \left(\mathcal{U}_i(\hat{\mathbf{a}}_{\mathcal{N}_i^S}, \mathbf{a}_{\mathcal{N}(i) \setminus S}) - \mathcal{U}_i(\mathbf{a}_{\mathcal{N}(i)}) \right) \right| \\ &\leq \sum_{\substack{S \in \mathcal{S}: \\ i \in S}} \frac{1}{|S|} \sum_{\hat{\mathbf{a}}_{\mathcal{N}_i^S} \in \mathcal{A}_{\mathcal{N}_i^S}} \mu^{(t+1)}(S, \hat{\mathbf{a}}_{\mathcal{N}_i^S}) \left| \mathcal{U}_i(\hat{\mathbf{a}}_{\mathcal{N}_i^S}, \mathbf{a}_{\mathcal{N}(i) \setminus S}) - \mathcal{U}_i(\mathbf{a}_{\mathcal{N}(i)}) \right| \\ &\stackrel{(i)}{\leq} \sum_{\substack{S \in \mathcal{S}: \\ i \in S}} \frac{1}{|S|} \sum_{\hat{\mathbf{a}}_{\mathcal{N}_i^S} \in \mathcal{A}_{\mathcal{N}_i^S}} \mu^{(t+1)}(S, \hat{\mathbf{a}}_{\mathcal{N}_i^S}). \end{aligned}$$

(i) is because $\mathcal{U}_i(\mathbf{a}_{\mathcal{N}(i)}) \in [0, 1]$ for any $\mathbf{a}_{\mathcal{N}(i)} \in \mathcal{A}_{\mathcal{N}(i)}$.

Recall that by definition, $\mathbf{a}_{-B}^{(t)} = \operatorname{argmax}_{\mathbf{a}_{-B} \in \mathcal{A}_{-B}} \sum_{i \in \text{st}(B)} u_i^{(t)}(\mathbf{a}_{\mathcal{N}(i)})$. Then,

$$\begin{aligned} & d^{(t)}(B, \mathbf{a}_B) \\ &= \sum_{i \in \text{st}(B)} u_i^{(t)}\left(\left(\mathbf{a}_{\mathcal{N}(i) \cap B}, \mathbf{a}_{\mathcal{N}(i) \setminus B}^{(t)}\right)\right) \\ &\geq \sum_{i \in \text{st}(B)} u_i^{(t)}\left(\left(\mathbf{a}_{\mathcal{N}(i) \cap B}, \mathbf{a}_{\mathcal{N}(i) \setminus B}^{(t+1)}\right)\right) \\ &= \sum_{i \in \text{st}(B)} u_i^{(t+1)}\left(\left(\mathbf{a}_{\mathcal{N}(i) \cap B}, \mathbf{a}_{\mathcal{N}(i) \setminus B}^{(t+1)}\right)\right) + \left(u_i^{(t+1)}\left(\left(\mathbf{a}_{\mathcal{N}(i) \cap B}, \mathbf{a}_{\mathcal{N}(i) \setminus B}^{(t+1)}\right)\right) - u_i^{(t)}\left(\left(\mathbf{a}_{\mathcal{N}(i) \cap B}, \mathbf{a}_{\mathcal{N}(i) \setminus B}^{(t+1)}\right)\right) \right) \\ &\geq \sum_{i \in \text{st}(B)} \left(u_i^{(t+1)}\left(\left(\mathbf{a}_{\mathcal{N}(i) \cap B}, \mathbf{a}_{\mathcal{N}(i) \setminus B}^{(t+1)}\right)\right) - \sum_{\substack{S \in \mathcal{S}: \\ i \in S}} \frac{1}{|S|} \sum_{\hat{\mathbf{a}}_{\mathcal{N}_i^S} \in \mathcal{A}_{\mathcal{N}_i^S}} \mu^{(t+1)}(S, \hat{\mathbf{a}}_{\mathcal{N}_i^S}) \right) \\ &= d^{(t+1)}(B, \mathbf{a}_B) - \sum_{S \in \mathcal{S}} \frac{1}{|S|} \sum_{i \in S \cap \text{st}(B)} \sum_{\hat{\mathbf{a}}_{\mathcal{N}_i^S} \in \mathcal{A}_{\mathcal{N}_i^S}} \mu^{(t+1)}(S, \hat{\mathbf{a}}_{\mathcal{N}_i^S}) \\ &\geq d^{(t+1)}(B, \mathbf{a}_B) - 1. \end{aligned}$$

Similarly, we can get the upper bound that $d^{(t)}(B, \mathbf{a}_B) \leq d^{(t+1)}(B, \mathbf{a}_B) + 1$. Hence, the proof is completed. \square

G EFFICIENT UPDATE OF μ

This section provides the omitted details regarding the update procedure for μ and presents the complete proof of Theorem 5.3.

G.1 EFFICIENT UPDATE OF μ

The procedure for updating μ closely parallels that of π . Specifically, we iterate over all coalitions $S \in \mathcal{S}$ and, for each S , determine the optimal action $\mathbf{a}_S \in \mathcal{A}_S$. To achieve this, we maintain a dynamic programming vector $\mathbf{g}_S \in \mathbb{R}^{\times_{B \in \mathcal{T}} \mathcal{A}_B}$ for each $S \in \mathcal{S}$, which is updated according to

$$\begin{aligned} g_S^{(t+1)}(B, \hat{\mathbf{a}}_B) &= \frac{1}{|S|} \sum_{\tau=1}^t \sum_{\substack{i \in S: \\ i \text{ assigned to } B}} \sum_{\mathbf{a} \in \mathcal{A}} \pi^{(\tau)}(\mathbf{a}) (\mathcal{U}_i(\hat{\mathbf{a}}_{B \cap S}, \mathbf{a}_{B \setminus S}) - \mathcal{U}_i(\mathbf{a}_B)) \\ &+ \sum_{B' \in C(B)} \max_{\substack{\hat{\mathbf{a}}_{B'} \in \mathcal{A}_{B'}: \\ \hat{\mathbf{a}}_{B \cap B'} = \hat{\mathbf{a}}'_{B \cap B'}}} g^{(t+1)}(B', \hat{\mathbf{a}}_{B'}) + m^{(t+1)}(B, \hat{\mathbf{a}}_B). \end{aligned}$$

At first sight, the $\sum_{\mathbf{a} \in \mathcal{A}}$ appears computationally prohibitive, since \mathcal{A} is exponentially large. Fortunately, the update becomes tractable once we recall that $\pi^{(\tau)}$ is always a pure strategy for $\tau \geq 1$. Denote by $\mathbf{a}^{(\tau)}$ the joint action selected by $\pi^{(\tau)}$. Then (G.0) simplifies to

$$\begin{aligned} g_S^{(t+1)}(B, \hat{\mathbf{a}}_B) &= \frac{1}{|S|} \sum_{\tau=1}^t \sum_{\substack{i \in S: \\ i \text{ assigned to } B}} (\mathcal{U}_i(\hat{\mathbf{a}}_{B \cap S}, \mathbf{a}_{B \setminus S}^{(\tau)}) - \mathcal{U}_i(\mathbf{a}_B^{(\tau)})) \\ &+ \sum_{B' \in C(B)} \max_{\substack{\hat{\mathbf{a}}_{B'} \in \mathcal{A}_{B'}: \\ \hat{\mathbf{a}}_{B \cap B'} = \hat{\mathbf{a}}'_{B \cap B'}}} g^{(t+1)}(B', \hat{\mathbf{a}}_{B'}) + m_S^{(t+1)}(B, \hat{\mathbf{a}}_B). \end{aligned} \tag{G.1}$$

After completing the dynamic programming updates, we focus on the root bag B^r of the tree decomposition. The selected coalition is then $S^{(t+1)} = \operatorname{argmax}_{S \in \mathcal{S}} \max_{\hat{\mathbf{a}}_{B^r} \in \mathcal{A}_{B^r}} g_S^{(t+1)}(B^r, \hat{\mathbf{a}}_{B^r})$. Next, we apply the reconstruction procedure in (5.5) on $\mathbf{g}_{S^{(t+1)}}^{(t+1)}$ to extract a joint action $\hat{\mathbf{a}}^{(t+1)} \in \mathcal{A}$. Finally, we update $\mu^{(t+1)}(S^{(t+1)}, \hat{\mathbf{a}}^{(t+1)}) = 1$.

This procedure ensures that μ can be updated efficiently while maintaining consistency with the tree decomposition structure. Analogous to the update of π , the regret of this process can be bounded.

G.2 PROOF OF THEOREM 5.3

Theorem 5.3. Consider the updates in (5.3). For any $\delta > 0$, with probability at least $1 - \delta$, the following holds:

$$\max_{\hat{\mu} \in \Delta^{\times_{S \in \mathcal{S}} \mathcal{A}_S}} \sum_{t=1}^T F(\pi^{(t)}, \hat{\mu}) - F(\pi^{(t)}, \mu^{(t)}) \leq 2|\mathcal{T}| \frac{1 + (\operatorname{tw}(\mathcal{G}) + 1) \log A}{\eta} + 2\eta |\mathcal{T}| T + \sqrt{2T \log \frac{1}{\delta}}.$$

Proof. The proof of Theorem 5.3 is similar to that of Theorem 5.2. By using a similar argument as the proof of Theorem 5.2, for any $\hat{\mu} \in \Delta^{\times_{S \in \mathcal{S}} \mathcal{A}_S}$, we have

$$\begin{aligned} &\sum_{t=1}^T (F(\pi^{(t)}, \hat{\mu}) - F(\pi^{(t)}, \mu^{(t)})) \\ &\leq \langle \tilde{\mathbf{m}}, \hat{\mu} - \mu^{(2)} \rangle + \sum_{t=1}^T (F(\pi^{(t)}, \mu^{(t+1)}) - F(\pi^{(t)}, \mu^{(t)})). \end{aligned}$$

Next, by introducing the counterpart of Lemma F.3 in the following, the rest of the proof follows that of Theorem 5.2.

Lemma G.1. Consider the update rule (5.3). For any timestep $t = 1, 2, \dots, T$, bag $B \in \mathcal{T}$, joint action $\hat{\mathbf{a}}_B \in \mathcal{A}_B$, coalition $S \in \mathcal{S}$, and noise $\tilde{\mathbf{m}} \in \mathbb{R}^{\times_{B \in \mathcal{T}} \mathcal{A}_B}$, we have

$$\left| g_S^{(t+1)}(B, \hat{\mathbf{a}}_B) - g_S^{(t)}(B, \hat{\mathbf{a}}_B) \right| \leq 1.$$

The proof is postponed to the end of this section. Let $S^{(t)}, \widehat{\mathbf{a}}^{(t)}$ denote the coalition and action the deviator picks at timestep t , *i.e.*, $\mu^{(t)}(S^{(t)}, \widehat{\mathbf{a}}^{(t)}) = 1$. Then,

$$\begin{aligned} & \mathbb{E} \left[F \left(\pi^{(t)}, \mu^{(t+1)} \right) - F \left(\pi^{(t)}, \mu^{(t)} \right) \right] \\ & \leq \Pr_{\widehat{\mathbf{m}}} \left(\mu^{(t+1)} \neq \mu^{(t)} \right) \\ & = \Pr_{\widehat{\mathbf{m}}} \left(S^{(t)} = S^{(t+1)} \right) \prod_{B \in \mathcal{T}} \Pr_{\widehat{\mathbf{m}}} \left(\widehat{\mathbf{a}}_{B \setminus fa(B)}^{(t)} = \widehat{\mathbf{a}}_{B \setminus fa(B)}^{(t+1)} \mid \widehat{\mathbf{a}}_{fa(B)}^{(t)} = \widehat{\mathbf{a}}_{fa(B)}^{(t+1)}, S^{(t)} = S^{(t+1)} \right) \\ & \stackrel{(i)}{\leq} 1 - \exp(-2\eta|\mathcal{T}|). \end{aligned}$$

(i) is because choosing $S^{(t)}$ is equivalent to adding a new player in the root bag B^r , whose action is to select the coalition. Finally, for any $\delta > 0$, with probability at least $1 - \delta$,

$$\begin{aligned} & \sum_{t=1}^T \left(F \left(\pi^{(t)}, \widehat{\mu} \right) - F \left(\pi^{(t)}, \mu^{(t)} \right) \right) \\ & \leq 2|\mathcal{T}| \frac{1 + (\text{tw}(\mathcal{G}) + 1) \log A}{\eta} + 2\eta|\mathcal{T}|T + \sqrt{2T \log \frac{1}{\delta}}. \quad \square \end{aligned}$$

G.3 PROOF OF AUXILIARY LEMMAS

Lemma G.1. Consider the update rule (5.3). For any timestep $t = 1, 2, \dots, T$, bag $B \in \mathcal{T}$, joint action $\widehat{\mathbf{a}}_B \in \mathcal{A}_B$, coalition $S \in \mathcal{S}$, and noise $\widehat{\mathbf{m}} \in \mathbb{R}^{\times_{B \in \mathcal{T}} \mathcal{A}_B}$, we have

$$\left| g_S^{(t+1)}(B, \widehat{\mathbf{a}}_B) - g_S^{(t)}(B, \widehat{\mathbf{a}}_B) \right| \leq 1.$$

Proof. For any $S \in \mathcal{S}$, the upper bound of $\left| g_S^{(t+1)}(B, \widehat{\mathbf{a}}_B) - g_S^{(t)}(B, \widehat{\mathbf{a}}_B) \right|$ can be obtained similarly to the proof of Lemma F.3 by choosing

$$u_i^{(t)}(\mathbf{a}_{\mathcal{N}(i)}) = \frac{1}{|S|} \sum_{\tau=1}^t \sum_{\substack{i \in S: \\ i \text{ assigned to } B}} \sum_{\mathbf{a} \in \mathcal{A}} \pi^{(\tau)}(\mathbf{a}) \left(\mathcal{U}_i(\widehat{\mathbf{a}}_{B \cap S}, \mathbf{a}_{B \setminus S}) - \mathcal{U}_i(\mathbf{a}_B) \right) \text{ for } i \in [N]$$

$$u_i^{(t)}(\mathbf{a}_{\mathcal{N}(i)}) = m_S(B, \mathbf{a}_{\mathcal{N}(i)}) \text{ for } i \text{ as the noise player assigned to bag } B.$$

Then,

$$\begin{aligned} \left| u_i^{(t)}(\mathbf{a}_{\mathcal{N}(i)}) - u_i^{(t+1)}(\mathbf{a}_{\mathcal{N}(i)}) \right| & = \left| \frac{1}{|S|} \sum_{\substack{i \in S: \\ i \text{ assigned to } B}} \sum_{\mathbf{a} \in \mathcal{A}} \pi^{(t+1)}(\mathbf{a}) \left(\mathcal{U}_i(\widehat{\mathbf{a}}_{B \cap S}, \mathbf{a}_{B \setminus S}) - \mathcal{U}_i(\mathbf{a}_B) \right) \right| \\ & \stackrel{(i)}{\leq} \frac{1}{|S|} \sum_{\substack{i \in S: \\ i \text{ assigned to } B}} \sum_{\mathbf{a} \in \mathcal{A}} \pi^{(t+1)}(\mathbf{a}). \end{aligned}$$

(i) is by the fact that $\mathcal{U}_i(\widehat{\mathbf{a}}_{B \cap S}, \mathbf{a}_{B \setminus S}), \mathcal{U}_i(\mathbf{a}_B) \in [0, 1]$. The rest of the proof follows that of Lemma F.3, and thus we complete the proof. \square

H ADDITIONAL EXPERIMENTAL RESULTS

In this section, we detail our experimental setup and report additional results for two further games: the Chicken game (Bergstrom & Godfrey-Smith, 1998) and Pigou's network (Pigou, 2017). All experiments are conducted on a 13th Gen Intel(R) Core(TM) i7-13700K @ 3.40 GHz. Error bars for MASE and FTPL indicate $\pm 1\sigma$ over 100 random seeds (0, 1, ..., 99). Across all experiments, we set the learning rate to $\eta = 0.01$ and run for $T = 10,000$ timesteps.

H.1 EXPERIMENTAL DETAILS

For all baselines, we run each algorithm independently for each player, thus the average strategy converges to a CCE (Hazan et al., 2016). For **FTRL**, **Hedge**, and **OMD**, each player $i \in [N]$ is initialized with a uniform distribution $\pi_i^{(1)}$ over all actions. **MASE** and **FTPL** are initialized with a pure strategy chosen uniformly at random from all pure strategies.

H.2 UTILITY FUNCTIONS

This subsection specifies the utility functions for all four games.

Prisoner’s Dilemma. The utility matrix is shown in Table 2. If both prisoners confess, they receive reduced sentences. If one confesses while the other defects, the confessor is imprisoned and the defector is released immediately. If both defect, both are imprisoned for longer than in the mutual-confession case to penalize dishonesty.

	Confess (C)	Defect (D)
Confess (C)	(0.6, 0.6)	(0, 1)
Defect (D)	(1, 0)	(0.2, 0.2)

Table 2: Utility matrix of the Prisoner’s Dilemma. Each entry (a, b) denotes the payoffs to the row player (a) and the column player (b) .

Stag Hunt. The utility matrix is shown in Table 3. A stag yields a higher reward, but it can only be hunted successfully if both players choose Stag; a solo stag attempt yields nothing. A hare provides a smaller payoff but can be secured by a single player.

	Stag (S)	Hare (H)
Stag (S)	(1, 1)	(0.1, 0.8)
Hare (H)	(0.8, 0.1)	(0.5, 0.5)

Table 3: Utility matrix of the Stag Hunt. Each entry (a, b) denotes the payoff of the row player (a) and the column player (b) .

Chicken Game. Two drivers head toward each other and can either swerve or go straight. If one goes straight while the other swerves, the swerving player "loses." If both go straight, they crash.

	Swerve (Sw)	Straight (St)
Swerve (Sw)	(5/6, 5/6)	(2/3, 1)
Straight (St)	(1, 2/3)	(0, 0)

Table 4: Utility matrix of the Chicken game. Each entry (a, b) denotes the payoff of the row player (a) and the column player (b) .

Pigou Network. We use a three-player variant of Pigou’s network. Each player chooses a *fast* or *slow* route. The slow route yields a constant utility of 0.25. The fast route yields utility $1.5 - 0.5 \cdot$ (number of players choosing the fast route), reflecting congestion.

H.3 ADDITIONAL EXPERIMENTAL RESULTS

Figure 5 reports additional experiments on the Chicken game and Pigou’s network. **MASE** consistently outperforms the baselines in both coalition exploitability and social welfare. In Pigou’s network, purely self-interested players overuse the fast route, which in equilibrium becomes slow. By contrast, when players form coalitions and consider average utility within a coalition, they share the routes so that everyone is better off.

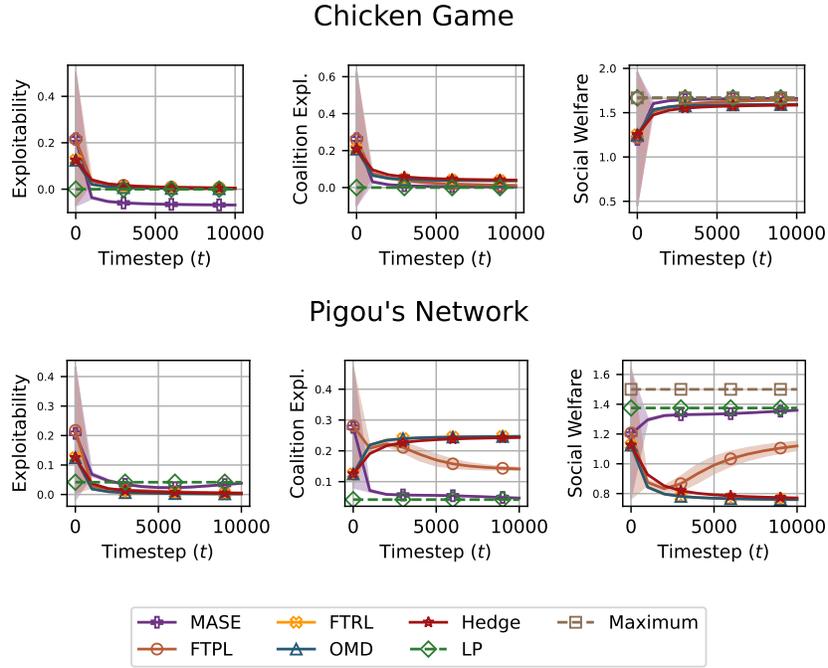


Figure 5: LP refers to the linear program in Section B. Maximum marks the maximum social welfare. MASE outperforms the baselines in both games in terms of coalition exploitability and social welfare.

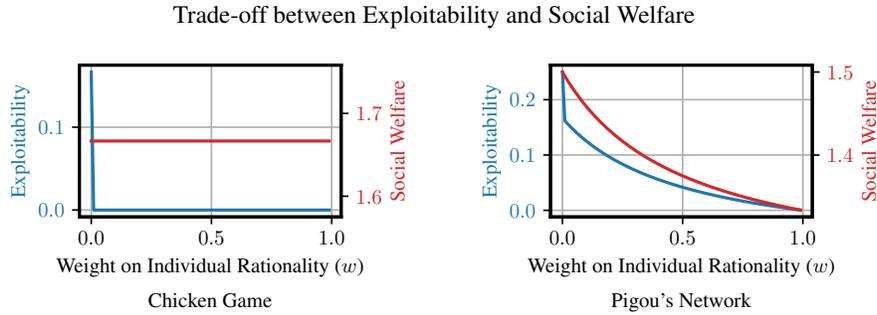


Figure 6: The trade-off between exploitability and social welfare in the Chicken game and the Pigou's network.

Figure 6 shows the trade-off between exploitability and social welfare in the Chicken game and Pigou's network.

Figure 7 reports the runtime of the algorithm for polymatrix games with varying numbers of players, action set sizes, and interaction densities.

H.4 PROOF OF LEMMA 6.1

Lemma 6.1. For any $\epsilon > 0$, computing the CCE with exploitability no more than ϵ that maximizes social welfare is equivalent to (6.1) by setting $\mathcal{S} = \{\{i\}\}_{i \in [N]} \cup \{[N]\}$ and using the weights:

$$w_S = \begin{cases} w & \text{if } |S| = 1 \\ 1 - w & \text{if } S = [N] \end{cases}$$

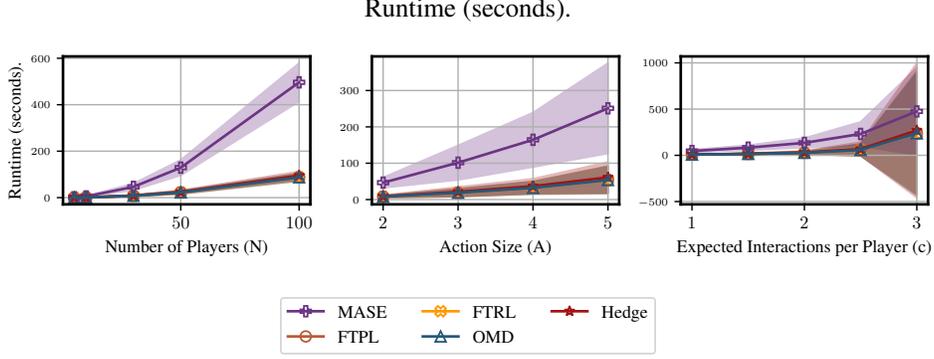


Figure 7: The runtime of the algorithm in different polymatrix games.

for some $w \in [0, 1)$. Conversely, solving (6.1) with these parameters corresponds to finding a point on the Pareto frontier of social welfare and exploitability.

Proof. For any $\epsilon > 0$, let π^* be the strategy that maximizes social welfare subject to its exploitability being at most ϵ . Let $g^* = \max_{\hat{\mathbf{a}} \in \mathcal{A}} \frac{1}{N} \sum_{i=1}^N \mathbb{E}_{\mathbf{a} \sim \pi^*} [\mathcal{U}_i(\hat{\mathbf{a}}) - \mathcal{U}_i(\mathbf{a})]$ be the maximum gain for the grand coalition $[N]$. Let $w = \frac{g^*}{\epsilon + g^*}$. Then, by construction, the objective value for π^* under (6.1) is:

$$\begin{aligned} \max_{S \in \mathcal{S}} \max_{\hat{\mathbf{a}}_S \in \mathcal{A}_S} \frac{w_S}{|S|} \sum_{i \in S} \mathbb{E}_{\mathbf{a} \sim \pi^*} [\dots] &= \max(w \cdot (\text{exploitability}), (1 - w) \cdot g^*) \\ &\leq \max(w\epsilon, (1 - w)g^*) = \frac{\epsilon g^*}{\epsilon + g^*}. \end{aligned}$$

Any strategy $\hat{\pi}$ with exploitability $> \epsilon$ would have an objective value $> w\epsilon = \frac{\epsilon g^*}{\epsilon + g^*}$, which is worse than the value π^* achieves. Therefore, any optimal solution to (6.1) must have exploitability at most ϵ . Since π^* by definition maximizes social welfare (i.e., minimizes the coalition gain g^*) among all strategies in this set, it must also be an optimal solution to (6.1).

Conversely, for any $w \in [0, 1)$, let π^* be the corresponding strategy that optimizes (6.1). Let its exploitability be

$$\epsilon = \max_{i \in [N]} \max_{\hat{\mathbf{a}}_i \in \mathcal{A}_i} \mathbb{E}_{\mathbf{a} \sim \pi^*} [\mathcal{U}_i(\hat{\mathbf{a}}_i, \mathbf{a}_{-i}) - \mathcal{U}_i(\mathbf{a})].$$

We will show by contradiction that no strategy π' exists such that $\text{exploitability}(\pi') \leq \epsilon$ and $SW(\pi') > SW(\pi^*)$ (which implies $g' < g^*$, where g' is the gain for the grand coalition under π').

Suppose such a π' exists. We analyze two cases:

Case 1: exploitability($\pi') < \epsilon$. Since π' has both strictly lower exploitability than π^* and $g' < g^*$ (higher social welfare), its objective value is $\max(w \cdot \text{exploitability}(\pi'), (1 - w)g')$. This is strictly less than $\max(w\epsilon, (1 - w)g^*)$, which is the objective value of π^* . This contradicts the optimality of π^* .

Case 2: exploitability($\pi') = \epsilon$. If $\epsilon > 0$, choose a small $\delta > 0$ and consider the mixed strategy $\pi_{new} = (1 - \delta)\pi' + \delta\pi''$, where π'' is an arbitrary CCE, which is guaranteed to exist (Nash, 1950). For any $i \in [N]$ and $\hat{\mathbf{a}}_i \in \mathcal{A}_i$, we have:

$$\begin{aligned} &\mathbb{E}_{\mathbf{a} \sim \pi_{new}} [\mathcal{U}_i(\hat{\mathbf{a}}_i, \mathbf{a}_{-i}) - \mathcal{U}_i(\mathbf{a})] \\ &= (1 - \delta)\mathbb{E}_{\mathbf{a} \sim \pi'} [\mathcal{U}_i(\hat{\mathbf{a}}_i, \mathbf{a}_{-i}) - \mathcal{U}_i(\mathbf{a})] + \delta\mathbb{E}_{\mathbf{a} \sim \pi''} [\mathcal{U}_i(\hat{\mathbf{a}}_i, \mathbf{a}_{-i}) - \mathcal{U}_i(\mathbf{a})] \\ &\stackrel{(i)}{\leq} (1 - \delta)\mathbb{E}_{\mathbf{a} \sim \pi'} [\mathcal{U}_i(\hat{\mathbf{a}}_i, \mathbf{a}_{-i}) - \mathcal{U}_i(\mathbf{a})] \\ &\leq (1 - \delta)\epsilon. \end{aligned}$$

Step (i) holds because π'' is a CCE, so its exploitability $\mathbb{E}_{\pi''}[\dots]$ is ≤ 0 . Since $\epsilon > 0$, the new strategy π_{new} has exploitability $(\pi_{new}) < \epsilon$. By continuity, for sufficiently small δ , $SW(\pi_{new})$ remains strictly higher than $SW(\pi^*)$ (since $SW(\pi') > SW(\pi^*)$). This puts us in Case 1, which leads to a contradiction.

If $\epsilon \leq 0$, then $\text{exploitability}(\pi^*) \leq 0$. The objective value for π^* is $\max(w\epsilon, (1-w)g^*) = (1-w)g^*$ (since $w\epsilon \leq 0$ and $(1-w)g^* \geq 0$ by definition). The hypothetical strategy π' has $\text{exploitability}(\pi') = \epsilon \leq 0$ and $g' < g^*$. Its objective value is $\max(w\epsilon, (1-w)g') = (1-w)g'$. Since $g' < g^*$ and $w < 1$, the objective value of π' is strictly less than that of π^* , which contradicts the optimality of π^* .

In all cases, the existence of such a π' leads to a contradiction. Thus, π^* must be a solution that maximizes social welfare for a given exploitability ϵ . \square

I POLYMATRIX GAMES

In this section, we present experimental details for MASE on games with a larger number of players. We select polymatrix games as the benchmark for these large-scale experiments. This choice is motivated by their inherent graphical structure, which allows for the efficient generation of instances with a low treewidth of their Utility Dependency Graph.

We begin with the formal definition. A polymatrix game has a corresponding undirected graph $\mathcal{G}^U = (\mathcal{V}^U, \mathcal{E}^U)$, with $\mathcal{V}^U = [N]$. For any joint action $\mathbf{a} \in \mathcal{A}$, the utility of any player i is defined as:

$$\mathcal{U}_i(\mathbf{a}) := \sum_{(i,j) \in \mathcal{E}^U} \mathcal{U}_{i,j}(a_i, a_j), \quad (\text{I.1})$$

where $\mathcal{U}_{i,j}: \mathcal{A}_i \times \mathcal{A}_j \rightarrow [0, 1]$ represents the interaction between players i and j . In other words, only players who are connected in \mathcal{G}^U interact, and a player's total utility is the summation of these pairwise interactions.

If we construct the Utility Dependency Graph directly, then the tree decomposition may explode unwillingly, *e.g.*, Figure 8 (a). We can see that the treewidth of \mathcal{G}^U is one while the treewidth of the Utility Dependency Graph is three.

Constructing the Utility Dependency Graph directly from the polymatrix game can cause its treewidth to explode. For example, in Figure 8 (a), the original graph \mathcal{G}^U has a treewidth of one, while the resulting Utility Dependency Graph has a treewidth of three.

To prevent this, we construct a strategically equivalent game (note that this new game is *not* a polymatrix game). This construction explicitly models the pairwise interactions as new players:

- For any original edge $(i, j) \in \mathcal{E}^U$, we introduce two edge players, $e_{i,j}$ and $e_{j,i}$.
- Each edge player $e_{i,j}$ has a singleton action set, $|\mathcal{A}_{e_{i,j}}| = 1$, meaning it has only a single strategy.
- The utility function of an edge player $e_{i,j}$ is defined as the original interaction utility: $\tilde{\mathcal{U}}_{e_{i,j}} = \mathcal{U}_{i,j}$.
- The utility function of an original vertex player i (one of the original N players) is now a constant zero: $\tilde{\mathcal{U}}_i \equiv 0$.

This transformation is illustrated in Figure 8 (b). The Utility Dependency Graph for this new game, shown on the right of Figure 8 (b), now has a treewidth of $\max(\text{tw}(\mathcal{G}^U), 2)$. This method effectively bounds the treewidth and avoids the undesirable explosion.

Next, we show that the new game and the original polymatrix game are strategically equivalent. In other words, for any joint strategy $\pi \in \Delta^{\mathcal{A}}$, the maximum average deviation gain, $\max_{S \in \mathcal{S}} \max_{\hat{\mathbf{a}}_S \in \mathcal{A}_S} \frac{1}{|S|} \sum_{i \in S} \mathbb{E}_{\mathbf{a} \sim \pi} [\mathcal{U}_i(\hat{\mathbf{a}}_S, \mathbf{a}_{-S}) - \mathcal{U}_i(\mathbf{a})]$, does not change. Recall that since the edge players have only a single action, $\pi \in \Delta^{\mathcal{A}}$ (a distribution over the original players' joint actions) is sufficient to specify the joint strategy in both games.

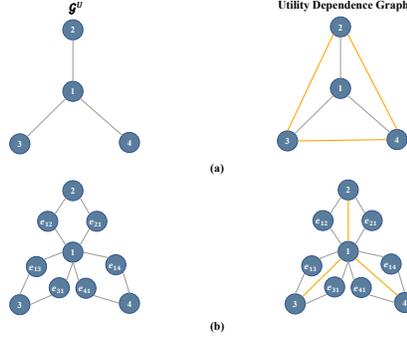


Figure 8: (a) The original graph \mathcal{G}^U corresponding to the polymatrix game (left) and the Utility Dependency Graph. (b) The strategically equivalent game and its Utility Dependency Graph.

Lemma I.1. The new game described above is equivalent to the original polymatrix game, when $\tilde{S} = \left\{ S \cup \{e_{i,j}\}_{i \in S \wedge (i,j) \in \mathcal{E}^U} \right\}_{S \in \mathcal{S}}$. Formally, for any joint strategy $\pi \in \Delta^{\mathcal{A}}$, we have

$$\begin{aligned} & \max_{S \in \mathcal{S}} \max_{\hat{\mathbf{a}}_S \in \mathcal{A}_S} \frac{1}{|S|} \sum_{i \in S} \mathbb{E}_{\mathbf{a} \sim \pi} [\mathcal{U}_i(\hat{\mathbf{a}}_S, \mathbf{a}_{-S}) - \mathcal{U}_i(\mathbf{a})] \\ &= \max_{\tilde{S} \in \tilde{\mathcal{S}}} \max_{\hat{\mathbf{a}}_{\tilde{S}} \in \mathcal{A}_{\tilde{S}}} \frac{1}{|\tilde{S} \cap [N]|} \sum_{i \in \tilde{S}} \mathbb{E}_{\mathbf{a} \sim \pi} [\tilde{\mathcal{U}}_i(\hat{\mathbf{a}}_{\tilde{S} \cap [N]}, \mathbf{a}_{-(\tilde{S} \cap [N])}) - \tilde{\mathcal{U}}_i(\mathbf{a})]. \end{aligned}$$

The proof is postponed to the end of this section. This equivalence allows us to solve the new game instead of the original one. Our algorithm can minimize a more general objective, $\max_{S \in \mathcal{S}} \max_{\hat{\mathbf{a}}_S \in \mathcal{A}_S} \mathbf{w}_S \sum_{i \in S} \mathbb{E}_{\mathbf{a} \sim \pi} [\mathcal{U}_i(\hat{\mathbf{a}}_S, \mathbf{a}_{-S}) - \mathcal{U}_i(\mathbf{a})]$, for any weight vector $\mathbf{w} \in \mathbb{R}^S$.⁹ We can therefore apply our algorithm to this new, strategically equivalent game.

I.1 EXPERIMENTAL DETAILS

We generate random polymatrix games using the following procedure:

- Each pair of players (i, j) is connected independently with probability $\frac{c}{N-1}$, where N is the total number of players. This results in an expected degree of c for each player in \mathcal{G}^U .
- For each connected pair (i, j) , the interaction utilities $\mathcal{U}_{i,j}(a_i, a_j)$ are sampled independently and uniformly from $[0, 1]$ for all action pairs $a_i \in \mathcal{A}_i, a_j \in \mathcal{A}_j$. These pairwise utilities are then normalized according to the formula:

$$\frac{\mathcal{U}_{i,j}(a_i, a_j) - \min_{k \in [N], \hat{\mathbf{a}} \in \mathcal{A}} \mathcal{U}_k(\hat{\mathbf{a}})}{\max_{k \in [N], \hat{\mathbf{a}} \in \mathcal{A}} \mathcal{U}_k(\hat{\mathbf{a}}) - \min_{k \in [N], \hat{\mathbf{a}} \in \mathcal{A}} \mathcal{U}_k(\hat{\mathbf{a}})}.$$

This process ensures that the final total utility $\mathcal{U}_i(\mathbf{a})$ for any player $i \in [N]$ and joint action $\mathbf{a} \in \mathcal{A}$ falls within the range $[0, 1]$.

Consistent with the experiments on small games, we average the results over *100 runs* for each hyper-parameter setting (using seeds 0–99). All algorithms use a learning rate of $\eta = 0.01$, and error bars represent 1σ . For these larger games, we set the number of timesteps to $T = 100,000$ and a uniform action set size $|\mathcal{A}_i| = A$ for all players $i \in [N]$.

The hyper-parameters for the ablation studies are as follows:

- **Ablation on N :** $A = 2$ and $c = 1$.
- **Ablation on A :** $N = 30$ and $c = 1$.

⁹Both the implementation and the proof only use the linearity of the objective. Hence, any weighted-sum can fit into the framework.

- **Ablation on c :** $N = 30$ and $A = 2$.

Furthermore, to accelerate the algorithm, without loss of generality, we only need to consider $\mathcal{S} = \{\{i\}\}_{i \in [N]} \cup \{\{i, j\} \mid i, j \in [N], (i, j) \in \mathcal{G}^U\}$ to minimize the coalition exploitability for any coalitions with no more than two players. In other words, for coalitions of two players, we only need to consider the case when they are connected in \mathcal{G}^U . As shown in the following lemma.

Lemma I.2. For any joint strategy $\pi \in \Delta^{\mathcal{A}}$, by letting $\mathcal{S} = \{\{i\}\}_{i \in [N]} \cup \{\{i, j\} \mid i, j \in [N] \wedge (i, j) \in \mathcal{G}^U\}$, we have

$$\begin{aligned} & \max_{S \in \mathcal{S}} \max_{\hat{\mathbf{a}}_S \in \mathcal{A}_S} \frac{1}{|S|} \sum_{i \in S} \mathbb{E}_{\mathbf{a} \sim \pi} [\mathcal{U}_i(\hat{\mathbf{a}}_S, \mathbf{a}_{-S}) - \mathcal{U}_i(\mathbf{a})] \\ &= \max_{S \in \{\{i\}\}_{i \in [N]} \cup \{\{i, j\} \mid i, j \in [N] \wedge i \neq j\}} \max_{\hat{\mathbf{a}}_S \in \mathcal{A}_S} \frac{1}{|S|} \sum_{i \in S} \mathbb{E}_{\mathbf{a} \sim \pi} [\mathcal{U}_i(\hat{\mathbf{a}}_S, \mathbf{a}_{-S}) - \mathcal{U}_i(\mathbf{a})]. \end{aligned}$$

The proof is postponed to the end of this section.

I.2 PROOF OF THE AUXILIARY LEMMA

Lemma I.1. The new game described above is equivalent to the original polymatrix game, when $\tilde{\mathcal{S}} = \left\{ S \cup \{e_{i,j}\}_{i \in S \wedge (i,j) \in \mathcal{E}^U} \right\}_{S \in \mathcal{S}}$. Formally, for any joint strategy $\pi \in \Delta^{\mathcal{A}}$, we have

$$\begin{aligned} & \max_{S \in \mathcal{S}} \max_{\hat{\mathbf{a}}_S \in \mathcal{A}_S} \frac{1}{|S|} \sum_{i \in S} \mathbb{E}_{\mathbf{a} \sim \pi} [\mathcal{U}_i(\hat{\mathbf{a}}_S, \mathbf{a}_{-S}) - \mathcal{U}_i(\mathbf{a})] \\ &= \max_{\tilde{S} \in \tilde{\mathcal{S}}} \max_{\hat{\mathbf{a}}_{\tilde{S}} \in \mathcal{A}_{\tilde{S}}} \frac{1}{|\tilde{S} \cap [N]|} \sum_{i \in \tilde{S}} \mathbb{E}_{\mathbf{a} \sim \pi} \left[\tilde{\mathcal{U}}_i(\hat{\mathbf{a}}_{\tilde{S} \cap [N]}, \mathbf{a}_{-(\tilde{S} \cap [N])}) - \tilde{\mathcal{U}}_i(\mathbf{a}) \right]. \end{aligned}$$

Proof. For any $S \in \mathcal{S}$, let \tilde{S} be its correspondence in $\tilde{\mathcal{S}}$. Then,

$$\begin{aligned} & \max_{\hat{\mathbf{a}}_S \in \mathcal{A}_S} \frac{1}{|S \cap [N]|} \sum_{i \in \tilde{S}} \mathbb{E}_{\mathbf{a} \sim \pi} \left[\tilde{\mathcal{U}}_i(\hat{\mathbf{a}}_{\tilde{S} \cap [N]}, \mathbf{a}_{-(\tilde{S} \cap [N])}) - \tilde{\mathcal{U}}_i(\mathbf{a}) \right] \\ & \stackrel{(i)}{=} \max_{\hat{\mathbf{a}}_S \in \mathcal{A}_S} \frac{1}{|S|} \sum_{i \in \tilde{S}} \mathbb{E}_{\mathbf{a} \sim \pi} \left[\tilde{\mathcal{U}}_i(\hat{\mathbf{a}}_S, \mathbf{a}_{-S}) - \tilde{\mathcal{U}}_i(\mathbf{a}) \right] \\ & \stackrel{(ii)}{=} \max_{\hat{\mathbf{a}}_S \in \mathcal{A}_S} \frac{1}{|S|} \sum_{i \in S} \sum_{j: (i,j) \in \mathcal{E}^U} \mathbb{E}_{\mathbf{a} \sim \pi} \left[\tilde{\mathcal{U}}_{e_{i,j}}(\hat{\mathbf{a}}_S, \mathbf{a}_{-S}) - \tilde{\mathcal{U}}_{e_{i,j}}(\mathbf{a}) \right] \\ & \stackrel{(iii)}{=} \max_{\hat{\mathbf{a}}_S \in \mathcal{A}_S} \frac{1}{|S|} \sum_{i \in S} \mathbb{E}_{\mathbf{a} \sim \pi} [\mathcal{U}_i(\hat{\mathbf{a}}_S, \mathbf{a}_{-S}) - \mathcal{U}_i(\mathbf{a})]. \end{aligned}$$

(i) uses the fact that $|\mathcal{A}_{e_{i,j}}| = 1$ and $\tilde{S} \cap [N] = S$. (ii) is because $\mathcal{U}_i \equiv 0$ for any $i \in [N]$. (iii) is by the definition of $\tilde{\mathcal{U}}_{e_{i,j}}$ and \tilde{S} . \square

Lemma I.2. For any joint strategy $\pi \in \Delta^{\mathcal{A}}$, by letting $\mathcal{S} = \{\{i\}\}_{i \in [N]} \cup \{\{i, j\} \mid i, j \in [N] \wedge (i, j) \in \mathcal{G}^U\}$, we have

$$\begin{aligned} & \max_{S \in \mathcal{S}} \max_{\hat{\mathbf{a}}_S \in \mathcal{A}_S} \frac{1}{|S|} \sum_{i \in S} \mathbb{E}_{\mathbf{a} \sim \pi} [\mathcal{U}_i(\hat{\mathbf{a}}_S, \mathbf{a}_{-S}) - \mathcal{U}_i(\mathbf{a})] \\ &= \max_{S \in \{\{i\}\}_{i \in [N]} \cup \{\{i, j\} \mid i, j \in [N] \wedge i \neq j\}} \max_{\hat{\mathbf{a}}_S \in \mathcal{A}_S} \frac{1}{|S|} \sum_{i \in S} \mathbb{E}_{\mathbf{a} \sim \pi} [\mathcal{U}_i(\hat{\mathbf{a}}_S, \mathbf{a}_{-S}) - \mathcal{U}_i(\mathbf{a})]. \end{aligned}$$

Proof. For any disconnected players i, j and $S = \{i, j\}$, we can see that

$$\begin{aligned}
& \max_{\hat{\mathbf{a}}_S \in \mathcal{A}_S} \frac{1}{|S|} \sum_{k \in S} \mathbb{E}_{\mathbf{a} \sim \pi} [\mathcal{U}_k(\hat{\mathbf{a}}_S, \mathbf{a}_{-S}) - \mathcal{U}_k(\mathbf{a})] \\
& \leq \max_{\hat{\mathbf{a}}_S \in \mathcal{A}_S} \max_{k \in S} \mathbb{E}_{\mathbf{a} \sim \pi} [\mathcal{U}_k(\hat{\mathbf{a}}_S, \mathbf{a}_{-S}) - \mathcal{U}_k(\mathbf{a})] \\
& = \max_{k \in S} \max_{\hat{\mathbf{a}}_S \in \mathcal{A}_S} \mathbb{E}_{\mathbf{a} \sim \pi} [\mathcal{U}_k(\hat{\mathbf{a}}_S, \mathbf{a}_{-S}) - \mathcal{U}_k(\mathbf{a})] \\
& = \max_{k \in S} \max_{\hat{\mathbf{a}}_S \in \mathcal{A}_S} \mathbb{E}_{\mathbf{a} \sim \pi} \left[\sum_{k' \in [N]: (k, k') \in \mathcal{G}^U} \mathcal{U}_{k, k'}(\hat{\mathbf{a}}_S, \mathbf{a}_{-S}) - \mathcal{U}_{k, k'}(\mathbf{a}) \right] \\
& \stackrel{(i)}{=} \max_{k \in S} \max_{\hat{\mathbf{a}}_S \in \mathcal{A}_S} \mathbb{E}_{\mathbf{a} \sim \pi} \left[\sum_{k' \in [N]: (k, k') \in \mathcal{G}^U} \mathcal{U}_{k, k'}(\hat{\mathbf{a}}_k, \mathbf{a}_{-k}) - \mathcal{U}_{k, k'}(\mathbf{a}) \right] \\
& = \max_{k \in S} \max_{\hat{\mathbf{a}}_S \in \mathcal{A}_S} \mathbb{E}_{\mathbf{a} \sim \pi} [\mathcal{U}_k(\hat{\mathbf{a}}_k, \mathbf{a}_{-k}) - \mathcal{U}_k(\mathbf{a})].
\end{aligned}$$

(i) is because $k' \notin S$ since $k \in S = \{i, j\}$ and i, j are not connected. Therefore, since the coalition exploitability of S is upper bounded by the maximum of that of coalitions $\{i\}$ and $\{j\}$, we do not need to consider $\{i, j\}$.

Actually, the argument can be generalized to coalitions of any size M . If we want to consider the coalition exploitability for coalitions no more than size M , then we only need to consider all connected coalitions of size no more than size M by an induction similar to the proof above. \square