Constrained Phi-Equilibria

Martino Bernasconi [1] **Matteo Castiglioni** [1] **Alberto Marchesi** [1] **Francesco Trovò** [1] **Nicola Gatti** [1]

Abstract

The computational study of equilibria involving constraints on players' strategies has been largely neglected. However, in real-world applications, players are usually subject to constraints ruling out the feasibility of some of their strategies, such as, *e.g.*, safety requirements and budget caps. Computational studies on constrained versions of the Nash equilibrium have lead to some results under very stringent assumptions, while finding constrained versions of the *correlated equilibrium* (CE) is still unexplored. In this paper, we introduce and computationally characterize *constrained Phi-equilibria*—a more general notion than constrained CEs—in *normal-form games*. We show that computing such equilibria is in general computationally intractable, and also that the set of the equilibria may *not* be convex, providing a sharp divide with unconstrained CEs. Nevertheless, we provide a polynomial-time algorithm for computing a constrained (approximate) Phi-equilibrium maximizing a given linear function, when either the number of constraints or that of players' actions is fixed. Moreover, in the special case in which a player's constraints do *not* depend on other players' strategies, we show that an exact, function-maximizing equilibrium can be computed in polynomial time, while *one* (approximate) equilibrium can be found with an efficient decentralized no-regret learning algorithm.

1. Introduction

Over the last years, equilibrium computation problems have received a terrific attention from AI and ML research (Brown & Sandholm, 2019; Bakhtin et al., 2022), as game-theoretical equilibrium notions provide a principled framework to deal with multi-player decision-making problems. Most of the works on equilibrium computation problems focus on classical solution concepts—such as the well-known *Nash equilibrium* (NE) (Nash, 1951) and *correlated equilibrium* (CE) (Aumann, 1974)—, thus neglecting the presence of constraints entirely. However, in most of the real-world applications, the players are usually subject to constraints that rule out the feasibility of some of their strategies, such as, *e.g.*, safety requirements and budget caps. Thus, addressing equilibrium notions involving constraints is a crucial step needed for the operationalization of game-theoretic concepts into real-world settings.

The study of equilibrium notions involving constraints was initiated by Arrow & Debreu (1954), who define the concept of *generalized* NE (GNE). The GNE can be interpreted as an NE of a game where players' strategies are subject to some constraints, which must be satisfied at the equilibrium and also determine which are the feasible players' deviations. However, given that computing a GNE is clearly PPAD-hard (Daskalakis et al., 2009), all the works dealing with the computation of GNEs (see, *e.g.*, (Facchinei & Kanzow, 2010)) provide efficient algorithms only in specific settings that require very stringent assumptions.

Most of the computationally challenges in finding GNEs are inherited from the NE. In settings in which constrained are *not* involved, the computational issues of NEs are usually circumvented by considering weaker equilibrium notions. Among them, those that have received most of the attention in the literature are the CE and its variations, which have been shown to be efficiently computable in several settings of interest (Papadimitriou & Roughgarden, 2008; Celli et al., 2020). Surprisingly, with the only exception of (Chen et al., 2022) (see Section 1.2 for a detailed discussion on it), no work has considered the problem of computing CEs in constrained settings. Thus, investigating whether the CE retains its nice computational properties when adding constraints on players' strategies is an open interesting question.

1.1. Original Contributions

In this paper, we introduce and computationally characterize *constrained Phi-equilibria*, starting, as it is customary, from the setting of *normal-form games*. Our equilibria include the constrained versions of the classical CE and all of its variations as special cases, by generalizing the notion of Phi-

[1]Dipartimento di Elettronica, Informazione e Bioingegneria, Politecnico di Milano, Milan, Italy. Correspondence to: Martino Bernasconi <martino.bernasconideluca@polimi.it>.

Proceedings of the 40^{th} International Conference on Machine Learning, Honolulu, Hawaii, USA. PMLR 202, 2023.

equilibria introduced by Greenwald & Jafari (2003) to constrained settings. In particular, constrained Phi-equilibria are defined as Phi-equilibria, but in games where players are subject to some constraints. Such constraints must be satisfied at the equilibrium, and, additionally, players are only allowed to undertake *safe deviations*, namely those that are feasible according to the constraints. Crucially, the set of safe deviations of a player does *not* only depend on the strategy of that player, but also on those of the others.

We start by showing that one of the most appealing computational properties of Phi-equilibria, namely that the set of the equilibria of a game is convex, is lost when moving to their constrained version. This raises considerable computational challenges in computing constrained Phi-equilibria. Indeed, we formally prove a strong intractability result: for any factor $\alpha > 0$, it is *not* possible, unless $\mathsf{P} = \mathsf{NP}$, to find in polynomial time a constrained (approximate) Phi-equilibrium which achieves a multiplicative approximation α of the optimal value of a given linear function. Then, in the rest of the paper, we show several ways in which such a negative result can be circumvented.

We prove that a constrained approximate Phi-equilibrium which maximizes a given linear function can be found in polynomial time, when either the number of constraints or that of players' actions is fixed. Our results are based on a general algorithm that employs a non-standard "Lagrangification" of the constraints defining the set of safe deviations of a player. Moreover, the algorithm needs a way of dealing with the non-convexity of the set of the equilibria, which we provide in the form of a clever discretization of the space of the Lagrange multipliers.

Finally, we focus on the special case in which the constraints defining the safe deviations of a player do *not* depend on the the strategies of the other players, but only on the strategy of that player. This includes constrained Phi-equilibria identifying a particular constrained version of the *coarse* CE by Moulin & Vial (1978a), in which the players' strategies are subject to *marginal cost constraints*. These arise in several real-world applications in which the players have bounded resources, such as, *e.g.*, budget-constrained bidding in auctions. In such a special case, we prove that a constrained (exact) Phi-equilibrium maximizing a given linear function can be computed in polynomial time, and we provide an efficient decentralized no-regret learning algorithm for finding *one* constrained (approximate) Phi-equilibrium.

1.2. Related Works

GNEs Rosen (1965) initiated the study of the computational properties of GNEs. After that, several other works addressed the problem of computing GNEs by mainly exploiting techniques based on *quasi-variational inequalities* (see (Facchinei & Kanzow, 2010) for a survey). More recently, some works (Kanzow & Steck, 2016; Bueno et al., 2019; Jordan et al., 2022; Goktas & Greenwald, 2022) also studied the convergence of iterative optimization algorithms to GNEs. In order to provide efficient algorithms, all these works need to introduce very stringent assumptions, which are even stronger than those required for the efficient computation of NEs.

Constrained Markov Games Equilibrium notions involving constraints have also been addressed in the literature on *Markov games*, with (Altman & Shwartz, 2000; Alvarez-Mena & Hernández-Lerma, 2006) being the first works introducing GNEs in such a field. More recently, Hakami & Dehghan (2015) defined a notion of constrained CE in Markov games. However, the incentive constraints in their notion of equilibrium only predicate on "pure" deviations, which, in presence of constraints, may lead to empty sets of safe deviations. Very recently, Chen et al. (2022) generalize the work of Hakami & Dehghan (2015) by considering "mixed" deviations. However, their algorithm provides rather weak convergence guarantees, as it only ensures that the returned solution satisfies incentive constraints in expectation. Indeed, as we show in Proposition 3.1, the set of constrained equilibria may *not* be convex (it is easy to see that Example 1 also applies to the setting studied by Chen et al. (2022)), and, thus, the fact that incentive constraints are only satisfied in expectation does *not* necessarily imply that the "true" incentive constraints defining the equilibrium are satisfied. We refer the reader to Appendix A for additional details on these aspects.

2. Preliminaries

In this section, we introduce all the preliminary definitions and results that are needed in the rest of the paper.

2.1. Cost-constrained Normal-form Games

In a *normal-form game*, there is a finite set $\mathcal{N} := \{1, \dots, n\}$ of n players. Each player $i \in \mathcal{N}$ has a finite set $\mathcal{A}_i$ of actions available, with $s := |\mathcal{A}_i|$ for $i \in \mathcal{N}$ being the number of players' actions.[1] We denote by $\boldsymbol{a} \in \mathcal{A} := \bigtimes_{i \in \mathcal{N}} \mathcal{A}_i$ an action profile specifying an action a_i for each player $i \in \mathcal{N}$. Moreover, for $i \in \mathcal{N}$, we let $\boldsymbol{a}_{-i} \in \mathcal{A}_{-i} := \bigtimes_{j \neq i \in \mathcal{N}} \mathcal{A}_i$ be an action profile of all players other than i, while $(a, \boldsymbol{a}_{-i})$ is the action profile obtained by adding $a \in \mathcal{A}_i$ to $\boldsymbol{a}_{-i}$. Finally, we let $u_i : \mathcal{A} \to [0, 1]$ be the utility function of player $i \in \mathcal{N}$, with $u_i(\boldsymbol{a})$ being the utility perceived by that player when the action profile $\boldsymbol{a} \in \mathcal{A}$ is played.

We extend classical normal-form games by considering the

[1] For ease of presentation, in this paper we assume that all the players have the same number of actions. All the results can be easily generalized to the case of different numbers of actions.

case in which each player $i \in \mathcal{N}$ has m_i *cost functions*, namely $c_{i,j} : \mathcal{A} \to [-1,1]$ for $j \in [m_i]$.[2] Each player $i \in \mathcal{N}$ is subject to m_i *constraints*, which require that all player i's costs are less than or equal to zero.[3] For ease of notation, we assume w.l.o.g. that all players have the same number of constraints, namely $m := m_i$ for all $i \in \mathcal{N}$. Moreover, we encode the costs of player $i \in \mathcal{N}$ by a vector-valued function $\boldsymbol{c}_i : \mathcal{A} \to [-1,1]^m$ such that, for every $\boldsymbol{a} \in \mathcal{A}$, the j-th component of the vector $\boldsymbol{c}_i(\boldsymbol{a})$ is $c_{i,j}(\boldsymbol{a})$.

Correlated Strategies In this paper, we deal with solution concepts defined by correlated strategies. A *correlated strategy* $\boldsymbol{z} \in \Delta_{\mathcal{A}}$ is a probability distribution defined over the set of actions profiles, with $\boldsymbol{z}[\boldsymbol{a}]$ denoting the probability assigned to $\boldsymbol{a} \in \mathcal{A}$.[4] With an abuse of notation, for every player $i \in \mathcal{N}$, we let $u_i(\boldsymbol{z})$ be player i's expected utility when the action profile played by the players is drawn from $\boldsymbol{z} \in \Delta_{\mathcal{A}}$. In particular, it holds $u_i(\boldsymbol{z}) := \sum_{\boldsymbol{a} \in \mathcal{A}} u_i(\boldsymbol{a})\boldsymbol{z}[\boldsymbol{a}]$. Similarly, we let $\boldsymbol{c}_i(\boldsymbol{z}) := \sum_{\boldsymbol{a} \in \mathcal{A}} \boldsymbol{c}_i(\boldsymbol{a})\boldsymbol{z}[\boldsymbol{a}]$ be the vector of player i's expected costs, so that player i's constraints can be compactly written as $\boldsymbol{c}_i(\boldsymbol{z}) \preceq \boldsymbol{0}$. Finally, we define $\mathcal{S} \subseteq \Delta_{\mathcal{A}}$ as the set of *safe* correlated strategies, which are those satisfying the cost constraints of all players. Formally:

$$\mathcal{S} := \{\boldsymbol{z} \in \Delta_{\mathcal{A}} \mid \boldsymbol{c}_i(\boldsymbol{z}) \preceq \boldsymbol{0} \quad \forall i \in \mathcal{N}\}.$$

In the following, we assume w.l.o.g. that $\mathcal{S} \neq \varnothing$.

2.2. Constrained Phi-equilibria

We generalize the notion of Phi-equilibria (Greenwald & Jafari, 2003) to cost-constrained normal-form games. Such equilibria are defined as correlated strategies $\boldsymbol{z} \in \Delta_{\mathcal{A}}$ that are robust against a given set Φ of players' deviations, in the sense that, if a mediator draws an action profile $\boldsymbol{a} \in \mathcal{A}$ according to $\boldsymbol{z}$ and recommends to play action a_i to each player $i \in \mathcal{N}$, then no player has an incentive to deviate from their recommendation by selecting a deviation in Φ.

For every $i \in \mathcal{N}$, we let Φ_i be the set of player i's *deviations*, *i.e.*, linear transformations $\phi_i : \mathcal{A}_i \to \Delta_{\mathcal{A}_i}$ that prescribe a probability distribution over player i's actions for every possible action recommendation. For ease of notation, we encode a deviation ϕ_i by means of its matrix representation. Formally, an entry $\phi_i[b,a]$ of the matrix represents the probability assigned to action $a \in \mathcal{A}_i$ by $\phi_i(b)$. We denote the set of all the possible deviations by $\Phi := \{\Phi_i\}_{i \in \mathcal{N}}$.

Given a correlated strategy $\boldsymbol{z} \in \Delta_{\mathcal{A}}$ and a deviation $\phi_i \in \Phi_i$, we define $\phi_i \diamond \boldsymbol{z}$ as the *modification* of $\boldsymbol{z}$ induced by ϕ_i, which is a linear transformation that can be expressed as follows in terms of matrix representation:

$$(\phi_i \diamond \boldsymbol{z})[a_i, \boldsymbol{a}_{-i}] := \sum_{b \in \mathcal{A}_i} \phi_i[b, a_i]\boldsymbol{z}[b, \boldsymbol{a}_{-i}],$$

for every $a_i \in \mathcal{A}_i$ and $\boldsymbol{a}_{-i} \in \mathcal{A}_{-i}$. Moreover, given a set Φ_i of deviations of player $i \in \mathcal{N}$, in the following we denote by $\Phi_i^{\mathcal{S}}(\boldsymbol{z}) := \{\phi_i \in \Phi_i \mid \phi_i \diamond \boldsymbol{z} \in \mathcal{S}\}$ the set of *safe deviations* for player i at a given correlated strategy $\boldsymbol{z} \in \Delta_{\mathcal{A}}$.

We are now ready to provide our definition of *constrained Phi-equilibria* in cost-constrained normal-form games.

Definition 2.1 (Constrained ϵ-Phi-equilibria)**.** Given a set $\Phi := \{\Phi_i\}_{i \in \mathcal{N}}$ of deviations and an $\epsilon > 0$, a *constrained ϵ-Phi-equilibrium* is a safe correlated strategy $\boldsymbol{z} \in \mathcal{S}$ such that, for all $i \in \mathcal{N}$, the following holds:

$$u_i(\boldsymbol{z}) \geq u_i(\phi_i \diamond \boldsymbol{z}) - \epsilon \quad \forall \phi_i \in \Phi_i^{\mathcal{S}}(\boldsymbol{z}).$$

A *constrained Phi-equilibrium* is defined for $\epsilon = 0$.

2.3. Computing Constrained Phi-equilibria

In the following, we formally introduce the computational problem that we study in the rest of the paper.

We denote by $I := (\Gamma, \Phi)$ an instance of the problem, where the tuple $\Gamma := (\mathcal{N}, \mathcal{A}, \{u_i\}_{i \in \mathcal{N}}, \{c_{i,j}\}_{i \in \mathcal{N}, j \in [m]})$ is a cost-constrained normal-form game and $\Phi := \{\Phi_i\}_{i \in \mathcal{N}}$ is a set of deviations. Moreover, we let $|I|$ be the size (in terms of number of bits) of the instance I. We assume that the number n of players is fixed, so that $|I|$ does *not* grow exponentially in n.[5] We also make the following assumption on how the sets of deviations are represented:

Assumption 1. For every $i \in \mathcal{N}$, the set Φ_i is a polytope encoded by a finite of linear inequalities.[6]

Let us remark that, in games without constraints, this assumption is met by all the sets Φ which determine the classical notions of Phi-equilibria (Greenwald & Jafari, 2003).

Next, we formally define our computational problem:

Definition 2.2 (APXCPE(α, ϵ))**.** For any $\alpha, \epsilon > 0$, we define APXCPE(α, ϵ) as the problem of finding, given an instance $I := (\Gamma, \Phi)$ and a linear function $\ell : \Delta_{\mathcal{A}} \to \mathbb{R}$ as input, a constrained ϵ-Phi-equilibrium $\boldsymbol{z} \in \Delta_{\mathcal{A}}$ such that $\ell(\boldsymbol{z}) \geq \alpha \ell(\boldsymbol{z}')$ for all constrained Phi-equilibria $\boldsymbol{z}' \in \Delta_{\mathcal{A}}$.

---

[2] In this paper, given some $x \in \mathbb{N}_{>0}$, we let $[x] := \{1, \dots, x\}$ be the set of the first x natural numbers.

[3] Since $\boldsymbol{z} \in \Delta_{\mathcal{A}}$, we can assume w.l.o.g. that all the constraints are of the form ≤ 0, as any constraint can always be cast in such a form by suitably manipulating the cost function $c_{i,j}$.

[4] In this paper, given a finite set X, we denote by Δ_X the set of all the probability distributions defined over the elements of X.

[5] Notice that the size of the representation of a normal-form game is $O(s^n)$, and, thus, exponential in n. Any algorithm that runs in time polynomial in such instance size is *not* computationally appealing, as even its input has size exponential in n. For this reason, we focus on the case in which n is fixed, and, thus, the instance size does *not* grow exponentially with n.

[6] Notice that, since each $\phi_i \in \Phi_i$ is represented as a matrix, a linear inequality is expressed as $\sum_{b,a \in \mathcal{A}_i} M[b,a]\phi_i[b,a] \leq d$, for some matrix M and scalar d.

Intuitively, $\text{APXCPE}(\alpha, \epsilon)$ asks to compute a constrained ϵ-Phi-equilibrium whose value for the linear function ℓ is at least a fraction α of the maximum value which can be achieved by an (exact) constrained Phi-equilibrium.

In order to ensure that an instance of our problem is well defined, we make the following "Slater-like" assumption on how the players' cost constraints are defined.

Assumption 2. For every correlated strategy $z \in \Delta_{\mathcal{A}}$, player $i \in \mathcal{N}$, and index $j \in [m]$, there exists $\phi_i^\circ \in \Phi_i^{\mathcal{S}}(z)$:

$$c_{i,j}(\phi_i^\circ \diamond z) \leq -\rho,$$

where $\rho > 0$ and $1/\rho$ is $O(\text{poly}(|I|))$, with $\text{poly}(|I|)$ being a polynomial function of the instance size $|I|$.

In Assumption 2, the condition $\rho > 0$ is required to guarantee the existence of a constrained Phi-equilibrium (see Theorem 2.1) and that the sets $\Phi_i^{\mathcal{S}}(z)$ are non-empty (otherwise our solution concept would be ill defined). Moreover, the second condition on ρ in Assumption 2 is equivalent to requiring that our algorithms run in time polynomial in $\frac{1}{\rho}$.

Assumption 2 also allows us to prove the existence of our equilibria, by showing that the constrained Nash equilibria introduced by Altman & Shwartz (2000), which always exist under Assumption 2, are also constrained Phi-equilibria.

Theorem 2.1. *Given a cost-constrained normal-form game Γ and a set Φ of deviations, if Assumption 2 is satisfied, then Γ admits a constrained Phi-equilibrium.*

2.4. Relation with Unconstrained Phi-equilibria

We conclude the section by discussing the relation between our constrained Phi-equilibria and classical equilibrium concepts for unconstrained games.

Correlated Equilibrium When there are no constraints, the *correlated equilibrium* (CE) (Aumann, 1974) is a special case of Phi-equilibrium. As shown by Greenwald & Jafari (2003), the CE is obtained when the sets Φ_i contain all the possible deviations. Formally, the CE is defined by the set $\Phi_{\text{ALL}} := \{\Phi_{i,\text{ALL}}\}$ of deviations such that:

$$\Phi_{i,\text{ALL}} := \left\{ \phi_i \,\middle|\, \sum_{a \in \mathcal{A}_i} \phi_i[b,a] = 1 \quad \forall b \in \mathcal{A}_i \right\}.$$

Coarse Correlated Equilibrium The *coarse correlated equilibrium* (CCE) (Moulin & Vial, 1978b) is a special (unconstrained) Phi-equilibrium whose set of deviations is $\Phi_{\text{CCE}} := \{\Phi_{i,\text{CCE}}\}_{i \in \mathcal{N}}$ such that:

$$\Phi_{i,\text{CCE}} := \left\{ \phi_i \,\middle|\, \exists \boldsymbol{h} \in \Delta_{\mathcal{A}_i} : \phi_i[b,a] = \boldsymbol{h}[a] \ \ \forall b, a \in \mathcal{A}_i \right\}.$$

Intuitively, such sets contain all the possible deviations that prescribe the same probability distribution independently of the received action recommendation.

Thus, our constrained Phi-equilibria include the generalization of CEs and CCEs to cost-constrained games.

Our definition of constrained Phi-equilibrium needs to employ "mixed" deviations that map action recommendations to probability distributions over actions. This is necessary in presence of constraints. Instead, without them, one can simply consider "pure" deviations that map recommendations to actions deterministically (Greenwald & Jafari, 2003).

3. Challenges of Constrained Phi-equilibria

In this section, we show that, in cost-constrained normal-form games, Phi-equilibria loose the nice computational properties that they exhibit in unconstrained settings. This is crucially determined by the fact that the set of constrained Phi-equilibria may *not* be convex in general.

Proposition 3.1. *Given any instance $I := (\Gamma, \Phi)$, the set of constrained Phi-equilibria may not be convex.*

Proposition 3.1 is proved by the following example.

Example 1. *Let Φ_{ALL} be the set of all the possible deviations in a two-player game in which each player has two actions, namely $\mathcal{A}_1 = \mathcal{A}_2 = \{a_0, a_1\}$. The first player's utility is such that $u_1(a, a') = 0$ for all $a \in \mathcal{A}_1$ and $a' \in \mathcal{A}_2$, while the second player's utility is such that $u_2(a_0, a_1) = 1$, and 0 otherwise. Both players share the same single cost constraint ($m = 1$). Their cost functions are defined as $c_i(a_0, a_1) = 1$, $c_i(a_0, a_0) = -\frac{1}{2}$, and $c_i(a_1, a) = -1$ for all $a \in \mathcal{A}_2$. Notice that the instance defined above satisfies Assumption 2 for $\rho = 1/2$. It is easy to see that the correlated strategy $z^1 \in \Delta_{\mathcal{A}}$ such that $z^1[a_0, a_0] = \frac{2}{3}$ and $z^1[a_0, a_1] = \frac{1}{3}$ is a constrained Phi-equilibrium. Moreover, the "pure" correlated strategy $z^2 \in \Delta_{\mathcal{A}}$ such that $z^2[a_1, a_0] = 1$ is also a constrained Phi-equilibrium. However, the combination $z^3 = \frac{1}{2}(z^1 + z^2)$ is not a constrained Phi-equilibrium. Indeed, the second player has an incentive to deviate by using a deviation ϕ_2 such that $\phi_2[a_0, a_1] = 1$ and $\phi_2[a_1, a_1] = 1$. Such a deviation prescribes to play action a_1 when a_0 is recommended, and to play action a_1 when the recommendation is a_1. Straightforward calculations show that, for every $a \in \mathcal{A}_1$:*

$$(\phi_2 \diamond z^3)[a, a'] = \begin{cases} \frac{1}{2} & \text{if} \quad a' = a_1 \\ 0 & \text{otherwise}, \end{cases}$$

and $u_2(\phi_2 \diamond z^3) = \frac{1}{2} > u_2(z^3) = \frac{1}{6}$. Moreover, the deviation is safe, since $\phi_2 \in \Phi_2^{\mathcal{S}}(z^3)$ as $c_2(\phi_2 \diamond z^3) = 0$.

In order to formally asses the computational challenges of computing constrained Phi-equilibria, we prove the following strong inapproximability result:

Theorem 3.1 (Hardness)**.** *For any constant $\alpha > 0$, the problem $\text{APXCPE}(\alpha, (\alpha/s)^2)$ is NP-hard, where s is the number of players' actions in the instance given as input.*

Intuitively, Theorem 3.1 states that, for every multiplicative approximation factor $\alpha > 0$, it is *not* possible to find a constrained ϵ-Phi-equilibrium having value of ℓ at least a fraction α of its optimal value in time polynomial in $\frac{1}{\epsilon}$. Moreover, as a byproduct of Theorem 3.1, we also get the inapproximability up to within any factor of the problem of computing an optimal constrained (exact) Phi-equilibrium.

Notice that the hardness result in Theorem 3.1 cannot hold for values of ϵ that are *independent* from the instance size. Indeed, as we prove in Corollary 4.3 in Section 4, problem $\text{APXCPE}(1, \epsilon)$ can be solved in quasi-polynomial time in the instance size whenever $\epsilon > 0$ is a given constant. Thus, any NP-hardness result for $\text{APXCPE}(\alpha, \epsilon)$ would contradict the *exponential-time hypothesis*.[7]

4. Computing Optimal Constrained ϵ-Phi-equilibria Efficiently

In this section, we show how to circumvent the negative result established by Theorem 3.1. In particular, we prove that, when the number of cost constraints is fixed, problem $\text{APXCPE}(1, \epsilon)$ can be solved in time polynomial in the instance size and $\frac{1}{\epsilon}$ for $\epsilon > 0$ (Corollary 4.2). Moreover, we also prove that, in general, for any constant $\epsilon > 0$ problem $\text{APXCPE}(1, \epsilon)$ admits a quasi-polynomial-time algorithm, whose running time becomes polynomial when the number of players' actions is fixed (Corollary 4.3).

First, in Section 4.1, we provide a general algorithm that is at the core of the two main results of this section. Then, in Section 4.1, we show how the algorithm can be suitably instantiated in order to prove each result. In the rest of this section, unless stated otherwise, we always assume that an $\epsilon > 0$ has been fixed, and that $I := (\Gamma, \Phi)$ and $\ell : \Delta_{\mathcal{A}} \to \mathbb{R}$ are the inputs of a given instance of problem $\text{APXCPE}(1, \epsilon)$.

4.1. General Algorithm

The main technical tool that we employ in order to design our algorithm is a "Lagrangification" of the constraints defining the sets $\Phi_i^{\mathcal{S}}(\boldsymbol{z})$ of safe deviations. First, we prove the following preliminary result, which shows that strong duality holds for the problem $\max_{\phi_i \in \Phi_i^{\mathcal{S}}(\boldsymbol{z})} u_i(\phi_i \diamond \boldsymbol{z})$ of finding the best safe deviation for player $i \in \mathcal{N}$ at $\boldsymbol{z} \in \Delta_{\mathcal{A}}$.

Lemma 4.1. *For every $\boldsymbol{z} \in \Delta_{\mathcal{A}}$ and $i \in \mathcal{N}$, it holds*

$$\begin{aligned}
&\max_{\phi_i \in \Phi_i^{\mathcal{S}}(\boldsymbol{z})} u_i(\phi_i \diamond \boldsymbol{z}) \\
&\qquad = \sup_{\phi_i \in \Phi_i} \inf_{\boldsymbol{\eta}_i \in \mathbb{R}_+^m} \left(u_i(\phi_i \diamond \boldsymbol{z}) - \boldsymbol{\eta}_i^\top \boldsymbol{c}_i(\phi_i \diamond \boldsymbol{z}) \right) \\
&\qquad = \inf_{\boldsymbol{\eta}_i \in \mathbb{R}_+^m} \sup_{\phi_i \in \Phi_i} \left(u_i(\phi_i \diamond \boldsymbol{z}) - \boldsymbol{\eta}_i^\top \boldsymbol{c}_i(\phi_i \diamond \boldsymbol{z}) \right).
\end{aligned}$$

[7]The exponential-time hypothesis conjectures that solving 3SAT requires at least exponential time.

Then, by exploiting Lemma 4.1, we can prove that, under Assumption 2, strong duality continues to hold even when restricting the Lagrange multipliers $\boldsymbol{\eta}_i$ to have ℓ_1-norm less than or equal to $1/\rho$. Formally:

Lemma 4.2. *Let $\mathcal{D} := \left\{ \boldsymbol{\eta} \in \mathbb{R}_+^m \mid ||\boldsymbol{\eta}||_1 \leq 1/\rho \right\}$. Then, for every $\boldsymbol{z} \in \Delta_{\mathcal{A}}$ and $i \in \mathcal{N}$, it holds:*

$$\begin{aligned}
&\max_{\phi_i \in \Phi_i^{\mathcal{S}}(\boldsymbol{z})} u_i(\phi_i \diamond \boldsymbol{z}) \\
&\qquad = \max_{\phi_i \in \Phi_i} \min_{\boldsymbol{\eta}_i \in \mathcal{D}} \left(u_i(\phi_i \diamond \boldsymbol{z}) - \boldsymbol{\eta}_i^\top \boldsymbol{c}_i(\phi_i \diamond \boldsymbol{z}) \right) \\
&\qquad = \min_{\boldsymbol{\eta}_i \in \mathcal{D}} \max_{\phi_i \in \Phi_i} \left(u_i(\phi_i \diamond \boldsymbol{z}) - \boldsymbol{\eta}_i^\top \boldsymbol{c}_i(\phi_i \diamond \boldsymbol{z}) \right).
\end{aligned}$$

Lemma 4.2 allows us to write player i's incentive constraints in the definition of constrained ϵ-Phi-equilibria as

$$u_i(\boldsymbol{z}) \geq \min_{\boldsymbol{\eta}_i \in \mathcal{D}} \max_{\phi_i \in \Phi_i} \left(u_i(\phi_i \diamond \boldsymbol{z}) - \boldsymbol{\eta}_i^\top \boldsymbol{c}_i(\phi_i \diamond \boldsymbol{z}) \right) - \epsilon. \quad (1)$$

This crucially allows us to show the following result: solving problem $\text{APXCPE}(1, \epsilon)$ is equivalent to computing $\max_{(\boldsymbol{\eta}_1, \ldots, \boldsymbol{\eta}_n) \in \mathcal{D}^n} F_\epsilon(\boldsymbol{\eta}_1, \ldots, \boldsymbol{\eta}_n)$, where $F_\epsilon(\boldsymbol{\eta}_1, \ldots, \boldsymbol{\eta}_n)$ is the optimal value of a suitable maximization problem parameterized by tuples of Lagrange multipliers $\boldsymbol{\eta}_i \in \mathcal{D}$, one per player $i \in \mathcal{N}$. Such a problem asks to compute a safe correlated strategy maximizing the linear function ℓ subject to players' incentive constraints that are re-formulated by means of Lemma 4.2. Formally, we define $F_\epsilon(\boldsymbol{\eta}_1, \ldots, \boldsymbol{\eta}_n)$ as the maximum of $\ell(\boldsymbol{z})$ over those $\boldsymbol{z} \in \mathcal{S}$ that additionally satisfy the following constraint for every $i \in \mathcal{N}$:

$$u_i(\boldsymbol{z}) \geq \max_{\phi_i \in \Phi_i} \left(u_i(\phi_i \diamond \boldsymbol{z}) - \boldsymbol{\eta}_i^\top \boldsymbol{c}_i(\phi_i \diamond \boldsymbol{z}) \right) - \epsilon. \quad (2)$$

Notice that the $\min$ operator that appears in the right-hand side of Constraints (1) is dropped by adding the outer maximization over the tuples $(\boldsymbol{\eta}_1, \ldots, \boldsymbol{\eta}_n) \in \mathcal{D}^n$, as the maximum of ℓ is always achieved when the right-hand side of such constraints is as small as possible.

Next, we show that $F_\epsilon(\boldsymbol{\eta}_1, \ldots, \boldsymbol{\eta}_n)$ can be computed in polynomial time by means of the *ellipsoid algorithm*.

Lemma 4.3. *For every tuple $(\boldsymbol{\eta}_1, \ldots, \boldsymbol{\eta}_n) \in \mathcal{D}^n$, the value of $F_\epsilon(\boldsymbol{\eta}_1, \ldots, \boldsymbol{\eta}_n)$ can be computed in time polynomial in the instance size $|I|$ and $\frac{1}{\epsilon}$.*

Proof. We show that $F_\epsilon(\boldsymbol{\eta}_1, \ldots, \boldsymbol{\eta}_n)$ can be solved in polynomial time by means of the ellipsoid algorithm. Let us notice that Constraints (2) can be equivalently encoded by a set of linear inequalities, one for each player $i \in \mathcal{N}$ and deviation $\phi_i \in vert(\Phi_i)$, where $vert(\Phi_i)$ denotes the set of vertexes of the polytope Φ_i (recall Assumption 1). Thus, solving $F_\epsilon(\boldsymbol{\eta}_1, \ldots, \boldsymbol{\eta}_n)$ is equivalent to solving an LP with a (possibly) exponential number of constraints, but polynomially-many variables. Such an LP can be solved in polynomial time by means of the ellipsoid algorithm, provided that a polynomial-time separation oracle for the linearized version of Constraints (2) is available. Such an oracle can be implemented by solving the maximization in the right-hand side of Constraints (2) for a correlated strategy $\boldsymbol{z} \in \Delta_{\mathcal{A}}$ given as input. Formally, the separation oracle solves the following problem for each player $i \in \mathcal{N}$:

$$\phi_i^\star \in \arg\max_{\phi_i \in \Phi_i} \left(u_i(\phi_i \diamond \boldsymbol{z}) - \boldsymbol{\eta}_i^\top \boldsymbol{c}_i(\phi_i \diamond \boldsymbol{z}) \right),$$

which can be done efficiently thanks to Assumption 1. Then, if the separation oracle finds any $\phi_i^\star$ such that:

$$u_i(\boldsymbol{z}) \geq u_i(\phi_i^\star \diamond \boldsymbol{z}) - \boldsymbol{\eta}_i^\top \boldsymbol{c}_i(\phi_i^\star \diamond \boldsymbol{z}),$$

it outputs the above inequality as a separating hyperplane to be used in the ellipsoid algorithm. □

Lemma 4.3 is *not* enough to complete our algorithm, since we need an efficient way of optimizing $F_\epsilon(\boldsymbol{\eta}_1, \dots, \boldsymbol{\eta}_n)$ over all the tuples of Lagrange multipliers. This problem is non-trivial, since $F_\epsilon(\boldsymbol{\eta}_1, \dots, \boldsymbol{\eta}_n)$ is non-concave in $\boldsymbol{\eta}_i$. Nevertheless, we show that, by restricting the domain $\mathcal{D}$ of the Lagrange multipliers to a suitably-defined finite "small" subset, we can still find a constrained ϵ-Phi-equilibrium whose value of ℓ is at least as large as that of any constrained (exact) Phi-equilibrium. This is enough to solve $\textsc{ApxCPE}(1, \epsilon)$. In particular, we need a finite subset of "good" Lagrange multipliers, in the sense of the following definition.

Definition 4.1. Given any $\delta > 0$, a set $\tilde{\mathcal{D}} \subseteq \mathcal{D}$ is δ-*optimal* if, for every $\boldsymbol{z} \in \Delta_{\mathcal{A}}$ and $i \in \mathcal{N}$, the following holds:

$$\min_{\boldsymbol{\eta}_i \in \tilde{\mathcal{D}}} \max_{\phi_i \in \Phi_i} \left(u_i(\phi_i \diamond \boldsymbol{z}) - \boldsymbol{\eta}_i^\top \boldsymbol{c}_i(\phi_i \diamond \boldsymbol{z}) \right) \leq \max_{\phi_i \in \Phi_i^{\mathcal{S}}(\boldsymbol{z})} u_i(\phi_i \diamond \boldsymbol{z}) + \delta.$$

Intuitively, thanks to Lemma 4.2, if we optimize the Lagrange multipliers over a δ-optimal set $\tilde{\mathcal{D}} \subseteq \mathcal{D}$, instead of optimizing them over $\mathcal{D}$, then we are allowing the players to violate incentive constraints by at most δ.

In the following, we assume that a finite δ-optimal set $\tilde{\mathcal{D}} \subseteq \mathcal{D}$ is available. In Section 4.2, se show how to design two particular δ-optimal sets that allow to prove our main results. For ease of presentation, we let

$$L_{\tilde{\mathcal{D}},\epsilon} := \max_{(\boldsymbol{\eta}_1, \dots, \boldsymbol{\eta}_n) \in \tilde{\mathcal{D}}^n} F_\epsilon(\boldsymbol{\eta}_1, \dots, \boldsymbol{\eta}_n)$$

be the optimal value of $F_\epsilon(\boldsymbol{\eta}_1, \dots, \boldsymbol{\eta}_n)$ when the Lagrange multipliers are constrained to be in a δ-optimal set $\tilde{\mathcal{D}} \subseteq \mathcal{D}$. Next, we show that, given any δ-optimal set $\tilde{\mathcal{D}}$ with $\delta \leq \epsilon$, the value of $L_{\tilde{\mathcal{D}},\epsilon}$ is at least that achieved by constrained (exact) Phi-equilibria, namely $L_{\mathcal{D},0}$. Formally:

Lemma 4.4. *Given any $0 < \delta \leq \epsilon$ and a δ-optimal set $\tilde{\mathcal{D}} \subseteq \mathcal{D}$, the following holds: $L_{\tilde{\mathcal{D}},\epsilon} \geq L_{\mathcal{D},0}$.*

Intuitively, Lemma 4.4 is proved by noticing that, provided that $\delta \leq \epsilon$, the incentive constraints violation introduced by using $\tilde{\mathcal{D}}$ instead of $\mathcal{D}$ is at most ϵ. Moreover, the set of feasible correlated strategies can only expand by allowing incentive constraints to be violated, and, thus, the value of the objective ℓ can only increase.

Lemma 4.4 suggests a way of solving $\textsc{ApxCPE}(1, \epsilon)$. Indeed, given a finite δ-optimal set $\tilde{\mathcal{D}} \subseteq \mathcal{D}$ with $\delta \leq \epsilon$, by enumerating over all the tuples of Lagrange multipliers $\boldsymbol{\eta}_i \in \tilde{\mathcal{D}}$, one per player $i \in \mathcal{N}$, we can find the desired constrained ϵ-Phi-equilibrium. The following theorem shows that this procedure gives an algorithm for $\textsc{ApxCPE}(1, \epsilon)$ that runs in time polynomial in the instance size, $|\tilde{\mathcal{D}}|$, and $\frac{1}{\epsilon}$.

Theorem 4.1. *Given a finite δ-optimal set $\tilde{\mathcal{D}} \subseteq \mathcal{D}$ with $\delta \leq \epsilon$, there exists an algorithm that solves* $\textsc{ApxCPE}(1, \epsilon)$ *and runs in time polynomial in the instance size $|I|$, the number $|\tilde{\mathcal{D}}|$ of elements in $\tilde{\mathcal{D}}$, and $\frac{1}{\epsilon}$ for every $\epsilon > 0$.*

Proof. The algorithm works by enumerating over all the possible tuples of Lagrange multipliers $\boldsymbol{\eta}_i \in \tilde{\mathcal{D}}$, one per player $i \in \mathcal{N}$. These are polynomially many in the size $|\tilde{\mathcal{D}}|$ when the number of players n is fixed. For every tuple $(\boldsymbol{\eta}_1, \dots, \boldsymbol{\eta}_n) \in \tilde{\mathcal{D}}^n$, the algorithm solves $F_\epsilon(\boldsymbol{\eta}_1, \dots, \boldsymbol{\eta}_n)$, which can be done in time polynomial in $|I|$ and $\frac{1}{\epsilon}$ thanks to Lemma 4.3. Finally, the algorithm returns the correlated strategy $\boldsymbol{z} \in \Delta_{\mathcal{A}}$ with the highest value of ℓ among those computed while solving $F_\epsilon(\boldsymbol{\eta}_1, \dots, \boldsymbol{\eta}_n)$. It is easy to see that the returned solution solves problem $\textsc{ApxCPE}(1, \epsilon)$ by applying Lemma 4.4. This concludes the proof. □

4.2. Instantiating the General Algorithm

Next, we show how to build δ-optimal sets $\tilde{\mathcal{D}}$ that, when they are plugged in the algorithm in Theorem 4.1, allow us to derive our results. In particular, we consider the set:

$$\mathcal{D}_\tau := \left\{ \boldsymbol{\eta} \in \mathcal{D} \,\middle|\, \eta_j = k\tau, k \in \{0, \dots, \lfloor 1/\tau\rho \rfloor\} \; \forall j \in [m] \right\},$$

which is a discretization of $\mathcal{D}$ with a regular lattice of step $\tau \in \mathbb{R}_+$ (notice that η_j is the j-th component of $\boldsymbol{\eta}$). By a simple stars and bars combinatorial argument, we have that $|\mathcal{D}_\tau| = \binom{\lfloor 1/\tau\rho \rfloor + m}{m}$. Thus, since it holds that $|\mathcal{D}_\tau| = O((1/\tau\rho)^m)$, if the number of constraints m is fixed, $|\mathcal{D}_\tau|$ is bounded by a polynomial in $1/\tau\rho$. Moreover, simple combinatorial arguments show that $|\mathcal{D}_\tau| \leq (1+m)^{\lfloor 1/\tau\rho \rfloor}$.[8] Thus, it also holds that $|\mathcal{D}_\tau| = O(m^{1/\tau\rho})$. Notice that the two bounds on $|\mathcal{D}_\tau|$ are non-comparable, and, thus, they give rise to two distinct results, as we show in the following.

[8] See Appendix D for a formal proof.

By using the first bound on $|\mathcal{D}_\tau|$, we can show that the set $\mathcal{D}_\tau$ is δ-optimal for $\delta = m\tau$. Formally:

Lemma 4.5. *For any $\tau > 0$, the set $\mathcal{D}_\tau$ is (τm)-optimal.*

Thus, whenever the number m of cost constraints is fixed, Lemma 4.5, together with Theorem 4.1, allows us to provide a polynomial-time algorithm. Indeed, it is sufficient to apply Theorem 4.1 for the (τm)-optimal set $\mathcal{D}_\tau$ with $\tau := \epsilon/m$ to obtain the following first main result:

Corollary 4.2. *There exists an algorithm that solves problem $\text{APXCPE}(1, \epsilon)$ in time polynomial in $|I|$ and $\frac{1}{\epsilon}$ for every $\epsilon > 0$, when the number m of cost constraints is fixed.*

On the other hand, by using the second bound on $|\mathcal{D}_\tau|$, we can show that $\mathcal{D}_\tau$ is δ-optimal for δ depending logarithmically on the number of players' actions. Formally:

Lemma 4.6. *For any $\tau > 0$, the set $\mathcal{D}_\tau$ is δ-optimal for $\delta = 2\sqrt{2\tau \log s/\rho}$, where s is the number of players' actions.*

Lemma 4.6 (together with Theorem 4.1) immediately gives us a quasi-polynomial-time for solving $\text{APXCPE}(1, \epsilon)$ for a given constant $\epsilon > 0$. Moreover, its running time becomes polynomial when the number of players' actions is fixed.

Corollary 4.3. *For any constant $\epsilon > 0$, there exists an algorithm that solves $\text{APXCPE}(1, \epsilon)$ in time $O(|I|^{\log s})$. Moreover, when the number s of players' actions is fixed, the algorithm runs in time polynomial in $|I|$.*

Notice that it is in general *not* possible to design an algorithm that runs in time polynomial in $\frac{1}{\epsilon}$, since this would contradict the hardness result in Theorem 3.1.

5. A Special Case: Deviation-dependent Costs

We complete our computational study of constrained Phi-equilibria by considering a special case in which player i's costs associated to a deviation ϕ_i only depend on ϕ_i and *not* on the (overall) modified correlated strategy $\phi_i \diamond \boldsymbol{z}$.

We consider instances satisfying the following assumption.

Assumption 3. For every player $i \in \mathcal{N}$ and player i's deviation $\phi_i \in \Phi_i$, there exists a function $\tilde{\boldsymbol{c}}_i : \Phi_i \to [-1, 1]^m$ such that $\tilde{\boldsymbol{c}}_i(\phi_i) := \boldsymbol{c}_i(\phi_i \diamond \boldsymbol{z})$ for every $\boldsymbol{z} \in \Delta_{\mathcal{A}}$.

Notice that, whenever Assumption 3 holds, the set $\Phi_i^{\mathcal{S}}(\boldsymbol{z})$ of safe deviations does *not* depend on $\boldsymbol{z}$. Thus, in the rest of this section, we write w.l.o.g. $\Phi_i^{\mathcal{S}}$ rather than $\Phi_i^{\mathcal{S}}(\boldsymbol{z})$.

A positive effect of Assumption 3 is that it recovers the convexity of the set of constrained Phi-equilibria, rendering them more akin to unconstrained ones. Formally:

Proposition 5.1. *For instances $I := (\Gamma, \Phi)$ satisfying Assumption 3, the set of constrained ϵ-Phi-equilibria is convex.*

Proposition 5.1 suggests that constrained Phi-equilibria are much more computationally appealing under Assumption 3 than in general, as we indeed show in the rest of this section.

First, in Section 5.1, we show that $\text{APXCPE}(1, 0)$ admits a polynomial-time algorithm under Assumption 3. Then, in Section 5.2, we design a no-regret learning algorithm that efficiently computes *one* constrained ϵ-Phi equilibrium with $\epsilon = O(1/\sqrt{T})$ as the number of rounds T grows. Finally, in Section 5.3, we provide a natural example of constrained Phi-equilibria satisfying Assumption 3.

5.1. A Poly-time Algorithm for Optimal Equilibria

We prove that, whenever Assumption 3 holds, the problem of computing an (exact) Phi-equilibrium maximizing a given linear function can be solved in polynomial time. This is done by formulating the problem as an LP with polynomially-many variables and exponentially-many constraints, which can be solved by means of the ellipsoid method, similarly to how we compute $F_\epsilon(\boldsymbol{\eta}_1, \dots \boldsymbol{\eta}_n)$ in Section 4 (see the proof of Lemma 4.3). Formally:

Theorem 5.1. *Restricted to instances $I := (\Gamma, \Phi)$ which satisfy Assumption 3, $\text{APXCPE}(1, 0)$ admits a polynomial-time algorithm.*

5.2. An Efficient No-regret Learning Algorithm

Next, we show how Assumption 3 allows us to find a constrained ϵ-Phi-equilibrium by means of a polynomial-time decentralized no-regret learning algorithm. Our algorithm is based on the Phi-regret minimization framework introduced by Greenwald & Jafari (2003), which needs to be extended in order to be able to work with polytopal sets $\Phi_i^{\mathcal{S}}$ of safe deviations, rather than finite sets of "pure" deviations.

Algorithm 1 Learning a Constrained ϵ-Phi-equilibria

Require: Regret minimizers $\mathfrak{R}_i$ for the sets $\Phi_i^{\mathcal{S}}$, for $i \in \mathcal{N}$

Initialize the regret minimizers $\mathfrak{R}_i$
for $t = 1, \dots, T$ **do**
  for each player $i \in \mathcal{N}$ **do**
    $\phi_{i,t} \leftarrow \mathfrak{R}_i.\text{RECOMMEND}()$
    Play according to a distribution $\boldsymbol{x}_{i,t} \in \Delta_{\mathcal{A}_i}$ s.t.

$$\boldsymbol{x}_{i,t}[a] = \sum_{b \in \mathcal{A}_i} \phi_{i,t}[b, a] \boldsymbol{x}_{i,t}[b] \quad \forall a \in \mathcal{A}_i$$

  end for
  $\boldsymbol{z}_t \leftarrow \otimes_{i \in \mathcal{N}} \boldsymbol{x}_{i,t}$
  $\mathfrak{R}_i.\text{OBSERVE}(\phi_i \mapsto u_i(\phi_i \diamond \boldsymbol{z}_t))$
end for
return $\bar{\boldsymbol{z}}_T := \frac{1}{T} \sum_{t=1}^{T} \boldsymbol{z}_t$

Algorithm 1 outlines our no-regret algorithm. It instantiates a regret minimizer $\mathfrak{R}_i$ for the polytope $\Phi_i^{\mathcal{S}}$ for each $i \in \mathcal{N}$. $\mathfrak{R}_i$ is an object that, at each round $t \in [T]$, recommends a safe deviation $\phi_{i,t} \in \Phi_i^{\mathcal{S}}$ to player i (Line 4 of Algorithm 1), and, then, observes a function $\phi_i \mapsto u_i(\phi_i \diamond \boldsymbol{z}_t)$ that specifies the utility that would have been obtained by selecting any safe deviation $\phi_i \in \Phi_i^{\mathcal{S}}$ at round t (Line 8 of Algorithm 1). $\mathfrak{R}_i$ guarantees that the regret R_i^T cumulated by player i over $[T]$ grows sublinearly, *i.e.,* $R_i^T = o(T)$, where:

$$R_i^T := \max_{\phi_i \in \Phi_i} \sum_{t=1}^{T} u_i(\phi_i \diamond \boldsymbol{z}_t) - \sum_{t=1}^{T} u_i(\phi_{i,t} \diamond \boldsymbol{z}_t),$$

which is how much player i loses by selecting $\phi_{i,t}$ at each t rather than choosing the same best-in-hindsight deviation at all rounds. Notice that, by taking inspiration from the Phi-regret framework (Greenwald & Jafari, 2003), given a recommended deviation $\phi_{i,t}$, player i actually plays according to a probability distribution $\boldsymbol{x}_{i,t} \in \Delta_{\mathcal{A}_i}$, which is a stationary distribution of the matrix representing $\phi_{i,t}$. This is crucial in order to implement the algorithm in a decentralized fashion and to provide convergence guarantees to constrained ϵ-Phi-equilibria (see Theorem 5.2). All the distributions $\boldsymbol{x}_{i,t}$ jointly determine a correlated strategy $\boldsymbol{z}_t \in \Delta_{\mathcal{A}}$ at each round $t \in [T]$, defined as $\boldsymbol{z}_t := \otimes_{i \in \mathcal{N}} \boldsymbol{x}_{i,t}$, where $\otimes$ denotes the product among distributions; formally, $\boldsymbol{z}_t[\boldsymbol{a}] := \prod_{i \in \mathcal{N}} \boldsymbol{x}_{i,t}[a_i]$ for all $\boldsymbol{a} \in \mathcal{A}$.

Algorithm 1 provides the following guarantees:

Theorem 5.2. *Given an instance $I := (\Gamma, \Phi)$ satisfying Assumption 3, after $T \in \mathbb{N}_{>0}$ rounds, Algorithm 1 returns a correlated strategy $\bar{\boldsymbol{z}}_T \in \Delta_{\mathcal{A}}$ that is a constrained ϵ_T-Phi-equilibrium with $\epsilon_T = O(1/\sqrt{T})$. Moreover, each round of Algorithm 1 runs in polynomial time.*

Let us remark that the crucial property which allows us to design Algorithm 1 is that the sets $\Phi_i^{\mathcal{S}}$ of safe deviations do *not* depend on players other than i. Finally, from Theorem 5.2, the following result follows:

Corollary 5.3. *In instances $I := (\Gamma, \Phi)$ satisfying Assumption 3, a constrained ϵ-Phi-equilibrium can be computed in time polynomial in the instance size and $\frac{1}{\epsilon}$ by means of a decentralized learning algorithm.*

5.3. Marginally-constrained CCE

We conclude the section by introducing a particular (natural) notion of constrained ϵ-Phi-equilibrium for which Assumption 3 is satisfied. This is a constrained version of the classical CCE in cost-constrained normal-form games where a player's costs *only* depend on the action of that player. We call it *marginally-constrained* ϵ-CCE. Formally, such an equilibrium is defined for games in which, for every player $i \in \mathcal{N}$, it holds $\boldsymbol{c}_i(\boldsymbol{a}) = \boldsymbol{c}_i(\boldsymbol{a}')$ for all $\boldsymbol{a}, \boldsymbol{a}' \in \mathcal{A}$ such that $a_i = a'_i$, and for the set Φ_{CCE} of CCE deviations that we have previously introduced in Section 2.4. Next, we prove that, with the definition above, Assumption 3 is satisfied.

Theorem 5.4. *For instances $I := (\Gamma, \Phi_{\text{CCE}})$ such that $\boldsymbol{c}_i(\boldsymbol{a}) = \boldsymbol{c}_i(\boldsymbol{a}')$ for every player $i \in \mathcal{N}$ and action profiles $\boldsymbol{a}, \boldsymbol{a}' \in \mathcal{A} : a_i = a'_i$, Assumption 3 holds.*

Thanks to Theorem 5.4, we readily obtain the two following corollaries of Theorems 5.1 and 5.1.

Corollary 5.5. *The problem of computing a marginally-constrained (exact) CCE that maximizes a linear function $\ell : \Delta_{\mathcal{A}} \to \mathbb{R}$ can be solved in polynomial time.*

Corollary 5.6. *A marginally-constrained ϵ-CCE can be computed in time polynomial in the instance size and $\frac{1}{\epsilon}$ by means of a decentralized learning algorithm.*

6. Discussion and Open Problems

The main positive results that we provide in this paper (Corollaries 4.2 and 4.3) show that a constrained ϵ-Phi equilibrium maximizing a given linear function can be computed in time polynomial in the instance size and $\frac{1}{\epsilon}$, when either the number of constraints or that of players' actions is fixed. Clearly, this implies that, under the same assumptions, *a* constrained ϵ-Phi-equilibrium can be found efficiently. Moreover, in Section 5, we designed an efficient no-regret learning algorithm that finds *a* constrained ϵ-Phi-equilibrium in settings enjoying special properties (Corollary 5.3). However, the problem of efficiently computing *a* constrained ϵ-Phi-equilibrium remains open in general. Formally:

Definition 6.1 (Open Problem)**.** Given any instance $I := (\Gamma, \Phi)$, find a constrained ϵ-Phi-equilibrium in time polynomial in the instance size and $\frac{1}{\epsilon}$.

Solving the problem above is non-trivial. Proposition 3.1 in Section 3 proves that the set of constrained ϵ-Phi-equilibria is non-convex, and, thus, solving the problem in Definition 6.1 is out of scope for most of the known equilibrium computation techniques. On the other hand, it is unlikely that such a problem is NP-hard. Indeed, a constrained ϵ-Phi-equilibrium always exists and, given any $\boldsymbol{z} \in \Delta_{\mathcal{A}}$, it is possible to verify whether $\boldsymbol{z}$ is an equilibrium or not in polynomial time. Formally, such a problem is said to belong to the TFNP complexity class, and, thus, standard arguments show that, if the problem is NP-hard, then NP = coNP (Megiddo & Papadimitriou, 1991). Thus, one should try to reduce the problem in Definition 6.1 to problems in TFNP, such as that of computing a Nash equilibrium. However, while the problem in Definition 6.1 shares some properties with that of computing a Nash equilibrium, such as the non-convexity of the set of the equilibria, the former is fundamentally different from the latter, since it exhibits correlation among the players. Thus, a reduction from such a problem to that of computing Nash equilibria would require a gadget to break the correlation among the players, and doing that is highly non-trivial as cost constraints are expressed by linear functions.

Acknowledgements

This paper is supported by FAIR (Future Artificial Intelligence Research) project, funded by the NextGenerationEU program within the PNRR-PE-AI scheme (M4C2, Investment 1.3, Line on Artificial Intelligence).

References

Altman, E. and Shwartz, A. Constrained markov games: Nash equilibria. In *Advances in dynamic games and applications*, pp. 213–221. Springer, 2000.

Alvarez-Mena, J. and Hernández-Lerma, O. Existence of nash equilibria for constrained stochastic games. *Mathematical Methods of Operations Research*, 63(2):261–285, 2006.

Arrow, K. J. and Debreu, G. Existence of an equilibrium for a competitive economy. *Econometrica: Journal of the Econometric Society*, pp. 265–290, 1954.

Aumann, R. J. Subjectivity and correlation in randomized strategies. *Journal of mathematical Economics*, 1(1): 67–96, 1974.

Bakhtin, A., Brown, N., Dinan, E., Farina, G., Flaherty, C., Fried, D., Goff, A., Gray, J., Hu, H., Jacob, A., et al. Human-level play in the game of diplomacy by combining language models with strategic reasoning. *Science (New York, NY)*, pp. eade9097–eade9097, 2022.

Brown, N. and Sandholm, T. Superhuman ai for multiplayer poker. *Science*, 365(6456):885–890, 2019.

Bueno, L. F., Haeser, G., and Rojas, F. N. Optimality conditions and constraint qualifications for generalized nash equilibrium problems and their practical implications. *SIAM Journal on Optimization*, 29(1):31–54, 2019.

Celli, A., Marchesi, A., Farina, G., and Gatti, N. No-regret learning dynamics for extensive-form correlated equilibrium. *Advances in Neural Information Processing Systems*, 33:7722–7732, 2020.

Chen, Z., Ma, S., and Zhou, Y. Finding correlated equilibrium of constrained markov game: A primal-dual approach. In *Advances in Neural Information Processing Systems*, 2022.

Daskalakis, C., Goldberg, P. W., and Papadimitriou, C. H. The complexity of computing a nash equilibrium. *SIAM Journal on Computing*, 39(1):195–259, 2009.

Ekeland, I. and Temam, R. *Convex analysis and variational problems*. SIAM, 1999.

Facchinei, F. and Kanzow, C. Generalized nash equilibrium problems. *Annals of Operations Research*, 175(1):177–211, 2010.

Goktas, D. and Greenwald, A. Exploitability minimization in games and beyond. *Advances in Neural Information Processing Systems*, 2022.

Greenwald, A. and Jafari, A. A general class of no-regret learning algorithms and game-theoretic equilibria. In *Learning theory and kernel machines*, pp. 2–12. Springer, 2003.

Hakami, V. and Dehghan, M. Learning stationary correlated equilibria in constrained general-sum stochastic games. *IEEE Transactions on Cybernetics*, 46(7):1640–1654, 2015.

Håstad, J. Clique is hard to approximate within $n^{1-\epsilon}$. *Acta Mathematica*, 182(1):105–142, 1999.

Hazan, E. et al. Introduction to online convex optimization. *Foundations and Trends® in Optimization*, 2(3-4):157–325, 2016.

Jordan, M. I., Lin, T., and Zampetakis, M. First-order algorithms for nonlinear generalized nash equilibrium problems. *arXiv preprint arXiv:2204.03132*, 2022.

Kanzow, C. and Steck, D. Augmented lagrangian methods for the solution of generalized nash equilibrium problems. *SIAM Journal on Optimization*, 26(4):2034–2058, 2016.

Megiddo, N. and Papadimitriou, C. H. On total functions, existence theorems and computational complexity. *Theoretical Computer Science*, 81(2):317–324, 1991.

Moulin, H. and Vial, J.-P. Strategically zero-sum games: the class of games whose completely mixed equilibria cannot be improved upon. *INT J GAME THEORY*, 7(3): 201–221, 1978a.

Moulin, H. and Vial, J.-P. Strategically zero-sum games: the class of games whose completely mixed equilibria cannot be improved upon. *International Journal of Game Theory*, 7(3):201–221, 1978b.

Nash, J. Non-cooperative games. *Annals of mathematics*, pp. 286–295, 1951.

Papadimitriou, C. H. and Roughgarden, T. Computing correlated equilibria in multi-player games. *Journal of the ACM (JACM)*, 55(3):1–29, 2008.

Rosen, J. B. Existence and uniqueness of equilibrium points for concave n-person games. *Econometrica: Journal of the Econometric Society*, pp. 520–534, 1965.

Zuckerman, D. Linear degree extractors and the inapproximability of max clique and chromatic number. *Theory of Computing*, 3(6):103–128, 2007.

A. On the Weaknesses of the Guarantees of the Algorithm of Chen et al. (2022)

The Algorithm of Chen et al. (2022) finds a distribution $\mu \in \Delta_{\Delta_\mathcal{A}}$ over correlated strategies such that:

$$\mathbb{E}_{\boldsymbol{z}\sim\mu}\left[\max_{\phi_i \in \Phi_i^{\mathcal{S}}(\boldsymbol{z})} u_i(\phi_i \diamond \boldsymbol{z}) - u_i(\boldsymbol{z})\right] \leq 0. \tag{3}$$

However, here we claim that this solution concept inherits some weaknesses from the non-convexity of the equilibria set that we proved in Theorem 5.1. Indeed, consider the same instance of Theorem 5.1 and consider the uniform distribution μ over $\{\boldsymbol{z}_1, \boldsymbol{z}_2\}$. In Theorem 5.1 we proved that $\max_{\phi_i \in \Phi_i^{\mathcal{S}}(\boldsymbol{z}^1)} u_i(\phi_i \diamond \boldsymbol{z}^1) - u_i(\boldsymbol{z}^1) \leq 0$ for all $i \in \{1, 2\}$ and $\max_{\phi_i \in \Phi_i^{\mathcal{S}}(\boldsymbol{z}^2)} u_i(\phi_i \diamond \boldsymbol{z}^2) - u_i(\boldsymbol{z}^2) \leq 0$ for all $i \in \{1, 2\}$ and thus Equation (3) holds over the distribution μ.

However we show that the expected correlated strategy $\boldsymbol{z}^3$ derived from distribution μ, *i.e.,* $\boldsymbol{z}^3 = \mathbb{E}_{\boldsymbol{z}\sim\mu}[\boldsymbol{z}] = \frac{1}{2}\boldsymbol{z}^1 + \frac{1}{2}\boldsymbol{z}^2$, it is not a feasible equilibrium, or an approximate one.

Indeed, in Theorem 5.1, we proved that $\max_{\phi_2 \in \Phi_2^{\mathcal{S}}(\boldsymbol{z}^3)} u_2(\phi_2 \diamond \boldsymbol{z}^3) - u_2(\boldsymbol{z}^3) \geq \frac{1}{3}$, showing that the average correlated strategies returned by their Algorithm is not an equilibrium nor close to it.

This comes from the peculiar fact about Constrained Phi-equilibria that exhibit non-convex set of solutions, which is in striking contrast with the unconstrained case. Indeed the guarantees of Equilibria (3) would imply that $\mathbb{E}_{\boldsymbol{z}\sim\mu}[\boldsymbol{z}]$ is a equilibrium in the unconstrained case in which the set of equilibria is convex.

B. Proofs Omitted from Section 2

Theorem 2.1. *Given a cost-constrained normal-form game Γ and a set Φ of deviations, if Assumption 2 is satisfied, then Γ admits a constrained Phi-equilibrium.*

Proof. With assumption 2 Altman & Shwartz (2000, Theorem 2.1) proves the existence of a constrained Nash equilibrium. In our setting this is equivalent to a product distribution $\boldsymbol{z} = \otimes_{i\in[N]}\boldsymbol{x}_i$ so that it is a Constrained Phi-equilibrium for any set of deviations Φ_i.[9] This is easily seen by observing that a Constrained Nash Equilibria is defined as:

$$\sum_{\boldsymbol{a}\in\mathcal{A}} u_i\left(\prod_{j\in[N]} \boldsymbol{x}_j(a_j)\right) \geq \sum_{\boldsymbol{a}\in\mathcal{A}} u_i\left(\tilde{\boldsymbol{x}}_i(a_j) \prod_{j\in[N]\setminus\{i\}} \boldsymbol{x}_i(a_i)\right)$$

for all $\tilde{\boldsymbol{x}}_i \in \Delta(\mathcal{A}^i)$ s.t. $\boldsymbol{x}_i \otimes \boldsymbol{x}_{-i} \in \mathcal{S}$.

On the other hand it easily seen that for all $\phi_i \in \Phi_i(\boldsymbol{z})$ there exists some $\tilde{\boldsymbol{x}}_i \in \Delta(\mathcal{A}^i)$ such that

$$\phi_i \diamond \left(\otimes_{j\in[N]}\boldsymbol{x}_j\right) = \tilde{\boldsymbol{x}}_i \otimes \boldsymbol{x}_{-i}$$

and $\tilde{\boldsymbol{x}}_i \otimes \boldsymbol{x}_{-i} \in \mathcal{S}$.

This is proved by the following calculations:

$$\phi_i \diamond \left(\otimes_{j\in[N]}\boldsymbol{x}_j\right)[a_i, \boldsymbol{a}_{-i}] := \sum_{b\in\mathcal{A}^i} \phi_i[b, a_i]\boldsymbol{x}_i(b)\boldsymbol{x}_{-i}(\boldsymbol{a}_{-i}) \tag{4}$$

$$= \tilde{\boldsymbol{x}}_i(a_i) \otimes \boldsymbol{x}_{-i}(\boldsymbol{a}_{-i}), \tag{5}$$

where $\tilde{\boldsymbol{x}}_i(a_i) := \sum_{b\in\mathcal{A}^i} \phi_i[b, a_i]\boldsymbol{x}_i(b)$ and $\tilde{\boldsymbol{x}}_i \in \Delta(A^i)$ since, by definition, $\sum_{a_i\in\mathcal{A}^i} \phi_i[b, a_i] = 1$ for all $b \in \mathcal{A}^i$.

This proves that a Constrained Nash Equilibrium is a Phi-Constrained Equilibrium for all Φ. □

C. Proofs Omitted from Section 3

Theorem 3.1 (Hardness)**.** *For any constant $\alpha > 0$, the problem* $\text{ApxCPE}(\alpha, (\alpha/s)^2)$ *is* NP*-hard, where s is the number of players' actions in the instance given as input.*

[9] As common in the normal form game literature, for any distribution $\boldsymbol{x} \in \Delta(X)$ and $\boldsymbol{y} \in \Delta(Y)$, $\boldsymbol{x} \otimes \boldsymbol{y} \in \Delta(X \times Y)$ is the product distribution defined as $(\boldsymbol{x} \otimes \boldsymbol{y})[a, b] = \boldsymbol{x}[a]\boldsymbol{y}[b]$ for $a \in X$ and $b \in Y$.

Proof. We reduce from GAP-INDEPENDENT-SET, which is a promise problem that formally reads as follows: given an $\delta > 0$ and a graph $G = (V, E)$, with set of nodes V and set of edges E, determine whether G admits an independent set of size at least $|V|^{1-\delta}$ or all the independent sets of G have size smaller than $|V|^{\delta}$. GAP-INDEPENDENT-SET is NP-hard for every $\delta > 0$ (Håstad, 1999; Zuckerman, 2007).

Let $\ell = |V|$ and $\alpha > 0$ be the desired approximation factor. Given an instance of GAP-INDEPENDENT-SET, we build an instance such that if there exists an independent set of size $\ell^{1-\delta}$, then there exists a Constrained Phi-equilibrium with social welfare 1. Otherwise, if all the independent sets have size at most ℓ^{δ}, all the Constrained ϵ-Phi-equilibria have social welfare at most $\alpha/2$. We can use any $\delta > 0$, since we simply need $\ell^{\delta} < \ell^{1-\delta}$. Moreover, we take $\epsilon = \frac{\alpha^2}{128\ell^2}$. As we will see, ℓ will be smaller than the number of action of the players, satisfying the condition in the statement.

Construction. The first player has a set of actions $\mathcal{A}_1$ that includes actions a_0, a_1, a_2 and an action a_v for each $v \in V$. Moreover, the first player has an action a_F.[10] The second player has a set of actions $\mathcal{A}_2$ that includes actions a_v and $\bar{a}_v$ for each $v \in V$. Moreover, the second player has an action a_F. Let $\gamma = \eta = \alpha/8$. The utility of the first agent is as follows:

- $u_1(a_0, a) = \gamma + \frac{1}{2}\eta$ for all $a \in \mathcal{A}_2 \setminus \{a_F\}$,
- $u_1(a_1, a_v) = \gamma + \eta$ and $u_1(a_2, a_v) = \gamma$ for all $v \in V$.
- $u_1(a_1, \bar{a}_v) = \gamma$ and $u_1(a_2, \bar{a}_v) = \gamma + \eta$ for all $v \in V$.
- $u_1(a_v, a_v) = u_1(a_v, \bar{a}_v) = \gamma$ for all $v \in V$
- $u_1(a_v, a_{v'}) = \gamma$ and $u_1(a_v, \bar{a}_{v'}) = \gamma + \frac{\ell - \ell^{1-\delta}}{\ell - \ell^{1-\delta} - 1}\eta$ for all $v' \neq v$.
- $u_1(a_F, a) = 0$ for each $a \in \mathcal{A}_2$.
- $u_1(a, a_F) = 0$ for each $a \in \mathcal{A}_1$.

The utility of the second agent is $u^2(a_0, a) = 1$ for each $a \in \mathcal{A}_2 \setminus \{a_F\}$ and 0 otherwise.

There is a cost function c_v for each $v \in V$, which is common to both the agents. For each $v \in V$, the cost function c_v is such that

- $c_v(a_v, a_{v'}) = -1$ for each $v' \neq v$, $(v, v') \in E$,
- $c_v(a_v, a_{v'}) = 0$ for each $v' \neq v, (v, v') \notin E$,
- $c_v(a_v, a_v) = 1$ for each $v \in V$.
- $c_v(a_F, a) = -\frac{1}{4\ell^2}$ for each $a \in \mathcal{A}_2$.
- $c_v(a, a_F) = -\frac{1}{4\ell^2}$ for each $a \in \mathcal{A}_1$.
- For every other action profile the cost is 0.

We dropped the player index from the cost functions c as they are equal to both players.

Moreover, we set of deviations $\Phi_i = \Phi_{i,\text{ALL}}$ for both players $i \in \{1, 2\}$.

Notice that the instance satisfies Assumption 2. Indeed, the deviation ϕ_i such that $\phi_i[a, a_F] = 1$ for all $a \in \mathcal{A}_i$ for $i \in \{1, 2\}$, that deviates deterministically to a_F is always strictly feasible for both player 1 and player 2. Moreover, its cost is polynomial in the instance size.

[10] This action is needed only to satisfy the strictly feasibility assumption.

Completeness. We show that if there exists an independent set of size $\ell^{1-\delta}$, then the social welfare of an optimal Constrained Phi-equilibria is at least 1. Let V^* be an independent set of size $\ell^{1-\delta}$. We build a Constrained Phi-equilibria $\boldsymbol{z}$ with social welfare at least 1. Consider the correlated strategy such that $\boldsymbol{z}[a_0, a_v] = \frac{1}{2\ell^{1-\delta}}$ for all $v \in V^*$, while $\boldsymbol{z}[a_0, \bar{a}_v] = \frac{1}{2(\ell-\ell^{1-\delta})}$ for all $v \notin V^*$. All the other action profiles have probability 0.

It is easy to see that the correlated strategy has social welfare at least 1 since player 1 always plays action a_0 and $u^2(a_0, a) = 1$ for all $a \in \mathcal{A}_2$. Moreover, it is easy to verify that it is safe since $c_v(a_0, a) \le 0$ for each $a \in \mathcal{A}_2$. Hence, to show that z is an Constrained Phi-equilibria we only need to prove that it satisfies the incentive constraints. The incentive constraints of the second player are satisfied since they obtain the maximum possible utility, *i.e.*, 1.

Consider now a possible deviation of the first player $\phi_1 \in \Phi_1$. As a first step, we compute the expected utility of a strategy ϕ_1. Let us define the following quantities:

- $T^1 = \sum_{v \in V^*} \phi_1[a_0, a_v] \left[\left(\boldsymbol{z}[a_0, a_v] + \boldsymbol{z}[a_0, \bar{a}_v] + \sum_{v' \neq v} \boldsymbol{z}[a_0, a_{v'}] \right) \gamma + \left(\gamma + \frac{\ell-\ell^{1-\delta}}{\ell-\ell^{1-\delta}-1} \eta \right) \sum_{v' \neq v} \boldsymbol{z}[a_0, \bar{a}_v] \right]$
- $T^2 = \sum_{v \notin V^*} \phi_1[a_0, a_v] \left[\left(\boldsymbol{z}[a_0, a_v] + \boldsymbol{z}[a_0, \bar{a}_v] + \sum_{v' \neq v} \boldsymbol{z}[a_0, a_{v'}] \right) \gamma + \left(\gamma + \frac{\ell-\ell^{1-\delta}}{\ell-\ell^{1-\delta}-1} \eta \right) \sum_{v' \neq v} \boldsymbol{z}[a_0, \bar{a}_v] \right]$
- $T^3 = \left(\gamma + \frac{\eta}{2} \right) \phi_1[a_0, a_0] + \frac{\gamma+\eta}{2} (\phi_1[a_0, a_1] + \phi_1[a_0, a_2]) + \frac{\gamma}{2} (\phi_1[a_0, a_1] + \phi_1[a_0, a_2])$

We bound each component individually.

$$
\begin{aligned}
T^1 &= \sum_{v \in V^*} \phi_1[a_0, a_v] \left[\left(\boldsymbol{z}[a_0, a_v] + \boldsymbol{z}[a_0, \bar{a}_v] + \sum_{v' \neq v} \boldsymbol{z}[a_0, a_{v'}] \right) \gamma + \left(\gamma + \eta \frac{\ell - \ell^{1-\delta}}{\ell - \ell^{1-\delta} - 1} \right) \sum_{v' \neq v} \boldsymbol{z}[a_0, \bar{a}_v] \right] \\
&= \sum_{v \in V^*} \phi_1[a_0, a_v] \left[\frac{1}{2} \gamma + \frac{1}{2} \left(\gamma + \eta \frac{\ell - \ell^{1-\delta}}{\ell - \ell^{1-\delta} - 1} \right) \right] \\
&= \sum_{v \in V^*} \phi_1[a_0, a_v] \left(\gamma + \frac{\eta}{2} \frac{\ell - \ell^{1-\delta}}{\ell - \ell^{1-\delta} - 1} \right) \\
&\le \sum_{v \in V^*} \phi_1[a_0, a_v] (\gamma + \eta),
\end{aligned}
$$

where in the last inequality we use $\frac{\ell-\ell^{1-\delta}}{\ell-\ell^{1-\delta}-1} \le 2$ for ℓ large enough. while

$$
\begin{aligned}
T^2 &= \sum_{v \notin V^*} \phi_1[a_0, a_v] \left[\left(\boldsymbol{z}[a_0, a_v] + \boldsymbol{z}[a_0, \bar{a}_v] + \sum_{v' \neq v} \boldsymbol{z}[a_0, a_{v'}] \right) \gamma + \left(\gamma + \eta \frac{\ell - \ell^{1-\delta}}{\ell - \ell^{1-\delta} - 1} \right) \sum_{v' \neq v} \boldsymbol{z}[a_0, \bar{a}_v] \right] \\
&= \sum_{v \notin V^*} \phi_1[a_0, a_v] \left[\left(\frac{1}{2} + \frac{1}{2(\ell - \ell^{1-\delta})} \right) \gamma + \left(\frac{1}{2} - \frac{1}{2(\ell - \ell^{1-\delta})} \right) \left(\gamma + \eta \frac{\ell - \ell^{1-\delta}}{\ell - \ell^{1-\delta} - 1} \right) \right] \\
&= \sum_{v \notin V^*} \phi_1[a_0, a_v] \left[\gamma + \frac{\eta}{2} \left(\frac{\ell - \ell^{1-\delta}}{\ell - \ell^{1-\delta} - 1} - \frac{1}{\ell - \ell^{1-\delta} - 1} \right) \right] \\
&= \sum_{v \notin V^*} \phi_1[a_0, a_v] \left(\gamma + \frac{\eta}{2} \right).
\end{aligned}
$$

Finally,

$$
\begin{aligned}
T^3 &= [a_0, a_0] \left(\gamma + \frac{\eta}{2} \right) + \frac{\gamma + \eta}{2} ([a_0, a_1] + [a_0, a_2]) + \frac{\gamma}{2} ([a_0, a_1] + [a_0, a_2]) \\
&= \left(\gamma + \frac{\eta}{2} \right) ([a_0, a_0] + \phi_1[a_0, a_1] + \phi_1[a_0, a_2])
\end{aligned}
$$

Finally, the utility of a deviation ϕ_1 is

$$
\sum_{a^1 \in \mathcal{A}_1, a^2 \in \mathcal{A}_2} \sum_{a \in A_1} \phi_1[a^1, a] \boldsymbol{z}[a^1, a^2] u_1(a, a^2)
$$

$$
\begin{aligned}
&= \sum_{a\in\mathcal{A}_1, a^2\in\mathcal{A}_2} \phi_1[a_0,a]\boldsymbol{z}[a_0,a^2]u_1(a,a^2) \\
&= T^1+T^2+T^3 \\
&\le (\gamma+\eta)\sum_{v\in V^*}\phi_1[a_0,a_v] + \left(\gamma+\frac{\eta}{2}\right)\sum_{v\notin V^*}\phi_1[a_0,a_v] + \left(\gamma+\frac{\eta}{2}\right)(\phi[a_0,a_0]+\phi_1[a_0,a_1]+\phi_1[a_0,a_2]) \\
&= \frac{\eta}{2}\sum_{v\in V^*}\phi_1[a_0,a_v] + \left(\gamma+\frac{\eta}{2}\right)(1-\phi_1[a_0,a_F])
\end{aligned}
$$

Now, we show that no deviation $\phi_1 \in \Phi_1$ is both safe and increases player 1 utility. In particular, we show that if a strategy ϕ_1 increases the utility than it is not safe. Indeed, if ϕ_1 increases the utility, then

$$
\sum_{\substack{a^1\in\mathcal{A}_1,\\ a^2\in\mathcal{A}_2}}\sum_{a\in A_1}\phi_1[a^1,a]\boldsymbol{z}[a^1,a^2]u_1(a,a^2) > \gamma+\frac{\eta}{2}
$$

This implies that

$$
\frac{\eta}{2}\sum_{v\in V^*}\phi_1[a_0,a_v] + \left(\gamma+\frac{\eta}{2}\right)(1-\phi_1[a_0,a_F]) > \gamma+\frac{\eta}{2}
$$

and

$$
\sum_{v\in V^*}\phi_1[a_0,a_v] > \frac{1}{2}\phi_1[a_0,a_F] \tag{6}
$$

Next, we show that any ϕ_1 that increases the utility (and hence that satisfies Eq (6)) is not a feasible deviation. First, notice that equation (6) implies that there is a $\bar{v}\in V^*$ such that

$$
\phi_1[a_0,a_{\bar{v}}] > \frac{1}{2\ell}\phi_1[a_0,a_F]. \tag{7}
$$

Then, we show that the deviation ϕ_1 violates the constraint $c_{\bar{v}}$. In particular,

$$
\begin{aligned}
\sum_{a^1\in\mathcal{A}_1,a^2\in\mathcal{A}_2}\sum_{a\in\mathcal{A}_1}\phi_1[a^1,a]\boldsymbol{z}[a^1,a^2]c_v(a,a^2) &= \phi_1[a_0,a_{\bar{v}}]\boldsymbol{z}[a_0,a_{\bar{v}}]1 - \frac{1}{4\ell^2}\phi_1[a_0,a_F] - \sum_{v\in V^*:(v,\bar{v})\in E}\phi_1[a_0,a_{\bar{v}}]\boldsymbol{z}[a_0,a_v]1 \\
&= \phi_1[a_0,a_{\bar{v}}]\boldsymbol{z}[a_0,a_{\bar{v}}] - \frac{1}{4\ell^2}\phi_1[a_0,a_F] \\
&= \frac{1}{2\ell^{1-\frac{1}{\ell}}}\phi_1[a_0,a_{\bar{v}}] - \frac{1}{4\ell^2}\phi_1[a_0,a_F] \\
&> \phi_1[a_0,a_{\bar{v}}]\left(\frac{1}{2\ell^{1-\frac{1}{\ell}}} - \frac{1}{2\ell}\right) \ge 0,
\end{aligned}
$$

where the second inequality holds since V^* is an independent set, and the second-to-last inequality by Equation (7). Hence, there is no deviation ϕ_1 that increases players 1 utility and that is safe. This concludes the first part of the proof.

Soundness. We show that if there exists a Constrained w-Phi-equilibria with social welfare $\alpha/2$, then there exists an independent set of size strictly larger than ℓ^δ, reaching a contradiction. Suppose by contradiction that there exists a Constrained ϵ-Phi-equilibrium $\boldsymbol{z}$ with social welfare strictly greater than $\alpha/2$. Thus,

$$
\sum_{a'\in\mathcal{A}_2\setminus\{a_F\}}\boldsymbol{z}[a_0,a']\cdot 1 + \sum_{a\in\mathcal{A}_1,a'\in\mathcal{A}_2}(\gamma+\eta) \ge \sum_{a\in\mathcal{A}_1,a'\in\mathcal{A}_2}\boldsymbol{z}[a,a'](u_1(a,a')+u^2(a,a')) \ge \alpha/2,
$$

where the first inequality comes from $u_2(a_0,a')=1$ for each $a'\in\mathcal{A}_2\setminus\{a_F\}$ and 0 otherwise, and $u_1(a,a')\le\gamma+\eta$ for each $a\in\mathcal{A}_1$ and $a'\in\mathcal{A}_2$. This implies

$$
\sum_{a'\in\mathcal{A}_2}\boldsymbol{z}[a_0,a'] \ge \alpha/4. \tag{8}
$$

Then, we show that $\boldsymbol{z}$ assigns similar probabilities on the set of action profiles $\{a_0, a_v\}_{v\in V}$ and $\{a_0, \bar{a}_v\}_{v\in V}$ Given an $a \in \mathcal{A}_1$, let $\phi_a \in \Phi_1$ be a deviation of the first player such that $\phi_a[a_0, a] = 1$ and $\phi_a[a', a'] = 1$ for each $a' \neq a_0$. Since $\boldsymbol{z}$ is an Constrained ϵ-Phi-equilibrium there is no feasible deviation ϕ_a that increases the utility of player 1 by more than ϵ. This implies that

$$\left|\sum_{v\in V} \boldsymbol{z}[a_0, a_v] - \sum_{v\in V} \boldsymbol{z}[a_0, \bar{a}_v]\right| \leq \frac{2\epsilon}{\eta}. \tag{9}$$

Indeed, if

$$\sum_{v\in V} \boldsymbol{z}[a_0, a_v] > \sum_{v\in V} \boldsymbol{z}[a_0, \bar{a}_v] + \frac{2\epsilon}{\eta}, \tag{10}$$

then the deviation ϕ_{a_1} has utility at least

$$\begin{aligned}
&\sum_{v\in V} \boldsymbol{z}[a_0, a_v]\phi_{a_1}[a_0, a_1](\gamma + \eta) + \boldsymbol{z}[a_0, \bar{a}_v]\phi_{a_1}[a_0, a_1]\gamma + \sum_{a\in\mathcal{A}_1\setminus\{a_0\}} \sum_{a'\in\mathcal{A}_2} \boldsymbol{z}[a, a']\phi_{a_1}[a, a]u^i(a, a') \\
&\quad = \eta \sum_{v\in V} \boldsymbol{z}[a_0, a_v] + \gamma \sum_{v\in V} (\boldsymbol{z}[a_0, a_v] + \boldsymbol{z}[a_0, \bar{a}_v]) + \sum_{a\in\mathcal{A}_1\setminus\{a_0\}} \sum_{a'\in\mathcal{A}_2} \boldsymbol{z}[a, a']\phi_{a_1}[a, a]u^i(a, a') \\
&\quad > \frac{\eta}{2}\left(\frac{2\epsilon}{\eta} + \sum_{v\in V}(\boldsymbol{z}[a_0, a_v] + \boldsymbol{z}[a_0, \bar{a}_v])\right) + \gamma \sum_{v\in V} (\boldsymbol{z}[a_0, a_v] + \boldsymbol{z}[a_0, \bar{a}_v]) \\
&\qquad\quad + \sum_{a\in\mathcal{A}_1\setminus\{a_0\}} \sum_{a'\in\mathcal{A}_2} \boldsymbol{z}[a, a']\phi_{a_1}[a, a]u^i(a, a') \\
&\quad \geq \epsilon + \left(\frac{\eta}{2} + \gamma\right) \sum_{v\in V}(\boldsymbol{z}[a_0, a_v] + \boldsymbol{z}[a_0, \bar{a}_v]) + \sum_{a\in\mathcal{A}_1\setminus\{a_0\}} \sum_{a'\in\mathcal{A}_2} \boldsymbol{z}[a, a']\phi_{a_1}[a, a]u^i(a, a') \\
&\quad \geq u_1(z) + \epsilon,
\end{aligned}$$

where the first inequality comes from adding $\sum_{v\in V} \boldsymbol{z}[a_0, a_v]$ to both sides of Equation (10). Moreover, ϕ_{a_1} is feasible since for each constraint $c_{\bar{v}}$, $\bar{v} \in V$, it has cost

$$\begin{aligned}
&\sum_{v\in V} (\boldsymbol{z}[a_0, a_v]\phi_{a_1}[a_0, a_1]c_{\bar{v}}(a_1, a_v) + \boldsymbol{z}[a_0, \bar{a}_v]\phi_{a_1}[a_0, a_1]c_{\bar{v}}(a_1, a_v)) \\
&\qquad + \sum_{a\in\mathcal{A}_1\setminus\{a_0\}} \sum_{a'\in\mathcal{A}_2} \boldsymbol{z}[a, a']\phi_{a_1}[a, a]c_{\bar{v}}(a, a') \\
&= \sum_{v\in V} (\boldsymbol{z}[a_0, a_v]\phi_{a_1}[a_0, a_1]c_{\bar{v}}(a_0, a_v) + \boldsymbol{z}[a_0, \bar{a}_v]\phi_{a_1}[a_0, a_1]c_{\bar{v}}(a_0, \bar{a}_v)) \\
&\qquad + \sum_{a\in\mathcal{A}_1\setminus\{a_0\}} \sum_{a'\in\mathcal{A}_2} \boldsymbol{z}[a, a']\phi_{a_1}[a, a]c_{\bar{v}}(a, a') \\
&= c_{\bar{v}}(z) \leq 0.
\end{aligned}$$

A similar argument shows that if $\sum_{v\in V} \boldsymbol{z}[a_0, a_v] < \sum_{v\in V} \boldsymbol{z}[a_0, \bar{a}_v] - \frac{2\epsilon}{\eta}$ then the deviation ϕ_{a_2} is safe and increases the utility. As a consequence of Equation (9), it holds

$$2\sum_{v\in V} \boldsymbol{z}[a_0, \bar{a}_v] \geq \sum_{v\in V}(\boldsymbol{z}[a_0, a_v] + \boldsymbol{z}[a_0, \bar{a}_v]) - \frac{\epsilon}{\eta} = \sum_{a\in\mathcal{A}_2\setminus\{a_F\}} \boldsymbol{z}[a_0, a] - \frac{2\epsilon}{\eta}, \tag{11}$$

where the first inequality comes from adding $\sum_{v\in V} \boldsymbol{z}[a_0, \bar{a}_v]$ to both sides of $\sum_{v\in V} \boldsymbol{z}[a_0, \bar{a}_v]| \geq \sum_{v\in V} \boldsymbol{z}[a_0, a_v] - \frac{2\epsilon}{\eta}$

The next step is to show that it is if there is no safe deviation ϕ_{a_v}, $v \in V$, that increases the utility, then there exists an independent set of size larger than ℓ^δ. Since z is an Constrained ϵ-Phi-equilibrium, for each a_v, $v \in V$ one of the following two conditions holds: i) $\phi_{a_v} \notin \Phi_1^{\mathcal{S}}(z)$ or ii) $u_1(\phi_{a_v} \diamond z) \leq u_1(z) + \epsilon$. Let $V^1 \subseteq V$ be the set of vertexes v such that ϕ_{a_v}

is not safe, *i.e.*, $\phi_{a_v} \notin \Phi_1^{\mathcal{S}}(z)$, and $V^2 = V \setminus V^1$ be the set of v such that ϕ_{a_v} does not increase the utility by more than ϵ and are not in V^1, *i.e.*, $u_1(\phi_{a_v} \diamond z) \leq u_1(z)$ and $\phi_{a_v} \in \Phi_1^{\mathcal{S}}(z)$. We show that $|V^2| \leq \ell - \ell^{1-\delta}$. Indeed, for each $v \in V^2$, deviation ϕ_{a_v} does not increase the utility and hence it holds:

$$\begin{aligned}
\gamma \sum_{a \in \mathcal{A}_2 \setminus \{a_F\}} \boldsymbol{z}[a_0, a] + \eta \frac{\ell - \ell^{1-\delta}}{\ell - \ell^{1-\delta} - 1} \sum_{v' \neq v} \boldsymbol{z}[a_0, \bar{a}_{v'}] + \sum_{a \in A_1 \setminus \{a_0\}} \sum_{a' \in \mathcal{A}_2} \boldsymbol{z}[a, a'] \phi_{a_1}[a, a] u^i(a, a') \\
= \left(\sum_{v' \in V} \boldsymbol{z}[a_0, a] + \boldsymbol{z}[a_0, \bar{a}_v] \right) \phi[a_0, a_v] \gamma + \sum_{v' \neq v} \phi[a_0, a_v] \boldsymbol{z}[a_0, \bar{a}_{v'}] \left(\gamma + \eta \frac{\ell - \ell^{1-\delta}}{\ell - \ell^{1-\delta} - 1} \right) \\
+ \sum_{a \in \mathcal{A}_1 \setminus \{a_0\}} \sum_{a' \in \mathcal{A}_2} \boldsymbol{z}[a, a'] \phi_{a_1}[a, a] u^i(a, a') \\
\leq u_1(z) + \epsilon \\
= \left(\gamma + \frac{\eta}{2} \right) \sum_{a \in \mathcal{A}_2 \setminus \{a_F\}} \boldsymbol{z}[a_0, a] + \sum_{a \in \mathcal{A}_1 \setminus \{a_0\}} \sum_{a' \in \mathcal{A}_2} \boldsymbol{z}[a, a'] u^i(a, a') + \epsilon,
\end{aligned}$$

where the inequality holds since the lhs is the utility of the deviation ϕ_{a_v}.

This implies

$$\left(\sum_{v'} \boldsymbol{z}[a_0, \bar{a}_{v'}] - \boldsymbol{z}[a_0, \bar{a}_v] \right) \eta \frac{\ell - \ell^{1-\delta}}{\ell - \ell^{1-\delta} - 1} \leq \frac{\eta}{2} \sum_{a \in \mathcal{A}_2 \setminus \{a_F\}} \boldsymbol{z}[a_0, a] + \epsilon \leq \eta \sum_{v \in V} \boldsymbol{z}[a_0, \bar{a}_v] + 2\epsilon,$$

where the last inequality holds by Equation (11). Hence,

$$\boldsymbol{z}[a_0, \bar{a}_v] \frac{\ell - \ell^{1-\delta}}{\ell - \ell^{1-\delta} - 1} \geq \left(\frac{\ell - \ell^{1-\delta}}{\ell - \ell^{1-\delta} - 1} - 1 \right) \sum_{v'} \boldsymbol{z}[a_0, \bar{a}_{v'}] - 2\epsilon/\eta,$$

and

$$\bar{\boldsymbol{z}}[a_0, a_v] \geq \frac{1}{\ell - \ell^{1-\delta}} \sum_{v'} \boldsymbol{z}[a_0, \bar{a}_{v'}] - 2\epsilon/\eta. \tag{12}$$

Suppose that $|V^2| > \ell - \ell^{1-\delta}$, and hence Equation (12) is satisfied by at least $|V^2| \geq \ell - \ell^{1-\delta} + 1$ vertexes. We need the following inequality.

$$\frac{1}{\ell} \sum_{v'} \boldsymbol{z}[a_0, \bar{a}_{v'}] \geq \frac{1}{\ell} \sum_{a \in \mathcal{A}_2 \setminus \{a_F\}} \boldsymbol{z}[a_0, a] - \frac{2\epsilon}{\ell\eta} \geq \frac{\alpha}{4\ell} - \frac{2\epsilon}{\ell\eta} = \frac{\alpha}{4\ell} - \frac{\alpha}{8\ell^3} \geq \frac{\alpha}{8\ell} = \frac{2\ell}{\eta} \left(\frac{\alpha^2}{16\ell^2} \right) = \frac{2\ell}{\eta} \epsilon \tag{13}$$

where the first inequality comes from Equation (11), and the second one by Equation (8). Then, summing over the $|V^2|$ inequalities we get

$$\begin{aligned}
\sum_{v \in V^2} \bar{\boldsymbol{z}}[a_0, a_v] \geq (\ell - \ell^{1-\delta} + 1) \left(\frac{1}{\ell - \ell^{1-\delta}} \sum_{v'} \boldsymbol{z}[a_0, \bar{a}_{v'}] - 2\epsilon/\eta \right) \\
\geq \sum_{v'} \boldsymbol{z}[a_0, \bar{a}_{v'}] + \frac{1}{\ell} \sum_{v'} \boldsymbol{z}[a_0, \bar{a}_{v'}] - 2\ell \frac{\epsilon}{\eta} \\
> \sum_{v'} \boldsymbol{z}[a_0, \bar{a}_{v'}],
\end{aligned}$$

where the last inequality follows from equation (13). Hence, we reach a contradiction and $|V^2| \leq \ell - \ell^{1-\delta}$.

To conclude the proof, we show that V^1 is an independent set. Since $|V^1| \geq |V| - |V^2| = \ell^{1-\delta}$ we reach a contradiction. Let v and v' be two nodes in V^1 and w.l.o.g. let $\boldsymbol{z}[a_0, a_v] \geq \boldsymbol{z}[a_0, a_{v'}]$. We show that $(v, v') \notin E$. Since $v' \in V^1$, ϕ_{a_v} is not a safe deviation for player 1 with respect to constraint $c^{v'}$. if $(v, v') \in E$, then

$$
\begin{aligned}
&\sum_{a^1\in\mathcal{A}_1, a^2\in\mathcal{A}_2} \sum_{a\in\ \mathcal{A}_1} \phi[a^1, a]\boldsymbol{z}[a^1, a^2]c_v(a, a^2) \\
&\qquad = \boldsymbol{z}[a_0, a'_v] - \sum_{v'':(v'',v')\in E} \boldsymbol{z}[a_0, a_{v''}] - \frac{1}{4\ell}\boldsymbol{z}[a_0, a_F]c_v(a, a^2) \\
&\qquad\qquad + \sum_{a^1\in\mathcal{A}_1\setminus\{a_0\}, a^2\in\mathcal{A}_2} \sum_{a\in\mathcal{A}_1} \phi[a^1, a]\boldsymbol{z}[a^1, a^2]c_v(a, a^2) \\
&\qquad \leq \boldsymbol{z}[a_0, a'_v] - \boldsymbol{z}[a_0, a_v] - \frac{1}{4\ell}\boldsymbol{z}[a_0, a_F]c_v(a, a^2) + \\
&\qquad\qquad + \sum_{a^1\in\mathcal{A}_1\setminus\{a_0\}, a^2\in\mathcal{A}_2} \sum_{a\in\mathcal{A}_1} \phi[a^1, a]\boldsymbol{z}[a^1, a^2]c_v(a, a^2) \\
&\qquad \leq -\frac{1}{4\ell}\boldsymbol{z}[a_0, a_F]c_v(a, a^2) + \sum_{a^1\in\mathcal{A}_1\setminus\{a_0\}, a^2\in\mathcal{A}_2} \sum_{a\in\mathcal{A}_1} \phi[a^1, a]\boldsymbol{z}[a^1, a^2]c_v(a, a^2) \\
&\qquad = c_v(z) \leq 0.
\end{aligned}
$$

Hence, $(v, v') \notin E$. Since V^1 is an independent set of size at least $\ell^{1-\delta}$ we reach a contradiction. This concludes the proof. □

D. Proofs Omitted from Section 4

Lemma 4.1. *For every $\boldsymbol{z} \in \Delta_{\mathcal{A}}$ and $i \in \mathcal{N}$, it holds*

$$
\begin{aligned}
&\max_{\phi_i\in\Phi_i^{\mathcal{S}}(\boldsymbol{z})} u_i(\phi_i \diamond \boldsymbol{z}) \\
&\qquad = \sup_{\phi_i\in\Phi_i} \inf_{\boldsymbol{\eta}_i\in\mathbb{R}_+^m} \left(u_i(\phi_i \diamond \boldsymbol{z}) - \boldsymbol{\eta}_i^\top \boldsymbol{c}_i(\phi_i \diamond \boldsymbol{z})\right) \\
&\qquad = \inf_{\boldsymbol{\eta}_i\in\mathbb{R}_+^m} \sup_{\phi_i\in\Phi_i} \left(u_i(\phi_i \diamond \boldsymbol{z}) - \boldsymbol{\eta}_i^\top \boldsymbol{c}_i(\phi_i \diamond \boldsymbol{z})\right).
\end{aligned}
$$

Proof. First, it is easy to see that

$$
\sup_{\phi_i\in\Phi_i^{\mathcal{S}}(\boldsymbol{z})} u_i(\phi_i \diamond \boldsymbol{z}) = \sup_{\phi_i\in\Phi_i} \inf_{\boldsymbol{\eta}_i\in\mathbb{R}_+^m} \left(u_i(\phi_i \diamond \boldsymbol{z}) - \boldsymbol{\eta}_i^\top \boldsymbol{c}_i(\phi_i \diamond \boldsymbol{z})\right).
$$

Indeed, for every $\phi_i \notin \Phi_i^{\mathcal{S}}(\boldsymbol{z})$, it holds that the vector $\boldsymbol{c}_i(\phi_i \diamond \boldsymbol{z})$ has at least one positive component, and, thus, the vector of Lagrange multipliers $\boldsymbol{\eta}_i$ can be selected so that $u_i(\phi_i \diamond \boldsymbol{z}) - \boldsymbol{\eta}_i^\top \boldsymbol{c}_i(\phi_i \diamond \boldsymbol{z})$ goes to $-\infty$. This implies that the supremum over Φ_i cannot be attained in $\Phi_i^{\mathcal{S}}(\boldsymbol{z})$. On the other hand, for every $\phi_i \in \Phi_i^{\mathcal{S}}(\boldsymbol{z})$, all the components of $\boldsymbol{c}_i(\phi_i \diamond \boldsymbol{z})$ are negative, and, thus, the $\inf$ is achieved by $\boldsymbol{\eta}_i = \boldsymbol{0}$. This proves the first equality.

Then, the second equality directly follows from the generalization of the max-min theorem for unbounded domains (see (Ekeland & Temam, 1999, Proposition 2.3)), which allows us to swap the $\sup$ and the $\inf$. □

Lemma D.1. *For any two real-valued functions $f(x)$ and $g(x)$ with $g(x) \leq c$ then $\min(f(x), g(x)) \leq \min(f(x), c)$.*

Proof. We can identify three sets I_1, I_2 and I_3 defined as follows:

$$I_1 \coloneqq \{x \, s.t. \, f(x) \geq c\} \tag{14}$$

$$I_2 \coloneqq \{x \, s.t. \, g(x) \leq f(x) \leq c\} \tag{15}$$

$$I_3 \coloneqq \{x \, s.t. \, f(x) \leq g(x) \leq c\}. \tag{16}$$

Then for all $x \in I_1$ we have that $\min(f(x), c) = c \geq \min(f(x), g(x)) = g(x)$, while for all $x \in I_2$ we have that $\min(f(x), c) = f(x) \geq \min(f(x), g(x)) = g(x)$. Finally for all $x \in I_3$ we have $\min(f(x), c) = f(x) = \min(f(x), g(x)) = f(x)$. In all three sets we have that $\min(f(x), c) \geq \min(f(x), g(x))$. □

Lemma D.2. *For all $\boldsymbol{\eta}_i \in \mathcal{D}^c$ it holds that*

$$\sup_{\phi_i \in \Phi_i} \left(u_i(\phi_i \diamond \boldsymbol{z}) - \boldsymbol{\eta}_i^\top \boldsymbol{c}_i(\phi_i \diamond \boldsymbol{z}) \right) \geq 1$$

Proof. Thanks to Assumption 2 we have that for all $\boldsymbol{z} \in \Delta(\mathcal{A})$ we have that there exists $\tilde{\phi}_i \in \Phi_i^{\mathcal{S}}(\boldsymbol{z})$ such that $\boldsymbol{c}_i(\tilde{\phi}_i \diamond \boldsymbol{z}) \preceq -\rho\mathbf{1}$. Then, for all $\boldsymbol{\eta}_i \in \mathcal{D}^c$ we have:

$$\boldsymbol{\eta}_i^\top \boldsymbol{c}_i(\tilde{\phi}_i \diamond \boldsymbol{z}) \leq -\rho\|\boldsymbol{\eta}_i\|_1 \leq -1.$$

This easily concludes the proof of the statement

$$\sup_{\phi_i \in \Phi_i} \left(u_i(\phi_i \diamond \boldsymbol{z}) - \boldsymbol{\eta}_i^\top \boldsymbol{c}_i(\phi_i \diamond \boldsymbol{z}) \right) \geq u_i(\tilde{\phi}_i \diamond \boldsymbol{z}) - \boldsymbol{\eta}_i^\top \boldsymbol{c}_i(\tilde{\phi}_i \diamond \boldsymbol{z}) \geq 1,$$

as u_i is positive. □

Lemma D.3. *For all $\boldsymbol{\eta}_i \in \mathcal{D}$ we have*

$$\inf_{\boldsymbol{\eta}_i \in \mathcal{D}} \sup_{\phi_i \in \Phi_i} \left(u_i(\phi_i \diamond \boldsymbol{z}) - \boldsymbol{\eta}_i^\top \boldsymbol{c}_i(\phi_i \diamond \boldsymbol{z}) \right) \leq 1$$

Proof. Since $u_i \leq 1$ we have that

$$\inf_{\boldsymbol{\eta}_i \in \mathcal{D}} \sup_{\phi_i \in \Phi_i} \left(u_i(\phi_i \diamond \boldsymbol{z}) - \boldsymbol{\eta}_i^\top \boldsymbol{c}_i(\phi_i \diamond \boldsymbol{z}) \right) \leq 1 - \sup_{\boldsymbol{\eta}_i \in \mathcal{D}} \inf_{\phi_i \in \Phi_i} \boldsymbol{\eta}_i^\top \boldsymbol{c}_i(\phi_i \diamond \boldsymbol{z}).$$

Next we claim that $\sup_{\boldsymbol{\eta}_i \in \mathcal{D}} \inf_{\phi_i \in \Phi_i} \boldsymbol{\eta}_i^\top \boldsymbol{c}_i(\phi_i \diamond \boldsymbol{z}) \geq 0$. This follows from the fact that for all negative components of $\boldsymbol{c}_i(\phi_i \diamond \boldsymbol{z})$ then the corresponding components of $\boldsymbol{\eta}_i$ will be 0. This concludes the statement. □

Lemma 4.2. *Let $\mathcal{D} := \left\{\boldsymbol{\eta} \in \mathbb{R}_+^m \mid ||\boldsymbol{\eta}||_1 \leq 1/\rho\right\}$. Then, for every $\boldsymbol{z} \in \Delta_{\mathcal{A}}$ and $i \in \mathcal{N}$, it holds:*

$$\begin{aligned} &\max_{\phi_i \in \Phi_i^{\mathcal{S}}(\boldsymbol{z})} u_i(\phi_i \diamond \boldsymbol{z}) \\ &\qquad = \max_{\phi_i \in \Phi_i} \min_{\boldsymbol{\eta}_i \in \mathcal{D}} \left(u_i(\phi_i \diamond \boldsymbol{z}) - \boldsymbol{\eta}_i^\top \boldsymbol{c}_i(\phi_i \diamond \boldsymbol{z}) \right) \\ &\qquad = \min_{\boldsymbol{\eta}_i \in \mathcal{D}} \max_{\phi_i \in \Phi_i} \left(u_i(\phi_i \diamond \boldsymbol{z}) - \boldsymbol{\eta}_i^\top \boldsymbol{c}_i(\phi_i \diamond \boldsymbol{z}) \right). \end{aligned}$$

Proof. In Lemma 4.1 we already showed that:

$$\begin{aligned} \sup_{\phi_i \in \Phi_i^{\mathcal{S}}(\boldsymbol{z})} u_i(\phi_i \diamond \boldsymbol{z}) &= \sup_{\phi_i \in \Phi_i} \inf_{\boldsymbol{\eta}_i \in \mathbb{R}_+^m} \left(u_i(\phi_i \diamond \boldsymbol{z}) - \boldsymbol{\eta}_i^\top \boldsymbol{c}_i(\phi_i \diamond z) \right) = \\ &\inf_{\boldsymbol{\eta}_i \in \mathbb{R}_+^m} \sup_{\phi_i \in \Phi_i} \left(u_i(\phi_i \diamond \boldsymbol{z}) - \boldsymbol{\eta}_i^\top \boldsymbol{c}_i(\phi_i \diamond \boldsymbol{z}) \right). \end{aligned}$$

Note that to prove the statement it is enough to prove that:

$$\inf_{\boldsymbol{\eta}_i \in \mathbb{R}_+^m} \sup_{\phi_i \in \Phi_i} \left(u_i(\phi_i \diamond \boldsymbol{z}) - \boldsymbol{\eta}_i^\top \boldsymbol{c}_i(\phi_i \diamond \boldsymbol{z}) \right) = \inf_{\boldsymbol{\eta}_i \in \mathcal{D}} \sup_{\phi_i \in \Phi_i} \left(u_i(\phi_i \diamond \boldsymbol{z}) - \boldsymbol{\eta}_i^\top \boldsymbol{c}_i(\phi_i \diamond \boldsymbol{z}) \right)$$

and more specifically that:

$$\inf_{\boldsymbol{\eta}_i \in \mathbb{R}_+^m} \sup_{\phi_i \in \Phi_i} \left(u_i(\phi_i \diamond \boldsymbol{z}) - \boldsymbol{\eta}_i^\top \boldsymbol{c}_i(\phi_i \diamond \boldsymbol{z}) \right) \geq \inf_{\boldsymbol{\eta}_i \in \mathcal{D}} \sup_{\phi_i \in \Phi_i} \left(u_i(\phi_i \diamond \boldsymbol{z}) - \boldsymbol{\eta}_i^\top \boldsymbol{c}_i(\phi_i \diamond \boldsymbol{z}) \right)$$

since the reverse inequality holds trivially. We can show this by the following inequalities:

$$\begin{aligned} &\inf_{\boldsymbol{\eta}_i \in \mathbb{R}_+^m} \sup_{\phi_i \in \Phi_i} \left(u_i(\phi_i \diamond \boldsymbol{z}) - \boldsymbol{\eta}_i^\top \boldsymbol{c}_i(\phi_i \diamond \boldsymbol{z}) \right) \\ &\qquad = \min \left(\inf_{\boldsymbol{\eta}_i \in \mathcal{D}} \sup_{\phi_i \in \Phi_i} \left(u_i(\phi_i \diamond \boldsymbol{z}) - \boldsymbol{\eta}_i^\top \boldsymbol{c}_i(\phi_i \diamond \boldsymbol{z}) \right), \inf_{\boldsymbol{\eta}_i \in \mathcal{D}^c} \sup_{\phi_i \in \Phi_i} \left(u_i(\phi_i \diamond \boldsymbol{z}) - \boldsymbol{\eta}_i^\top \boldsymbol{c}_i(\phi_i \diamond \boldsymbol{z}) \right) \right) \end{aligned}$$

$$\geq \min\left(\inf_{\boldsymbol{\eta}_i\in\mathcal{D}}\sup_{\phi_i\in\Phi_i}\left(u_i(\phi_i\diamond \boldsymbol{z})-\boldsymbol{\eta}_i^\top \boldsymbol{c}_i(\phi_i\diamond \boldsymbol{z})\right),1\right)$$
$$=\inf_{\boldsymbol{\eta}_i\in\mathcal{D}}\sup_{\phi_i\in\Phi_i}\left(u_i(\phi_i\diamond \boldsymbol{z})-\boldsymbol{\eta}_i^\top \boldsymbol{c}_i(\phi_i\diamond \boldsymbol{z})\right),$$

where the first inequality hold thanks to Lemma D.1 and Lemma D.2, while that last equation follows from Lemma D.3. □

Lemma 4.4. *Given any $0<\delta\leq\epsilon$ and a δ-optimal set $\tilde{\mathcal{D}}\subseteq\mathcal{D}$, the following holds: $L_{\tilde{\mathcal{D}},\epsilon}\geq L_{\mathcal{D},0}$.*

Proof. By definition we have that: $L_{\tilde{\mathcal{D}},\epsilon}=\ell(\tilde{\boldsymbol{z}}^\star)$, where $\tilde{\boldsymbol{z}}^\star$ is a solution to the problem

$$\text{P1}:=\begin{cases}\tilde{\boldsymbol{z}}^\star\in\arg\max_{\boldsymbol{z}\in\mathcal{S}} & \ell(\boldsymbol{z})\ s.t.\\ & \epsilon+u_i(\tilde{\boldsymbol{z}}^\star)\geq\max_{\phi_i\in\Phi_i}\left(u_i(\phi_i\diamond\tilde{\boldsymbol{z}}^\star)-\tilde{\boldsymbol{\eta}}_i^{\star,\top}\boldsymbol{c}_i(\phi_i\diamond\tilde{\boldsymbol{z}}^\star)\right)\\ & \tilde{\boldsymbol{\eta}}_i^\star\in\arg\inf_{\boldsymbol{\eta}_i\in\tilde{\mathcal{D}}}\sup_{\phi_i\in\Phi_i}\left(u_i(\phi_i\diamond\tilde{\boldsymbol{z}}^\star)-\boldsymbol{\eta}_i^\top\boldsymbol{c}_i(\phi_i\diamond\tilde{\boldsymbol{z}}^\star)\right)\end{cases}$$

On the other hand, call $\boldsymbol{z}^\star$ the optimal Constrained Phi-equilibrium. This is a solution to the problem:

$$\text{P2}:=\begin{cases}\boldsymbol{z}^\star\in\arg\max_{\boldsymbol{z}\in\mathcal{S}} & \ell(\boldsymbol{z})\ s.t.\\ & u_i(\boldsymbol{z}^\star)\geq\max_{\phi_i\in\Phi_i}\left(u_i(\phi_i\diamond\boldsymbol{z}^\star)-\boldsymbol{\eta}_i^{\star,\top}\boldsymbol{c}_i(\phi_i\diamond\boldsymbol{z}^\star)\right)\\ & \boldsymbol{\eta}_i^\star\in\arg\inf_{\boldsymbol{\eta}_i\in\mathcal{D}}\sup_{\phi_i\in\Phi_i}\left(u_i(\phi_i\diamond\boldsymbol{z}^\star)-\boldsymbol{\eta}_i^\top\boldsymbol{c}_i(\phi_i\diamond\boldsymbol{z}^\star)\right)\end{cases}$$

which has value $L_{\mathcal{D},0}=\ell(\boldsymbol{z}^\star)$.

Moreover, thanks to Lemma 4.2 and since $\tilde{\mathcal{D}}$ is δ-optimal we have that:

$$\max_{\phi_i\in\Phi_i}\left(u_i(\phi_i\diamond\tilde{\boldsymbol{z}}^\star)-\tilde{\boldsymbol{\eta}}_i^{\star,\top}\boldsymbol{c}_i(\phi_i\diamond\tilde{\boldsymbol{z}}^\star)\right)\leq\max_{\phi_i\in\Phi_i}\left(u_i(\phi_i\diamond\boldsymbol{z}^\star)-\boldsymbol{\eta}_i^{\star,\top}\boldsymbol{c}_i(\phi_i\diamond\boldsymbol{z}^\star)\right)+\delta$$

which implies that feasible correlated strategies of problem P2 are feasible correlated strategies of problem P1, and thus problem P1 as long as $\delta\geq\epsilon$. Thus problem P1 is the problem of maximizing the same objective function over a larger set then problem P2 and thus $L_{\tilde{\mathcal{D}},\epsilon}\geq L_{\mathcal{D},0}$. □

Lemma 4.5. *For any $\tau>0$, the set $\mathcal{D}_\tau$ is (τm)-optimal.*

Proof. By Lemma 4.2, we know that for each player there exists an $\boldsymbol{\eta}_i^\star\in\mathcal{D}$ such that $\max_{\phi\in\Phi_i^{\mathcal{S}}(\boldsymbol{z})}u_i(\phi_i\diamond\boldsymbol{z})=\max_{\phi_i\in\Phi_i}\left(u_i(\phi_i\diamond\boldsymbol{z})-\boldsymbol{\eta}_i^{\star,\top}\boldsymbol{c}_i(\phi_i\diamond\boldsymbol{z})\right)$. By construction of $\mathcal{D}_\epsilon$ there exists a $\bar{\boldsymbol{\eta}}_i\in\mathcal{D}_\epsilon$ such that $||\bar{\boldsymbol{\eta}}_i-\boldsymbol{\eta}_i^\star||_\infty\leq\epsilon$. Thus

$$\max_{\phi\in\Phi_i^{\mathcal{S}}(\boldsymbol{z})}u_i(\phi_i\diamond\boldsymbol{z})=\max_{\phi_i\in\Phi_i}\left(u_i(\phi_i\diamond\boldsymbol{z})-\boldsymbol{\eta}_i^{\star,\top}\boldsymbol{c}_i(\phi_i\diamond\boldsymbol{z})\right)$$
$$\leq\max_{\phi_i\in\Phi_i}\left(u_i(\phi_i\diamond\boldsymbol{z})-\bar{\boldsymbol{\eta}}_i^\top\boldsymbol{c}_i(\phi_i\diamond\boldsymbol{z})\right)+m\epsilon,$$

where the last inequality comes the fact that:

$$|(\boldsymbol{\eta}_i^\star-\bar{\boldsymbol{\eta}}_i)^\top\boldsymbol{c}_i(\phi_i\diamond\boldsymbol{z})|\leq\|\boldsymbol{c}_i(\phi_i\diamond\boldsymbol{z})\|_1\|\boldsymbol{\eta}_i^\star-\bar{\boldsymbol{\eta}}_i\|_\infty\leq m\epsilon$$

as $\boldsymbol{c}_i\in[-1,1]^m$. □

Lemma 4.6. *For any $\tau>0$, the set $\mathcal{D}_\tau$ is δ-optimal for $\delta=2\sqrt{2\tau\log s/\rho}$, where s is the number of players' actions.*

Proof. The proof exploits a probability interpretation of the Lagrange multipliers. Let $\boldsymbol{\eta}^\star$ be the optimal multipliers, *i.e.,* $\boldsymbol{\eta}^\star \in \operatorname{argmin}_{\boldsymbol{\eta}\in\mathcal{D}} \max_{\phi_i\in\Phi_i} \left(u_i(\phi_i \diamond \boldsymbol{z}) - \boldsymbol{\eta}^\top \boldsymbol{c}_i(\phi_i \diamond \boldsymbol{z})\right)$. Now consider a basis $\Gamma = \{\frac{1}{\rho}\boldsymbol{e}_j\}_{j\in[m]} \cup \{\boldsymbol{0}\}$ for $\mathcal{D}$. By Carathoedory's theorem there exists a distribution $\gamma \in \Delta(\Gamma)$ such that $\boldsymbol{\eta}^\star = \sum_{\boldsymbol{\eta}\in\Gamma} \gamma_{\boldsymbol{\eta}} \boldsymbol{\eta}$. Assume that ϵ and ρ are such that $1/\epsilon\rho$ is an integer and take $1/\rho\epsilon$ samples from the distribution γ and call $\tilde{\boldsymbol{\eta}}$ the resulting empirical mean.

First, we argue that $\tilde{\boldsymbol{\eta}} \in \mathcal{D}_\epsilon$. Indeed $\tilde{\boldsymbol{\eta}}_j = \frac{k_j}{1/\rho\epsilon}\frac{1}{\rho} = \epsilon\left(\frac{k_j}{1/\rho\epsilon}\frac{1}{\rho\epsilon}\right) = \epsilon k_j$ where $k_j \in \mathbb{N}$ and thus we have that $\tilde{\boldsymbol{\eta}} \in \mathcal{D}_\epsilon$.[11]

Now we show that with high probability $\tilde{\boldsymbol{\eta}} \in \mathcal{D}_\epsilon$ is close (in terms of utilities) to the optimal multiplier $\boldsymbol{\eta}^\star$. First observe that:

$$\begin{aligned}
\boldsymbol{\eta}_i^{\star,\top} \boldsymbol{c}_i(\phi_i \diamond \boldsymbol{z}) &:= \sum_{a_i\in\mathcal{A}^i, b_i\in\mathcal{A}^i} \left(\phi_i[b,a_i] \sum_{\boldsymbol{a}_{-i}\in\mathcal{A}^{-i}} \boldsymbol{\eta}^{\star,\top} \boldsymbol{c}_i(a_i,\boldsymbol{a}_{-i})\boldsymbol{z}[b,\boldsymbol{a}_{-i}]\right) && \text{(17a)}\\
&\le \sum_{a_i\in\mathcal{A}^i, b_i\in\mathcal{A}^i} \left(\phi_i[b,a_i]\left(\delta_{a_i,b} + \sum_{\boldsymbol{a}_{-i}\in\mathcal{A}^{-i}} \tilde{\boldsymbol{\eta}}^\top \boldsymbol{c}_i(a_i,\boldsymbol{a}_{-i})\boldsymbol{z}[b,\boldsymbol{a}_{-i}]\right)\right) && \text{(17b)}\\
&= \sum_{a_i\in\mathcal{A}^i, b_i\in\mathcal{A}^i} \left(\phi_i[b,a_i] \sum_{\boldsymbol{a}_{-i}\in\mathcal{A}^{-i}} \tilde{\boldsymbol{\eta}}^\top \boldsymbol{c}_i(a_i,\boldsymbol{a}_{-i})\boldsymbol{z}[b,\boldsymbol{a}_{-i}]\right) + \sum_{a_i\in\mathcal{A}^i, b_i\in\mathcal{A}^i} \phi_i[b,a_i]\delta_{a_i,b} && \text{(17c)}\\
&= \tilde{\boldsymbol{\eta}}_i^\top \boldsymbol{c}_i(\phi_i \diamond \boldsymbol{z}) + \sum_{a_i\in\mathcal{A}^i, b_i\in\mathcal{A}^i} \phi_i[b,a_i]\delta_{a_i,b} && \text{(17d)}
\end{aligned}$$

where the inequality comes from applying the Hoeffeding's inequality to every $a_i, b \in \mathcal{A}_i$:

$$\left|\sum_{a_{-i}\in\mathcal{A}^{-i}} \left(\tilde{\boldsymbol{\eta}} - \boldsymbol{\eta}^\star\right)^\top \boldsymbol{c}_i(a_i,a_{-i})\boldsymbol{z}[b,\boldsymbol{a}_{-i}]\right| \le \delta_{a_i,b}$$

where $\delta_{a_i,b} = \frac{2}{\rho}\sqrt{\frac{2}{1/\rho\epsilon}\log\left(\frac{2}{p_{a_i,b}}\right)}\left(\sum_{\boldsymbol{a}_{-i}\in\mathcal{A}^{-i}} \boldsymbol{z}[b,\boldsymbol{a}_{-i}]\right)$ since the range of the each sample is $\frac{1}{\rho}\left(\sum_{\boldsymbol{a}_{-i}\in\mathcal{A}^{-i}} \boldsymbol{z}[b,\boldsymbol{a}_{-i}]\right)$.

Moreover, for Hoeffeding's inequality, for every $a_i, b \in \mathcal{A}^i$ the above inequality holds with probability at least $1 - p_{a_i,b}$ and thus holds for all the $a_i, b \in \mathcal{A}^i$ simultaneously, with probability at least $p := \sum_{a_i,b\in\mathcal{A}^i} p_{a_i,b}$. If then we take $p_{a_i,b} := \frac{1}{2|\mathcal{A}^i|^2}$ for all $a_i, b \in \mathcal{A}^i$, then we have that $p = 1/2 > 0$ and $\delta := \delta_{a_i,b} = \frac{2}{\rho}\sqrt{\frac{2}{1/\rho\epsilon}\log\left(|\mathcal{A}^i|\right)}\left(\sum_{\boldsymbol{a}_{-i}\in\mathcal{A}^{-i}} \boldsymbol{z}[b,\boldsymbol{a}_{-i}]\right)$

Now the following holds with probability at lest $1/2$:

$$\left|\sum_{a_{-i}\in\mathcal{A}^{-i}} \left(\tilde{\boldsymbol{\eta}} - \boldsymbol{\eta}^\star\right)^\top \boldsymbol{c}_i(a_i,a_{-i})\boldsymbol{z}[b,\boldsymbol{a}_{-i}]\right| \le \delta\left(\sum_{\boldsymbol{a}_{-i}\in\mathcal{A}^{-i}} \boldsymbol{z}[b,\boldsymbol{a}_{-i}]\right), \quad \forall a_i, b \in \mathcal{A}^i$$

The proof is concluded by observing plugging this definition of $\delta = \delta_{a_i,b}$ in Equation (17) yields $\sum_{a_i\in\mathcal{A}^i, b_i\in\mathcal{A}^i} \phi_i[b,a_i]\delta_{a_i,b} = \delta$, and we can conclude that:

$$\boldsymbol{\eta}_i^{\star,\top} \boldsymbol{c}_i(\phi_i \diamond \boldsymbol{z}) \le \tilde{\boldsymbol{\eta}}_i^\top \boldsymbol{c}_i(\phi_i \diamond \boldsymbol{z}) + \delta.$$

This holds with positive probability, and thus shows the existence of such $\tilde{\boldsymbol{\eta}} \in \mathcal{D}_\epsilon$ for which the above inequality holds and thus $\mathcal{D}_\epsilon$ is $\left(2\sqrt{\frac{2\epsilon}{\rho}\log(|\mathcal{A}^i|)}\right)$-optimal. □

E. Proofs Omitted from Section 5

Proposition 5.1. *For instances $I := (\Gamma, \Phi)$ satisfying Assumption 3, the set of constrained ϵ-Phi-equilibria is convex.*

[11] If ϵ if not such that $1/\rho\epsilon \in \mathbb{N}$ then the one can take $\lceil 1/\rho\epsilon \rceil$ samples from $\gamma \in \Delta(\Gamma)$ and then the statement hold for a slightly smaller $\epsilon' < \epsilon$ defined as $\epsilon' := \frac{1}{\lceil 1/\rho\epsilon \rceil}\frac{1}{\rho}$.

Proof. Let $\boldsymbol{z}'$ and $\boldsymbol{z}''$ be Constrained ϵ-Phi-equilibria that is for all $i \in [N]$:

$$\epsilon + u_i(\boldsymbol{z}') \geq u_i(\phi_i' \diamond \boldsymbol{z}')$$

for $\phi' \in \arg\max_{\phi_i \in \Phi_i^{\mathcal{S}}} u_i(\phi_i \diamond \boldsymbol{z}')$. Equivalently it holds for all $i \in [N]$ that:

$$\epsilon + u_i(\boldsymbol{z}'') \geq u_i(\phi_i'' \diamond \boldsymbol{z}'')$$

where $\phi'' \in \arg\max_{\phi_i \in \Phi_i^{\mathcal{S}}} u_i(\phi_i \diamond \boldsymbol{z}'')$. For any $\boldsymbol{z} := \alpha \boldsymbol{z}' + (1-\alpha)\boldsymbol{z}''$ we have that:

$$\begin{aligned}
\epsilon + u_i(\boldsymbol{z}) &= \alpha\left(\epsilon + u_i(\boldsymbol{z}')\right) + (1-\alpha)\left(\epsilon + u_i(\boldsymbol{z}'')\right) \\
&\geq \alpha u_i(\phi_i' \diamond \boldsymbol{z}') + (1-\alpha) u_i(\phi_i'' \diamond \boldsymbol{z}'') \\
&\geq \max_{\phi_i \in \Phi_i^{\mathcal{S}}} u_i(\phi_i \diamond \boldsymbol{z}),
\end{aligned}$$

where the inequality holds for the linearity of u_i, the first inequality as both $\boldsymbol{z}'$ and $\boldsymbol{z}''$ are Constrained ϵ-Phi-equilibria and the last inequality holds since the $\max$ is a convex operator. □

Theorem 5.1. *Restricted to instances $I := (\Gamma, \Phi)$ which satisfy Assumption 3,* APXCPE$(1, 0)$ *admits a polynomial-time algorithm.*

Proof. APXCPE$(1, 0)$ amounts to solving the following problem:

$$\begin{aligned}
&\max_{\boldsymbol{z} \in \mathcal{S}} \ell(\boldsymbol{z}) \quad \text{s.t.} \\
&\qquad u_i(\boldsymbol{z}) \geq \max_{\phi_i \in \Phi_i^{\mathcal{S}}} u_i(\phi_i \diamond \boldsymbol{z}) \quad \forall i \in \mathcal{N},
\end{aligned}$$

which can be written as an LP with (possibly) exponentially-many constraints, by writing a constraint for each vertex of $\Phi_i^{\mathcal{S}}$. We can find an exact solution to such an LP in polynomial time by means of the ellipsoid algorithm that uses suitable separation oracle. Such an oracle solves the following optimization problem for every $i \in \mathcal{N}$:

$$\phi_i^\star \in \arg\max_{\phi_i \in \Phi_i^{\mathcal{S}}} u_i(\phi_i \diamond \boldsymbol{z}).$$

Then, the oracle returns as a separating hyperplane the incentive constraint corresponding to a $\phi_i^\star$ (if any) such that $u_i(\boldsymbol{z}) \geq u_i(\phi_i^\star \diamond \boldsymbol{z})$. Since all the steps of the separation oracle can be implemented in polynomial time, the ellipsoid algorithm runs in polynomial time, concluding the proof. □

Theorem 5.2. *Given an instance $I := (\Gamma, \Phi)$ satisfying Assumption 3, after $T \in \mathbb{N}_{>0}$ rounds, Algorithm 1 returns a correlated strategy $\bar{\boldsymbol{z}}_T \in \Delta_{\mathcal{A}}$ that is a constrained ϵ_T-Phi-equilibrium with $\epsilon_T = O(1/\sqrt{T})$. Moreover, each round of Algorithm 1 runs in polynomial time.*

Proof. Any regret minimizer $\mathfrak{R}_i$ for $\Phi_i^{\mathcal{S}}$ guarantees that, for every $\phi_i \in \Phi_i^{\mathcal{S}}$:

$$\sum_{t=1}^{T} u_i(\phi_i \diamond \boldsymbol{z}_t) - \sum_{t=1}^{T} u_i(\phi_{i,t} \diamond \boldsymbol{z}_t) \leq \epsilon_{i,T}\, T, \tag{19}$$

where $\epsilon_{i,T} = o(T)$. Since $\boldsymbol{x}_{i,t}[a] = \sum_{b \in \mathcal{A}_i} \phi_{i,t}[b,a] \boldsymbol{x}_{i,t}[b]$ for all $a \in \mathcal{A}_i$, for every $t \in [T]$ and $\boldsymbol{a} = (a_i, \boldsymbol{a}_{-i}) \in \mathcal{A}$:

$$\begin{aligned}
(\phi_{i,t} \diamond \boldsymbol{z}_t)[a_i, \boldsymbol{a}_{-i}] &= \sum_{b \in \mathcal{A}_i} \phi_{i,t}[b, a_i] \boldsymbol{z}[b, \boldsymbol{a}_{-i}] \\
&= \sum_{b \in \mathcal{A}_i} \phi_{i,t}[b, a_i] \Big(\boldsymbol{x}_{i,t}[b] \otimes \boldsymbol{x}_{-i,t}[\boldsymbol{a}_{-i}] \Big) \\
&= \left(\sum_{b \in \mathcal{A}_i} \phi_{i,t}[b, a_i] \boldsymbol{x}_{i,t}[b] \right) \otimes \boldsymbol{x}_{-i,t}[\boldsymbol{a}_{-i}]
\end{aligned}$$

$$
\begin{aligned}
&= \boldsymbol{x}_{i,t}[a_i] \otimes \boldsymbol{x}_{-i,t}[\boldsymbol{a}_{-i}] \\
&= \boldsymbol{z}_t[a_i, \boldsymbol{a}_{-i}],
\end{aligned}
$$

Plugging the equation above into Equation (19), we get:

$$
\sum_{t=1}^{T} u_i(\phi_i \diamond \boldsymbol{z}_t) - \sum_{t=1}^{T} u_i(\boldsymbol{z}_t) \leq \epsilon_{i,T}\, T.
$$

Now, since $\bar{\boldsymbol{z}}_T := \frac{1}{T}\sum_{t=1}^{T} \boldsymbol{z}_t$ and $u_i(\boldsymbol{z})$ is linear in $\boldsymbol{z}$, we can conclude that, for every $i \in \mathcal{N}$ and $\phi_i \in \Phi_i^{\mathcal{S}}$:

$$
u_i(\bar{\boldsymbol{z}}_T) \geq u_i(\phi_i \diamond \bar{\boldsymbol{z}}_T) - \epsilon_{i,T},
$$

and, thus, by letting $\epsilon_T := \max_{i\in\mathcal{N}} \epsilon_{i,T}$ we get that $\bar{\boldsymbol{z}}_T$ satisfies the incentivize constrained for being a constrained ϵ_T-Phi-equilibrium. We are left to verify that $\bar{\boldsymbol{z}}_T \in \mathcal{S}$, namely $\boldsymbol{c}_i(\bar{\boldsymbol{z}}_T) \leq \boldsymbol{0}$ for all $i \in \mathcal{N}$. This readily proved as:

$$
\begin{aligned}
\boldsymbol{c}_i(\bar{\boldsymbol{z}}_T) &= \frac{1}{T}\sum_{t=1}^{T} \boldsymbol{c}_i(\boldsymbol{z}_t) \\
&= \frac{1}{T}\sum_{t=1}^{T} \boldsymbol{c}_i(\phi_{i,t} \diamond \boldsymbol{z}_t) \\
&= \frac{1}{T}\sum_{t=1}^{T} \tilde{\boldsymbol{c}}_i(\phi_{i,t}) \\
&\leq \boldsymbol{0},
\end{aligned}
$$

where the first equality holds since $\boldsymbol{c}_i$ is linear, the second equality holds thanks to $\boldsymbol{z}_t = \phi_{i,t} \diamond \boldsymbol{z}_t$, the third one by Assumption 3, while the inequality holds since $\phi_{i,t} \in \Phi_i^{\mathcal{S}}$. This concludes the proof of the first part of the statement.

In conclusion, Algorithm 1 runs in polynomial time as finding $\boldsymbol{x}_{i,t}[a] = \sum_{b\in\mathcal{A}_i} \phi_{i,t}[b,a_i]\boldsymbol{x}_{i,t}[b]$ for all $a \in \mathcal{A}_i$ is equivalent to finding a stationary distribution of a Markov Chain, which can be done in polynomial time. Moreover, we can implement the regret minimizers $\mathfrak{R}_i$ over the polytopes $\Phi_i^{\mathcal{S}}$ so that their operations run in polynomial time, such as, *e.g.*, *online gradient descent*; see (Hazan et al., 2016). □

Theorem 5.4. *For instances $I := (\Gamma, \Phi_{\text{CCE}})$ such that $\boldsymbol{c}_i(\boldsymbol{a}) = \boldsymbol{c}_i(\boldsymbol{a}')$ for every player $i \in \mathcal{N}$ and action profiles $\boldsymbol{a}, \boldsymbol{a}' \in \mathcal{A} : a_i = a_i'$, Assumption 3 holds.*

Proof. Since the costs $\boldsymbol{c}_i(\boldsymbol{a})$ of player $i \in \mathcal{N}$ only depends on player i's action a_i and *not* on the actions of other players, it is possible to show that there exists $\tilde{\boldsymbol{c}}_i : \Phi_{\text{CCE}} \to [-1,1]^m$ such that the following holds for every $\boldsymbol{z} \in \Delta_{\mathcal{A}}$:

$$
\tilde{\boldsymbol{c}}_i(\phi_i) := \boldsymbol{c}_i(\phi_i \diamond \boldsymbol{z}).
$$

Indeed, for every $\phi_i \in \Phi_{\text{CCE}}$, by definition of Φ_{CCE} there exists a probability distribution $\boldsymbol{h} \in \Delta_{\mathcal{A}_i} : \phi_i[b,a] = \boldsymbol{h}[a]$ for all $b, a \in \mathcal{A}_i$. Then, for every $a_i \in \mathcal{A}_i$ and $\boldsymbol{a}_{-i} \in \mathcal{A}_{-i}$, we can write:

$$
\begin{aligned}
(\phi_i \diamond \boldsymbol{z})[a_i, \boldsymbol{a}_{-i}] &= \sum_{b\in\mathcal{A}_i} \phi_i[b, a_i]\boldsymbol{z}[b, \boldsymbol{a}_{-i}] \\
&= \sum_{b\in\mathcal{A}_i} \boldsymbol{h}[a_i]\boldsymbol{z}[b, \boldsymbol{a}_{-i}] \\
&= \boldsymbol{h}[a_i] \sum_{b\in\mathcal{A}_i} \boldsymbol{z}[b, \boldsymbol{a}_{-i}].
\end{aligned}
$$

Moreover, it holds:

$$
\boldsymbol{c}_i(\phi_i \diamond \boldsymbol{z})[a_i, \boldsymbol{a}_{-i}] = \sum_{\boldsymbol{a}\in\mathcal{A}} \boldsymbol{c}_i(\boldsymbol{a})(\phi_i \diamond \boldsymbol{z})[a_i, \boldsymbol{a}_{-i}]
$$

$$
\begin{aligned}
&= \sum_{\boldsymbol{a} \in \mathcal{A}} \boldsymbol{c}_i(\boldsymbol{a}) \boldsymbol{h}[a_i] \sum_{b \in \mathcal{A}_i} \boldsymbol{z}[b, \boldsymbol{a}_{-i}] \\
&= \sum_{a_i \in \mathcal{A}_i} \boldsymbol{c}_i(a_i, \cdot) \boldsymbol{h}[a_i] \sum_{\boldsymbol{a}_{-i} \in \mathcal{A}_{-i}} \sum_{b \in \mathcal{A}_i} \boldsymbol{z}[b, \boldsymbol{a}_{-i}] \\
&= \sum_{a_i \in \mathcal{A}_i} \boldsymbol{c}_i(a_i, \cdot) \boldsymbol{h}[a_i],
\end{aligned}
$$

which only depends on ϕ_i, as desired. Notice that, in the equations above, for every $a \in \mathcal{A}_i$ we let $\boldsymbol{c}_i(a, \cdot)$ be the (unique) value of $\boldsymbol{c}_i(\boldsymbol{a})$ for all $\boldsymbol{a} \in \mathcal{A} : a_i = a$. $\square$