

The Complexity of Nash Equilibria in Multiplayer Zero-Sum Games and Coordination Games

by

Yang Cai


Submitted to the Department of Electrical Engineering and Computer Science
in partial fulfillment of the requirements for the degree of

Master of Science in Computer Science and Engineering

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 2010



© Massachusetts Institute of Technology 2010. All rights reserved.

ARCHIVES

Author
Department of Electrical Engineering and Computer Science
September 1, 2010

Certified by
Constantinos Daskalakis
Assistant Professor
Thesis Supervisor

Accepted by
Terry P. Orlando
Chairman, EECS Committee on Graduate Students

The Complexity of Nash Equilibria in Multiplayer Zero-Sum Games and Coordination Games

by

Yang Cai

Submitted to the Department of Electrical Engineering and Computer Science
on September 1, 2010, in partial fulfillment of the
requirements for the degree of
Master of Science in Computer Science and Engineering

Abstract

We prove a generalization of von Neumann’s minmax theorem to the class of separable multiplayer zero-sum games, introduced in [Bregman and Fokin 1998]. These games are polymatrix—that is, graphical games in which every edge is a two-player game between its endpoints—in which every outcome has zero total sum of players’ payoffs. Our generalization of the minmax theorem implies convexity of equilibria, polynomial-time tractability, and convergence of no-regret learning algorithms to Nash equilibria. Given that three-player zero-sum games are already PPAD-complete, this class of games, i.e. with pairwise separable utility functions, defines essentially the broadest class of multi-player constant-sum games to which we can hope to push tractability results. Our result is obtained by establishing a certain game-class collapse, showing that separable constant-sum games are payoff equivalent to pairwise constant-sum polymatrix games—polymatrix games in which all edges are constant-sum games, and invoking a recent result of [Daskalakis, Papadimitriou 2009] for these games.

We also explore generalizations to classes of non-constant-sum multi-player games. A natural candidate is polymatrix games with strictly competitive games on their edges. In the two player setting, such games are minmax solvable and recent work has shown that they are merely affine transformations of zero-sum games [Adler, Daskalakis, Papadimitriou 2009]. Surprisingly we show that a polymatrix game comprising of strictly competitive games on its edges is PPAD-complete to solve, proving a striking difference in the complexity of networks of zero-sum and strictly competitive games. Finally, we look at the role of coordination in networked interactions, studying the complexity of polymatrix games with a mixture of coordination and zero-sum games. We show that finding a pure Nash equilibrium in coordination-only polymatrix games is PLS-complete; hence, computing a mixed Nash equilibrium is in $PLS \cap PPAD$, but it remains open whether the problem is in P . If, on the other hand, coordination and zero-sum games are combined, we show that the problem becomes PPAD-complete, establishing that coordination and zero-sum games achieve the full generality of PPAD.

This work is done in collaboration with Costis Daskalakis.

Thesis Supervisor: Constantinos Daskalakis
Title: Assistant Professor

Acknowledgments

I would like to thank my advisor, Costis Daskalakis, for all his guidance and encouragement. My last year has been a wonderful experience largely because of Costis. He has made himself available over these past two semesters whenever I want to have a discussion. I thank him for the hours we spent together thinking. In fact, my thesis topic was born in the first discussion after Costis arrived in MIT. Research can be frustrating, and depression is almost unavoidable. But Costis's optimism is so contagious that I can not help but think the same way as he does. Besides research, I must thank Costis for guiding me through my first teaching experience last semester. I learned a great deal on how to present and deliver the concepts and ideas to the audience from him.

I thank the many TOC students who met with me frequently for brainstorming, and their advices from taking courses to finding a research direction. Also, I would like to thank all my friends at MIT and back home in China (and scattered elsewhere) for all the love they have shown me in the time I have known them.

Last but not least, I am grateful to my parents. They always seem to understand what I have been going through. I want to thank them the most for always being supportive for all my choices, and offering wise advices and selfless helps through my toughest time. I have no way to thank them for their priceless love, care, inspiration, patience, and encouragement, but to dedicate this thesis to them.

Contents

1	Introduction	9
2	Definitions	17
2.1	Games	17
2.2	Nash Equilibria	18
2.3	Game Dynamics	19
2.4	Complexity Classes	20
3	Zero-sum Polymatrix Games	23
3.1	The Payoff Preserving Transformation	24
3.2	Algorithmic and Equilibrium Properties of Separable Constant-sum Games	27
3.3	A Direct Reduction to Linear Programming	28
3.4	A Constructive Proof of the Convergence of No-Regret Algorithms	30
4	Coordination Polymatrix Games	33
5	Combining Coordination and Zero-sum Games	37
5.1	Gadgets	38
5.2	Construction of \mathcal{G}^*	39
5.3	Correctness of the Reduction	40
6	Strictly Competitive Polymatrix Games	43
A	Omitted Proofs	49
A.1	Separable Zero-Sum Multiplayer Games	49
A.1.1	LP Formulation	49

A.2 Polymatrix Games with Coordination and Zero-Sum Edges	53
---	----

Chapter 1

Introduction

In the past few decades, computer scientists have constructed a rich theory for designing and analyzing systems, in which a central designer can control the behaviors of all components. In such systems, components are usually cooperating with each other to accomplish the designer's goal. However, with the recent advent of Internet, systems are increasingly being used by different parties with their own interests. Unfortunately, our knowledge for the global behaviors of such systems is quite limited, especially when the computational limitation of such parties are taken into account.

This raised new challenges and opportunities for computer science. Among the questions we are looking to understand are the following, how far is the performance of a system from optimality, due to the conflict of interest between users and administrators? Also, how can we design systems whose performances are robust under the potential conflict of interest inside the system? Game theory provides concepts and methodologies that appear to be powerful tools for designing and analyzing systems that consist of multiple parties who are only interested in maximizing their own welfare. To incorporate the game-theoretic tools into computer science, it is very critical to understand what portion of these tools is applicable to computer systems.

In game theory and economics in general, Nash equilibrium is perhaps the most commonly used tool for modeling the overall behavior of a system with interacting self-interested individuals. However, it has recently been shown that Nash equilibrium is computationally intractable. The hardness comes in the form of completeness for the class *PPAD* of the total search problems [21, 10, 7]. This intractability result severely hurts the prediction power of

this concept. Even though the concept is mathematically well-defined and universal, it is hard to compute. Hence, it is totally reasonable to be skeptical about the prediction made by the Nash equilibrium; if Nash equilibria are hard to compute, how can individuals find one and use it?

Although finding an equilibrium in most cases is intractable, an interesting and important class of games called two-player zero-sum games luckily circumvents the computational complexity barrier. A two player zero-sum game is a game such that in any outcome of the game, the two players' payoffs sum up to zero or a fixed constant number. This class of games can accurately model many scenarios of two players competing with each other, e.g. two-player chess and poker.

In 1928, von Neumann [22] established a beautiful theorem showing that every two-player zero-sum game satisfies the min-max property; thus we now know that an equilibrium can be computed efficiently with linear programming [9, 17]. In fact, according to Aumann [3], two-person strictly competitive games—these are affine transformations of two-player zero-sum games¹ [2]—are “one of the few areas in game theory, and indeed in the social sciences, where a fairly sharp, unique prediction is made.” The aforementioned intractability results on the computation of Nash equilibria [10, 7] can be viewed as complexity-theoretic support of Aumann's claim, steering research towards the following questions: In what classes of multiplayer games are equilibria tractable? And when equilibria *are* tractable, do there also exist decentralized, simple dynamics converging to equilibrium?

Recent work [11] has explored these questions on the following (network) generalization of two-player zero-sum games: The players are located at the nodes of a graph whose edges are zero-sum games between their endpoints; every player/node can choose a unique mixed strategy to be used in all games/edges she participates in, and her payoff is computed as the sum of her payoff from all adjacent edges. These games, called *pairwise zero-sum polymatrix games*, certainly contain two-player zero-sum games, which are amenable to linear programming and enjoy several important properties such as convexity of equilibria, uniqueness of values, and convergence of no-regret learning algorithms to equilibria [22]. Linear programming can also handle star topologies, but more complicated topologies introduce combinatorial structure that makes equilibrium computation harder. Indeed, the

¹See Chapter 2 for the precise statement of this equivalence result.

straightforward LP formulation that handles two-player games and star topologies breaks down already in the triangle topology (see discussion in [11]).

The class of pairwise zero-sum polymatrix games was studied in the early papers of Bregman and Fokin [5, 6], where the authors provide a linear programming formulation for finding equilibrium strategies. The size of their linear programs is exponentially large in both variables and constraints, albeit with a small rank, and a variant of the column-generation technique in the simplex method is provided for the solution of these programs. The work of [11] circumvents the large linear programs of [6] with a reduction to a polynomial-sized two-player zero-sum game, establishing the following properties for these games:

- (1) the set of Nash equilibria is convex;
- (2) a Nash equilibrium can be computed in polynomial-time using linear programming;
- (3) if the nodes of the network run any no-regret learning algorithm, the global behavior converges to a Nash equilibrium.²

In other words, pairwise zero-sum polymatrix games inherit several of the important properties of two-player zero-sum games.³ In particular, the third property above together with the simplicity, universality and distributed nature of the no-regret learning algorithms provide strong support on the plausibility of the Nash equilibrium predictions in this setting.

On the other hand, the hope for extending the positive results of [11] to larger classes of games imposing no constraints on the edge-games seems rather slim. Indeed it follows from the work of [10] that general polymatrix games are PPAD-complete. The same obstacle arises if we deviate from the polymatrix game paradigm. If our game is not the result of pairwise (i.e. two-player) interactions, the problem becomes PPAD-complete even for three-player zero-sum games. This is because every two-player game can be turned into a three-player zero-sum game by introducing a third player whose role is to balance the overall payoff to zero. Given these observations it appears that pairwise zero-sum polymatrix games are at the boundary of multi-player games with tractable equilibria.

Games That Are Globally Zero-Sum. The class of pairwise zero-sum polymatrix games was studied in the papers of Bregman and Fokin [5, 6] as a special case of *separable*

²The notion of a *no-regret learning algorithm*, and the type of convergence used here is quite standard in the learning literature and will be described in detail in chapter 2.

³If the game is non-degenerate (or perturbed) it can also be shown that the values of the nodes are unique. But, unlike two-player zero-sum games, there are examples of (degenerate) pairwise zero-sum polymatrix games with multiple Nash equilibria that give certain players different payoffs [12].

zero-sum multiplayer games. These are similar to pairwise zero-sum polymatrix games, albeit with no requirement that every edge is a zero-sum game; instead, it is only asked that the total sum of all players' payoffs is zero (or some other constant⁴) in every outcome of the game. Intuitively, these games can be used to model a broad class of competitive environments where there is a constant amount of wealth (resources) to be split among the players of the game, with no in-flow or out-flow of wealth that may change the total sum of players' wealth in an outcome of the game.

A simple example of this situation is the following game taking place in the wild west. A set of gold miners in the west coast need to transport gold to the east coast using wagons. Every miner can split her gold into a set of available wagons in whatever way she wants (or even randomize among partitions). Every wagon uses a specific path to go through the Rocky mountains. Unfortunately each of the available paths is controlled by a group of thieves. A group of thieves may control several of these paths and if they happen to wait on the path used by a particular wagon they can ambush the wagon and steal the gold being carried. On the other hand, if they wait on a particular path they will miss on the opportunity to ambush the wagons going through the other paths in their realm as all wagons will cross simultaneously. The utility of each miner in this game is the amount of her shipped gold that reaches her destination in the east coast, while the utility of each group of thieves is the total amount of gold they steal. Clearly, the total utility of all players in the wild west game is constant in every outcome of the game (it equals the total amount of gold shipped by the miners), but the pairwise interaction between every miner and group of thieves is not. In other words, the constant-sum property is a *global* rather than a local property of this game.

The reader is referred to [6] for further applications and a discussion of several special cases of these games, such as the class of pairwise zero-sum games discussed above. Given the positive results for the latter, explained earlier in this introduction, it is rather appealing to try to extend these results to the full class of separable zero-sum games, or at least to other special classes of these games. We show that this generalization is indeed possible, but for an unexpected reason that represents a game-class collapse. Namely,

Theorem 1. *There is a polynomial-time computable payoff preserving transformation from*

⁴In this case, the game is called *separable constant-sum multiplayer*.

every separable zero-sum multiplayer game to a pairwise constant-sum polymatrix game.⁵

In other words, given a separable zero-sum multiplayer game \mathcal{GG} , there exists a polynomial-time computable pairwise constant-sum multiplayer game \mathcal{GG}' such that, for any selection of strategies by the players, every player receives the same payoff in \mathcal{GG} and in \mathcal{GG}' . (Note that, for the validity of the theorem, it is important that we allow *constant-sum*—as opposed to only zero-sum—games on the edges of the game.) Theorem 1 implies that the class of separable zero-sum multiplayer games, suggested in [6] as a superset of pairwise zero-sum games, is only slightly larger, in that it is a subset, up to different representations of the game, of the class of pairwise constant-sum games. In particular, all the classes of games treated as special cases of separable zero-sum games in [6] can be reduced via payoff-preserving transformations to pairwise constant-sum polymatrix games. Since it is not hard to extend the results of [11] to pairwise constant-sum games, as a corollary we obtain:

Corollary 1. *Pairwise constant-sum polymatrix games and separable constant-sum multiplayer games are payoff preserving transformation equivalent, and satisfy properties (1), (2) and (3).*

We provide the payoff preserving transformation from separable zero-sum to pairwise constant-sum games in Section 3.1. The transformation is rather involved, but morally it works out by unveiling the local-to-global consistency constraints that the payoff tables of the game need to satisfy in order for the global zero-sum property to arise. Given our transformation, in order to obtain Corollary 1, we only need a small extension to the result of [11], establishing properties (1), (2) and (3) for pairwise constant-sum games. This can be done in an indirect way by subtracting the constants from the edges of a pairwise constant-sum game \mathcal{GG} to turn it into a pairwise zero-sum game \mathcal{GG}' , and then showing that the set of equilibria, as well as the behavior of no-regret learning algorithms in these two games are the same. We can then readily use the results of [11] to prove Corollary 1. The details of the proof are given in section 3.2.

We also present a direct reduction of separable zero-sum games to linear programming, i.e. one that does not go the round-about way of establishing our payoff-preserving transformation, and then using the result of [11] as a black-box. This poses interesting challenges as

⁵Pairwise constant-sum games are similar to pairwise zero-sum games, except that every edge can be constant-sum, for arbitrary constant.

the validity of the linear program proposed in [11] depended crucially on the pairwise zero-sum nature of the interactions between nodes in a pairwise zero-sum game. Surprisingly, we show that the same linear program works for separable zero-sum games by establishing an interesting kind of restricted zero-sum property satisfied by these games (Lemma 17). The resulting LP is simpler and more intuitive, albeit more intricate to argue about, than the one obtained the round-about way. The details are given in Section 3.3.

Finally, we provide a constructive proof of the validity of Property (3). Interestingly enough, the argument of [11] establishing this property used in its heart Nash’s theorem (for non zero-sum games), giving rise to a non-constructive argument. Here we rectify this by providing a constructive proof based on first principles. The details can be found in Section 3.4.

Allowing General Strict Competition. It is surprising that the properties (1)–(3) of 2-player zero-sum games extend to the network setting despite the combinatorial complexity that the networked interactions introduce. Indeed, zero-sum games are one of the few classes of well-behaved two-player games for which we could hope for positive results in the networked setting. A small variant of zero-sum games are *strictly competitive games*. These are two-player games in which, for every pair of mixed strategy profiles s and s' , if the payoff of one player is better in s than in s' , then the payoff of the other player is worse in s than in s' . These games were known to be solvable via linear programming [3], and recent work has shown that they are merely affine transformations of zero-sum games [2]. That is, if (R, C) is a strictly competitive game, there exists a zero-sum game (R', C') and constants $c_1, c_2 > 0$ and d_1, d_2 such that $R = c_1 R' + d_1 \mathbb{1}$ and $C = c_2 C' + d_2 \mathbb{1}$, where $\mathbb{1}$ is the all-ones matrix. Given the affinity of these classes of games, it is quite natural to suspect that Properties (1)–(3) should also hold for polymatrix games with strictly competitive games on their edges. Indeed, the properties *do* hold for the special case of pairwise constant-sum polymatrix games (Corollary 1).⁶ Surprisingly we show that if we allow arbitrary strictly competitive games on the edges, the full complexity of the PPAD class arises from this seemingly benign class of games.

Theorem 2. *Finding a Nash equilibrium in polymatrix games with strictly competitive games on their edges is PPAD-complete.*

⁶Pairwise constant-sum polymatrix games arise from this model if all c ’s in the strictly competitive games are chosen equal across the edges of the game, but the d ’s can be arbitrary.

The Role of Coordination. Another class of tractable and well-behaved two-player games that we could hope to understand in the network setting is the class of coordination games. If zero-sum games represent “perfect competition”, coordination games represent “perfect cooperation”, and they are trivial to solve in the two-player setting. Given the positive results on zero-sum polymatrix games, it is natural to investigate the complexity of polymatrix games containing both zero-sum and coordination games. In fact, this was the immediate question of Game Theorists (e.g. in [23]) in view of the earlier results of [11]. We explore this thoroughly in this thesis.

First, it is easy to see that coordination-only polymatrix games are (cardinal) potential games, so that a pure Nash equilibrium always exists. We show however that finding a pure Nash equilibrium is an intractable problem.

Theorem 3. *Finding a pure Nash equilibrium in coordination-only polymatrix games is PLS-complete.*

On the other hand, Nash’s theorem implies that finding a mixed Nash equilibrium is in PPAD. From this observation and the above, we obtain as a corollary the following interesting result.

Corollary 2. *Finding a Nash equilibrium in coordination-only polymatrix games is in $PLS \cap PPAD$.*

So finding a Nash equilibrium in coordination-only polymatrix games is probably neither PLS- nor PPAD-complete, and the above corollary may be seen as an indication that the problem is in fact tractable. Whether it belongs to P is left open by this work. Coincidentally, the problem is tantamount to finding a coordinate-wise local maximum of a multilinear polynomial of degree two on the hypercube⁷. Surprisingly no algorithm for this very basic and seemingly simple problem is known in the literature. . .

While we leave the complexity of coordination-only polymatrix games open for future work, we do give a definite answer to the complexity of polymatrix games with both zero-sum and coordination games on their edges, showing that the full complexity of PPAD can be obtained this way.

Theorem 4. *Finding a Nash equilibrium in polymatrix games with coordination or zero-sum games on their edges is PPAD-complete.*

⁷i.e. finding a point x where the polynomial cannot be improved by single coordinate changes to x .

It is quite remarkable that polymatrix games exhibit such a rich range of complexities depending on the types of games placed on their edges, from polynomial-time tractability when the edges are zero-sum to PPAD-completeness when general strictly competitive games or coordination games are also allowed. Moreover, it is surprising that even though non-polymatrix three-player zero-sum games give rise to PPAD-hardness, separable zero-sum multiplayer games with any number of players remain tractable...

The results described above sharpen our understanding of the boundary of tractability of multiplayer games. In fact, given the PPAD-completeness of three-player zero-sum games, we cannot hope to extend positive results to games with three-way interactions. But can we circumvent some of the hardness results shown above, e.g. the intractability result of Theorem 4, by allowing a limited amount of coordination in a zero-sum polymatrix game? A natural candidate class of games are *group-wise zero-sum polymatrix games*. These are polymatrix games in which the players are partitioned into groups so that the edges going across groups are zero-sum while those within the same group are coordination games. In other words, players inside a group are “friends” who want to coordinate their actions, while players in different groups are competitors. It is conceivable that these games are simpler (at least for a constant number of groups) since the zero-sum and the coordination interactions are not interleaved. We show however that the problem is intractable even for 3 groups of players.

Theorem 5. *Finding a Nash equilibrium in group-wise zero-sum polymatrix games with at most three groups of players is PPAD-complete.*

Chapter 2

Definitions

2.1 Games

- A *graphical polymatrix game* is defined in terms of an undirected graph $G = (V, E)$, where V is the set of players of the game and every edge is associated with a 2-player game between its endpoints. Assuming that the set of (pure) strategies of player $v \in V$ is $[m_v] := \{1, \dots, m_v\}$, where $m_v \in \mathbb{N}$, we specify the 2-player game along the edge $(u, v) \in E$ by providing a pair of payoff matrices: a $m_u \times m_v$ real matrix $A^{u,v}$ and another $m_v \times m_u$ real matrix $A^{v,u}$ specifying the payoffs of the players u and v along the edge (u, v) for different choices of strategies by the two players.

Now the aggregate payoff of the players is computed as follows. Let f be a pure strategy profile, that is $f(u) \in [m_u]$ for all u . The payoff of player $u \in V$ in the strategy profile f is $\mathcal{P}_u(f) = \sum_{(u,v) \in E} A_{f(u), f(v)}^{u,v}$. In other words, the payoff of u is the sum of the payoffs that u gets from all the 2-player games that u plays with her neighbors.

- A *separable zero-sum multiplayer game* is a graphical polymatrix game in which the sum of players' payoffs is zero in every outcome, i.e. in every pure strategy profile, of the game. Formally,

Definition 1 (Separable zero-sum multiplayer games). *A separable zero-sum multiplayer game \mathcal{GG} is a graphical polymatrix game in which, for any pure strategy profile f , the sum of all players' payoffs is zero. I.e., for all f , $\sum_{u \in V} \mathcal{P}_u(f) = 0$.*

If the sum of all players' payoffs is allowed to be an arbitrary constant, the corre-

sponding games would be called *separable constant-sum multiplayer games*.

A simple class of games with this property are those in which every edge is a zero-sum game. This special class of games, studied in [11], are called *pairwise zero-sum polymatrix games*, as the zero-sum property arises as a result of pairwise zero-sum interactions between the players. If the edges were allowed to be arbitrary constant-sum games, the corresponding games would be called *pairwise constant-sum polymatrix games*.

- Two player zero-sum games represent strict competition between the two participants, while there is another interesting class of games called *two-player coordination games*, in which the two players are cooperating each other to achieve higher payoffs. In a two-player coordination game, two players have the same payoff matrix. If u and v are the two player, $A^{u,v} = (A^{v,u})^T$. In chapter 4, we will explore another special case of graphical polymatrix game called *coordination polymatrix game*. A *coordination polymatrix game* is a graphical polymatrix game in which every edge is a two player coordination game. Formally,

Definition 2 (Coordination polymatrix games). *A coordination polymatrix game \mathcal{GG} , is a graphical polymatrix game in which, for every edge (u, v) , there is a two-player coordination game between u and v . I.e., $A^{u,v} = (A^{v,u})^T$.*

- *Two-player strictly competitive games* are a commonly used generalization of zero-sum games. A two-player game is *strictly competitive* if it has the following property [3] : if both players change their mixed strategies, then either their expected payoffs remain the same, or one player’s expected payoff increases and the other’s decreases. It was recently shown that strictly competitive games are merely affine transformations of two-player zero-sum games [2]. That is, if (R, C) is a strictly competitive game, there exists a zero-sum game (R', C') and constants $c_1, c_2 > 0$ and d_1, d_2 such that $R = c_1 R' + d_1 \mathbf{1}$ and $C = c_2 C' + d_2 \mathbf{1}$, where $\mathbf{1}$ is the all-ones matrix.

2.2 Nash Equilibria

A *(mixed) Nash equilibrium* is a collection of mixed—that is randomized—strategies for the players of the game, such that every pure strategy played with positive probability by a player is a best response in expectation for that player given the mixed strategies of the

other players. A *pure Nash equilibrium* is a special case of a mixed Nash equilibrium in which the players' strategies are pure, i.e. deterministic. Besides the concept of exact Nash equilibrium, there are several different—but related— notions of approximate equilibrium.

Two widely used notions of approximate Nash equilibrium are the following: (1) In an ϵ -*Nash equilibrium*, all pure strategies played with positive probability should give the corresponding player expected payoff that lies to within an additive ϵ from the expected payoff guaranteed by the best mixed strategy against the other players' mixed strategies. (2) A related, but weaker, notion of approximate equilibrium is the concept of an ϵ -*approximate Nash equilibrium*, in which the expected payoff achieved by every player through her mixed strategy lies to within an additive ϵ from the optimal payoff she could possibly achieve via any mixed strategy given the other players' mixed strategies. Clearly, an ϵ -Nash equilibrium is also a ϵ -approximate Nash equilibrium, but the opposite need not be true. Nevertheless, the two concepts are computationally equivalent as the following proposition suggests.

Proposition 1. [10] *Given an ϵ -approximate Nash equilibrium of an n -player game, we can compute in polynomial time a $\sqrt{\epsilon} \cdot (\sqrt{\epsilon} + 1 + 4(n - 1)\alpha_{\max})$ -Nash equilibrium of the game, where α_{\max} is the magnitude of the maximum in absolute value possible utility of a player in the game.*

In this thesis, we focus on exact mixed Nash equilibria. It is easy to see—and is well-known—that polymatrix games have mixed Nash equilibria in rational numbers and with polynomial description complexity in the size of the game.

2.3 Game Dynamics

Although mathematically well defined and universal, Nash equilibrium is a static notion of behavior. It only gives us a prediction of the final outcome, but provides no hint about how players have reached that. However, we are also interested in distributed procedures via which equilibria arise, since the existence of such procedures will increase the credibility of the prediction made by the Nash equilibrium. It is widely known that there is a family of such procedures, in which every player takes a *no-regret behavior*, converge to a Nash equilibrium for two-player zero-sum games.

Now let us formally define the notion of *no-regret behavior* first.

Definition 3 (No-Regret Behavior). *Let every node $u \in V$ of a graphical polymatrix game choose a mixed strategy $x_u^{(t)}$ at every time step $t = 1, 2, \dots$. We say that the sequence of strategies $\langle x_u^{(t)} \rangle$ is a no-regret sequence, if for every mixed strategy x of player u and at all times T*

$$\sum_{t=1}^T \left(\sum_{(u,v) \in E} (x_u^{(t)})^T \cdot A^{u,v} \cdot x_v^{(t)} \right) \geq \sum_{t=1}^T \left(\sum_{(u,v) \in E} x^T \cdot A^{u,v} \cdot x_v^{(t)} \right) - o(T),$$

where the function $o(T)$ could depend on the number strategies available to player u , the number of neighbors of u and magnitude of the maximum in absolute value entry in the matrices $A^{u,v}$. The function $o(T)$ is called the regret of player u at time T .

We note that obtaining a no-regret sequence of strategies is far from exotic. If a node uses any no-regret learning algorithm to select strategies (for a multitude of such algorithms see, e.g., [4]), the output sequence of strategies will constitute a no-regret sequence. A common such algorithm is the *multiplicative weights-update algorithm* (see, e.g., [13]). In this algorithm every player maintains a mixed strategy. At each period, each probability is multiplied by a factor exponential in the utility the corresponding strategy would yield against the opponents' mixed strategies (and the probabilities are renormalized).

2.4 Complexity Classes

In this thesis, we are going to study the computational complexity of finding a Nash equilibrium in various games. However, the mainstream complexity concepts and techniques developed by complexity theorist — chief among them *NP*-completeness — can not be directly apply for the study of complexity of finding a Nash equilibrium. The big difference between *NP*-complete problems and finding a Nash equilibrium is the following: the solution of an *NP*-complete problem is not always guaranteed to exist, and the difficulty for solving a *NP*-complete problem seems to heavily rely on the possibility that the solution may not exist. But Nash equilibrium always exists, because of Nash's theorem.

Motivated mainly by this question, Meggido and Papadimitriou [19] defined in the 1980s the complexity class *TFNP* (for *NP* total functions), consisting exactly of all search problems in *NP* for which every instance is guaranteed to have a solution.

To capture the complexity of finding a Nash equilibrium in $TFNP$, we need to take one step further. We will group $TFNP$ total functions whose proofs of totality are similar into subclasses. In fact, most of these proofs work by essentially constructing an exponentially large graph on the solution space (with edges that are computed by some algorithm), and then applying a simple combinatorics lemma establishing the existence of a particular kind of node. The node who is guaranteed to exist by the lemma is the desired solution of the given instance. Interestingly, essentially all known problems in $TFNP$ can be grouped into four classes: PLS , PPA , PPP and $PPAD$. In this thesis, we will only talk about PLS and $PPAD$ and they are defined as follows:

- *In any directed acyclic graph, there must be a sink.* The corresponding class, PLS for Polynomial Local Search, had already been defined in [14] and contains many important complete problems.
- *In any directed graph with one unbalanced node (node with outdegree different from its indegree), there must be another unbalanced node.* The corresponding class is called $PPAD$ for Polynomial Parity Argument for Directed graphs, and it contains finding a Nash equilibrium.

In fact finding a Nash equilibrium is not only contained in $PPAD$, it is shown in [10] that finding a Nash equilibrium in a graphical polymatrix game is indeed $PPAD$ -complete.

We will also introduce a PLS -complete problem — the *Max-Cut problem* with the flip neighborhood. A Max-Cut problem with the flip neighborhood is the following: Given a graph G , we want to find a local Max-Cut, which means if we move any node to the opposite side of the current cut, the size of the cut will not increase.

Chapter 3

Zero-sum Polymatrix Games

In this section, we are interested in understanding the equilibrium properties of separable zero-sum multiplayer games. By studying this class of games, we cover the full expanse of zero-sum polymatrix games, and essentially the broadest class of multi-player zero-sum games for which we could hope to push tractability results. Recall that if we deviate from edgewise separable utility functions the problem becomes PPAD-complete, as already 3-player zero-sum games are PPAD-complete.

We organize this section as follows: In section 3.1, we present a payoff-preserving transformation from separable zero-sum games to pairwise constant-sum games. This establishes Theorem 1, proving that separable zero-sum games are not much more general—as were thought to be [6]—than pairwise zero-sum games. This can easily be used to show Corollary 1. We proceed in Section 3.3 to provide a direct reduction from separable zero-sum games to linear programming, obviating the use of our payoff-preserving transformation. In a way, our linear program corresponds to the minmax program of a related two-player game. The resulting LP formulation is similar to the one suggested in (a footnote of) [11] for pairwise zero-sum games, except that now its validity seems rather slim as the resulting 2-player game is not zero-sum. Surprisingly we show that it *does* work by uncovering a restricted kind of zero-sum property satisfied by the game. Finally, in Section 3.4 we provide an alternative proof, i.e. one that does not go via the payoff-preserving transformation, that no-regret dynamics converge Nash equilibria in separable zero-sum games. The older proof of this fact for pairwise zero-sum games [11] was using Brouwer’s fixed point theorem, and was hence non-constructive. Our new proof rectifies this as it is based on first principles and is constructive.

3.1 The Payoff Preserving Transformation

Our goal in this section is to provide a payoff-preserving transformation from a separable zero-sum game \mathcal{GG} to a pairwise constant-sum polymatrix game \mathcal{GG}' . We start by establishing a surprising consistency property satisfied by the payoff tables of a separable zero-sum game. On every edge (u, v) , the sum of u 's and v 's payoffs on that edge when they play $(1, 1)$ and when they play (i, j) equals the sum of their payoffs when they play $(1, j)$ and when they play $(i, 1)$. Namely,

Lemma 1. *For any edge (u, v) of a separable zero-sum multiplayer game \mathcal{GG} , and for every $i \in [m_u]$, $j \in [m_v]$,*

$$(A_{1,1}^{u,v} + A_{1,1}^{v,u}) + (A_{i,j}^{u,v} + A_{j,i}^{v,u}) = (A_{1,j}^{u,v} + A_{j,1}^{v,u}) + (A_{i,1}^{u,v} + A_{1,i}^{v,u}).$$

Proof. Let all players except u and v fix their strategies to $S_{-\{u,v\}}$. For $w \in \{u, v\}$, $k \in [m_w]$, let

$$P_{(w:k)} = \sum_{r \in N(w) \setminus \{u,v\}} (s_w^T \cdot A^{w,r} \cdot s_r + s_r^T \cdot A^{r,w} \cdot s_w),$$

where in the above expression take s_w to simply be the deterministic strategy k . Using that the game is zero-sum, the following must be true:

- suppose u plays strategy 1, v plays strategy j ; then

$$P_{(u:1)} + P_{(v:j)} + A_{1,j}^{u,v} + A_{j,1}^{v,u} = \alpha \quad (1)$$

- suppose u plays strategy i , v plays strategy 1; then

$$P_{(u:i)} + P_{(v:1)} + A_{i,1}^{u,v} + A_{1,i}^{v,u} = \alpha \quad (2)$$

- suppose u plays strategy 1, v plays strategy 1; then

$$P_{(u:1)} + P_{(v:1)} + A_{1,1}^{u,v} + A_{1,1}^{v,u} = \alpha \quad (3)$$

- suppose u plays strategy i , v plays strategy j ; then

$$P_{(u:i)} + P_{(v:j)} + A_{i,j}^{u,v} + A_{j,i}^{v,u} = \alpha \quad (4)$$

In the above, $-\alpha$ represents the total sum of players' payoffs on all edges that do not involve u or v as one of their endpoints. Since $S_{-\{u,v\}}$ is held fixed here for our discussion, α is also fixed. By inspecting the above, we obtain that (1) + (2) = (3) + (4). If we cancel out the common terms in the equation, we obtain

$$(A_{1,1}^{u,v} + A_{1,1}^{v,u}) + (A_{i,j}^{u,v} + A_{j,i}^{v,u}) = (A_{1,j}^{u,v} + A_{j,1}^{v,u}) + (A_{i,1}^{u,v} + A_{1,i}^{v,u}).$$

□

Now for every ordered pair of players (u, v) , let us construct a new payoff matrix $B^{u,v}$ based on $A^{u,v}$ and $A^{v,u}$ as follows. First, we set $B_{1,1}^{u,v} = A_{1,1}^{u,v}$. Then $B_{i,j}^{u,v} = B_{1,1}^{u,v} + (A_{1,j}^{u,v} - A_{1,1}^{u,v}) + (A_{j,1}^{v,u} - A_{j,i}^{v,u})$. Notice that Lemma 1 implies: $(A_{1,j}^{u,v} - A_{1,1}^{u,v}) + (A_{j,1}^{v,u} - A_{j,i}^{v,u}) = (A_{1,1}^{v,u} - A_{1,i}^{v,u}) + (A_{i,j}^{u,v} - A_{i,1}^{u,v})$. So we can also write $B_{i,j}^{u,v} = B_{1,1}^{u,v} + (A_{1,1}^{v,u} - A_{1,i}^{v,u}) + (A_{i,j}^{u,v} - A_{i,1}^{u,v})$. Our construction satisfies two important properties. (a) If we use the second representation of $B^{u,v}$, it is easy to see that $B_{i,j}^{u,v} - B_{i,k}^{u,v} = A_{i,j}^{u,v} - A_{i,k}^{u,v}$. (b) If we use the first representation, it is easy to see that $B_{i,j}^{u,v} - B_{k,j}^{u,v} = A_{j,k}^{v,u} - A_{j,i}^{v,u}$. Given these observations we obtain the following:

Lemma 2. *For every edge (u, v) , $B^{u,v} + (B^{v,u})^T = c^{\{u,v\}} \mathbb{1}$, where $\mathbb{1}$ is the all-ones matrix.*

Proof. • Using the second representation for $B_{i,j}^{u,v}$,

$$B_{i,j}^{u,v} = B_{1,1}^{u,v} + (A_{1,1}^{v,u} - A_{1,i}^{v,u}) + (A_{i,j}^{u,v} - A_{i,1}^{u,v}).$$

- Using the first representation for $B_{j,i}^{v,u}$,

$$B_{j,i}^{v,u} = B_{1,1}^{v,u} + (A_{1,i}^{v,u} - A_{1,1}^{v,u}) + (A_{i,1}^{u,v} - A_{i,j}^{u,v}).$$

So we have $B_{i,j}^{u,v} + B_{j,i}^{v,u} = B_{1,1}^{u,v} + B_{1,1}^{v,u} =: c^{\{u,v\}}$. □

We are now ready to describe the pairwise constant-sum game \mathcal{GG}' resulting from \mathcal{GG} : We preserve the graph structure of \mathcal{GG} , and we assign to every edge (u, v) the payoff matrices

$B^{u,v}$ and $B^{v,u}$ (for the players u and v respectively). Notice that the resulting game is pairwise-constant sum (by Lemma 2), and at the same time separable zero-sum.¹ We show the following lemmas, concluding the proof of Theorem 1.

Lemma 3. *Suppose that there is a pure strategy profile S such that, for every player u , u 's payoff in \mathcal{GG} is the same as his payoff in \mathcal{GG}' under S . If we modify S to \hat{S} by changing a single player's pure strategy, then under \hat{S} every player's payoff in \mathcal{GG}' equals the same player's payoff in \mathcal{GG} .*

Proof. Suppose that, in going from S to \hat{S} , we modify player v 's strategy from i to j . Notice that for all players that are not in v 's neighborhood, their payoffs are not affected by this change. Now take any player u in the neighborhood of v and let u 's strategy be k in both S and \hat{S} . The change in u 's payoff when going from S to \hat{S} in \mathcal{GG} is $A_{k,j}^{u,v} - A_{k,i}^{u,v}$. According to property (a), this equals $B_{k,j}^{u,v} - B_{k,i}^{u,v}$, which is exactly the change in u 's payoff in \mathcal{GG}' . Since the payoff of u is the same in the two games before the update in v 's strategy, the payoff of u remains the same after the change. Hence, all players except v have the same payoffs under \hat{S} in both \mathcal{GG} and \mathcal{GG}' . Since both games have zero total sum of players' payoffs, v should also have the same payoff under \hat{S} in the two games. \square

Lemma 4. *In every pure strategy profile, every player has the same payoff in games \mathcal{GG} and \mathcal{GG}'*

Proof. Suppose that, in going from S to \hat{S} , we modify player v 's strategy from i to j . Notice that for all players that are not in v 's neighborhood, their payoffs are not affected by this change. Now take any player u in the neighborhood of v and let u 's strategy be k in both S and \hat{S} . The change in u 's payoff when going from S to \hat{S} in \mathcal{GG} is $A_{k,j}^{u,v} - A_{k,i}^{u,v}$. According to property (a), this equals $B_{k,j}^{u,v} - B_{k,i}^{u,v}$, which is exactly the change in u 's payoff in \mathcal{GG}' . Since the payoff of u is the same in the two games before the update in v 's strategy, the payoff of u remains the same after the change. Hence, all players except v have the same payoffs under \hat{S} in both \mathcal{GG} and \mathcal{GG}' . Since both games have zero total sum of players' payoffs, v should also have the same payoff under \hat{S} in the two games. \square

¹Indeed, let all players play strategy 1. Since $B_{1,1}^{u,v} = A_{1,1}^{u,v}$, for all u, v , the sum of all players' payoffs in \mathcal{GG}' is the same as the sum of all players' payoffs in \mathcal{GG} , i.e. 0. But \mathcal{GG}' is a constant-sum game. Hence in every other pure strategy profile the total sum of all players' payoffs will also be 0.

3.2 Algorithmic and Equilibrium Properties of Separable Constant-sum Games

First, it is easy to check that the payoff preserving transformation of Theorem 1 also works for transforming separable constant-sum multiplayer games to pairwise constant-sum games. It follows that the two classes of games are payoff preserving transformation equivalent.

Let now \mathcal{GG} be a separable constant-sum multiplayer game, and \mathcal{GG}' be \mathcal{GG} 's payoff-equivalent pairwise constant-sum game, with payoff matrices $B^{u,v}$. Then $B^{u,v} + (B^{v,u})^T = c^{\{u,v\}} \mathbb{1}$ (from Lemma 2). We create a new game, \mathcal{GG}'' , by assigning payoff tables $D^{u,v} = B^{u,v} - \frac{c^{u,v}}{2} \mathbb{1}$ on each edge (u, v) . The new game \mathcal{GG}'' is a pairwise zero-sum game. Moreover, it is easy to see that, under the same strategy profile S , for any player u , the difference between her payoff in games \mathcal{GG} , \mathcal{GG}' and the game \mathcal{GG}'' is a fixed constant. Hence, the three games share the same set of Nash equilibria. From this and the result of [11] Properties (1) and (2) follow.

Now let every node $u \in V$ of the original game \mathcal{GG} choose a mixed strategy $x_u^{(t)}$ at every time step $t = 1, 2, \dots$, and suppose that each player's sequence of strategies $\langle x_u^{(t)} \rangle$ is no-regret against the sequences of the other players.² It is not hard to see that the same no-regret property must also hold in the games \mathcal{GG}' and \mathcal{GG}'' , since for every player u her payoffs in these three games only differ by a fixed constant under any strategy profile. But \mathcal{GG}'' is a pairwise zero-sum game. Hence, we know from [11] that the round-average of the players' mixed strategy sequences are approximate Nash equilibria in \mathcal{GG}'' , with the approximation going to 0 with the number of rounds. But, since for every player u her payoffs in the three games only differ by a fixed constant under any strategy profile, it follows that the round-average of the players' mixed strategy sequences are also approximate Nash equilibria in \mathcal{GG} , with the same approximation guarantee. Property (3) follows. The precise quantitative guarantee of this statement can be found in Lemma 6 of Section 3.4, where we also provide a different, constructive, proof of this statement. The original proof in [11] was non-constructive.

²A reader who is not familiar with the definition of no-regret sequences is referred to Chapter 2.

3.3 A Direct Reduction to Linear Programming

We describe a direct reduction of separable zero-sum games to linear programming, which obviates the use of our payoff-preserving transformation from the previous section. Our reduction can be described in the following terms. Given an n -player zero-sum polymatrix game we construct a 2-player game, called *the lawyer game*. The lawyer game is *not* zero-sum, so we cannot hope to compute its equilibria efficiently. In fact, its equilibria may be completely unrelated to the equilibria of the underlying polymatrix game. Nevertheless, we show that a certain kind of “restricted equilibrium” of the lawyer game *can* be computed with linear programming; moreover, we show that we can map a “restricted equilibrium” of the lawyer game to a Nash equilibrium of the zero-sum polymatrix-game in polynomial time. We proceed to the details of the lawyer-game construction.

Let $\mathcal{GG} := \{A^{u,v}, A^{v,u}\}_{(u,v) \in E}$ be an n -player separable zero-sum multiplayer game, such that every player $u \in [n]$ has m_u strategies, and set $A^{u,v} = A^{v,u} = 0$ for all pairs $(u,v) \notin E$. Given \mathcal{GG} , we define the corresponding lawyer game $\mathcal{G} = (R, C)$ to be a symmetric $\sum_u m_u \times \sum_u m_u$ bimatrix game, whose rows and columns are indexed by pairs $(u : i)$, of players $u \in [n]$ and strategies $i \in [m_u]$. For all $u, v \in [n]$ and $i \in [m_u], j \in [m_v]$, we set

$$R_{(u:i),(v:j)} = A_{i,j}^{u,v} \quad \text{and} \quad C_{(u:i),(v:j)} = A_{j,i}^{v,u}.$$

Intuitively, each lawyer can choose a strategy belonging to any one of the nodes of \mathcal{GG} . If they happen to choose strategies of adjacent nodes, they receive the corresponding payoffs that the nodes would receive in \mathcal{GG} from their joint interaction. For a fixed $u \in V$, we call the strategies $\{(u : i)\}_{i \in [m_u]}$ the *block of strategies corresponding to u* , and proceed to define the concepts of a *legitimate strategy* and a *restricted equilibrium* in the lawyer game.

Definition 4 (Legitimate Strategy). *Let x be a mixed strategy for a player of the lawyer game and let $x_u := \sum_{i \in [m_u]} x_{u:i}$. If $x_u = 1/n$ for all u , we call x a legitimate strategy.*

Definition 5 (Restricted Equilibrium). *Let x, y be legitimate strategies for the row and column players of the lawyer game. If for any legitimate strategies x', y' : $x^T \cdot R \cdot y \geq x'^T \cdot R \cdot y$ and $x^T \cdot C \cdot y \geq x^T \cdot C \cdot y'$, we call (x, y) a restricted equilibrium of the lawyer game.*

Given that the lawyer game is symmetric, it has a symmetric Nash equilibrium [20]. We

observe that it also has a symmetric *restricted* equilibrium; moreover, that these are in one-to-one correspondence with the Nash equilibria of the polymatrix game.

Lemma 5. *If $S = (x^1; \dots; x^n)$ is a Nash equilibrium of \mathcal{GG} , where the mixed strategies x^1, \dots, x^n of nodes $1, \dots, n$ have been concatenated in a big vector, $(\frac{1}{n}S, \frac{1}{n}S)$ is a symmetric restricted equilibrium of \mathcal{G} , and vice versa.*

We now have the ground ready to give our linear programming formulation for computing a symmetric restricted equilibrium of the lawyer game and, by virtue of Lemma 5, a Nash equilibrium of the polymatrix game. Our proposed LP is the following. The variables x and z are $(\sum_u m_u)$ -dimensional, and \hat{z} is n -dimensional. We show how this LP implies tractability and convexity of the Nash equilibria of \mathcal{GG} in Appendix A.1.1 (Lemmas 19 and 20).

$$\begin{aligned} \max \quad & \frac{1}{n} \sum_u \hat{z}_u \\ \text{s.t.} \quad & x^T \cdot R \geq z^T; \\ & z_{u:i} = \hat{z}_u, \quad \forall u, i; \\ & \sum_{i \in [m_u]} x_{u:i} = \frac{1}{n}, \quad \forall u \quad \text{and} \quad x_{u:i} \geq 0, \quad \forall u, i. \end{aligned}$$

Remark 1. It is a priori not clear why the linear program shown above computes a restricted equilibrium of the lawyer game. The intuition behind its formulation is the following: The last line of constraints is just guaranteeing that x is a legitimate strategy. Exploiting the separable zero-sum property we can establish that, when restricted to legitimate strategies, the lawyer game is actually a zero-sum game. I.e., for every pair of legitimate strategies (x, y) , $x^T \cdot R \cdot y + x^T \cdot C \cdot y = 0$ (see Lemma 17 in Appendix A.1.1). Hence, if the row player fixed her strategy to a legitimate x , the best response for the column player would be to minimize $x^T \cdot R \cdot y$. But the minimization is over legitimate strategies y ; so the minimum of $x^T \cdot R \cdot y$ coincides with the maximum of $\frac{1}{n} \sum_u \hat{z}_u$, subject to the first two sets of constraints of the program; this justifies our choice of objective function.

Notice that our program looks similar to the standard program for zero-sum bimatrix games, except for a couple of important differences. First, it is crucial that we only allow legitimate strategies x ; otherwise the lawyer game would not be zero-sum and the hope to solve it efficiently would be slim. Moreover, we average out the payoffs from different blocks

of strategies in the objective function instead of selecting the worst payoff, as is done by the standard program.

3.4 A Constructive Proof of the Convergence of No-Regret Algorithms

As we mentioned in the introduction, an attractive property of 2-player zero-sum games is that a large variety of learning algorithms converge to a Nash equilibrium of the game. In [11], it was shown that pairwise zero-sum polymatrix games inherit this property. In this thesis, we have generalized this result to the class of separable zero-sum multiplayer games by employing the proof of [11] as a black box. Nevertheless, the argument of [11] had an undesired (and surprising) property, in that it was employing Brouwer's fixed point theorem as a non-constructive step. Our argument here is based on first principles and is constructive.

We give a constructive proof of the following.

Lemma 6. *Suppose that every node $u \in V$ of a separable zero-sum multiplayer game \mathcal{GG} plays a no-regret sequence of strategies $\langle x_u^{(t)} \rangle_{t=1,2,\dots}$, with regret $g(T) = o(T)$. Then, for all T , the set of strategies $\bar{x}_u^{(T)} = \frac{1}{T} \sum_{t=1}^T x_u^{(t)}$, $u \in V$, is a $\left(n \cdot \frac{g(T)}{T}\right)$ -approximate Nash equilibrium of \mathcal{GG} .*

Proof. We have the following

$$\begin{aligned} & \sum_{t=1}^T \left(\sum_{(u,v) \in E} x^T \cdot A^{u,v} \cdot x_v^{(t)} \right) \\ &= \sum_{(u,v) \in E} \left(x^T \cdot A^{u,v} \cdot \left(\sum_{t=1}^T x_v^{(t)} \right) \right) \\ &= T \cdot \sum_{(u,v) \in E} x^T \cdot A^{u,v} \cdot \bar{x}_v^{(T)}. \end{aligned}$$

Let z_u be the best response of u , if for all v in u 's neighborhood v plays strategy $\bar{x}_v^{(T)}$. Then for all u , and any mixed strategy x for u , we have,

$$\sum_{(u,v) \in E} z_u^T \cdot A^{u,v} \cdot \bar{x}_v^{(T)} \geq \sum_{(u,v) \in E} x^T \cdot A^{u,v} \cdot \bar{x}_v^{(T)}. \quad (1)$$

Using the No-Regret Property

$$\begin{aligned}
& \sum_{t=1}^T \left(\sum_{(u,v) \in E} (x_u^{(t)})^T \cdot A^{u,v} \cdot x_v^{(t)} \right) \\
& \geq \sum_{t=1}^T \left(\sum_{(u,v) \in E} z_u^T \cdot A^{u,v} \cdot x_v^{(t)} \right) - g(T) \\
& = T \cdot \sum_{(u,v) \in E} z_u^T \cdot A^{u,v} \cdot \bar{x}_v^{(T)} - g(T)
\end{aligned}$$

Let us take a sum over all $u \in V$ on both the left and the right hand sides of the above.

The LHS will be

$$\begin{aligned}
& \sum_{u \in V} \left(\sum_{t=1}^T \left(\sum_{(u,v) \in E} (x_u^{(t)})^T \cdot A^{u,v} \cdot x_v^{(t)} \right) \right) \\
& = \sum_{t=1}^T \left(\sum_{u \in V} \left(\sum_{(u,v) \in E} (x_u^{(t)})^T \cdot A^{u,v} \cdot x_v^{(t)} \right) \right) \\
& = \sum_{t=1}^T \left(\sum_{u \in V} P_u \right) = \sum_{t=1}^T 0 = 0 \quad (\text{by the zero-sum property})
\end{aligned}$$

The RHS is

$$T \cdot \sum_{u \in V} \left(\sum_{(u,v) \in E} z_u^T \cdot A^{u,v} \cdot \bar{x}_v^{(T)} \right) - n \cdot g(T)$$

The LHS is greater than the RHS, thus

$$\begin{aligned}
0 & \geq T \cdot \sum_{u \in V} \left(\sum_{(u,v) \in E} z_u^T \cdot A^{u,v} \cdot \bar{x}_v^{(T)} \right) - n \cdot g(T) \\
\Rightarrow n \cdot \frac{g(T)}{T} & \geq \sum_{u \in V} \left(\sum_{(u,v) \in E} z_u^T \cdot A^{u,v} \cdot \bar{x}_v^{(T)} \right).
\end{aligned}$$

Recall that the game is zero-sum. So if every player u plays $\bar{x}_u^{(T)}$, the sum of players' payoffs is 0. Thus

$$\sum_{u \in V} \left(\sum_{(u,v) \in E} (\bar{x}_u^{(T)})^T \cdot A^{u,v} \cdot \bar{x}_v^{(T)} \right) = 0.$$

Hence:

$$\begin{aligned}
& n \cdot \frac{g(T)}{T} \\
& \geq \sum_{u \in V} \left(\sum_{(u,v) \in E} z_u^T \cdot A^{u,v} \cdot \bar{x}_v^{(T)} \right. \\
& \quad \left. - \sum_{(u,v) \in E} (\bar{x}_u^{(T)})^T \cdot A^{u,v} \cdot \bar{x}_v^{(T)} \right).
\end{aligned}$$

But (1) implies that $\forall u$:

$$\sum_{(u,v) \in E} z_u^T \cdot A^{u,v} \cdot \bar{x}_v^{(T)} - \sum_{(u,v) \in E} (\bar{x}_u^{(T)})^T \cdot A^{u,v} \cdot \bar{x}_v^{(T)} \geq 0.$$

So we have that the sum of positive numbers is bounded by $n \cdot \frac{g(T)}{T}$. Hence $\forall u$,

$$\begin{aligned}
n \cdot \frac{g(T)}{T} & \geq \sum_{(u,v) \in E} z_u^T \cdot A^{u,v} \cdot \bar{x}_v^{(T)} \\
& \quad - \sum_{(u,v) \in E} (\bar{x}_u^{(T)})^T \cdot A^{u,v} \cdot \bar{x}_v^{(T)}.
\end{aligned}$$

So for all u , if all other players v play $\bar{x}_v^{(T)}$, the payoff given by the best response is at most $\left(n \cdot \frac{g(T)}{T}\right)$ better than payoff given by playing $(\bar{x}_u^{(T)})$. Thus, it is a $\left(n \cdot \frac{g(T)}{T}\right)$ -approximate Nash equilibrium for every player u to play $(\bar{x}_u^{(T)})$. \square

Chapter 4

Coordination Polymatrix Games

A pairwise constant-sum polymatrix game models a network of competitors. What if the endpoints of every edge are not competing, but coordinating? We model this situation by coordination polymatrix games. This class of games is useful for modeling the spread of ideas and technologies over social networks [16]. Clearly the modification changes the nature of the polymatrix game. We explore the result of this modification to the computational complexity of the new model.

Two-player coordination games are well-known to be potential games. We observe that coordination polymatrix games are also (cardinal) potential games (Proposition 2). Moreover, a pure Nash equilibrium of a two-player coordination game can be found trivially by inspection. We show instead that in coordination polymatrix games the problem becomes PLS-complete (Theorem 3); our reduction is from the Max-Cut Problem with the flip neighborhood. Because our games are potential games, best response dynamics converge to a pure Nash equilibrium, albeit potentially in exponential time. It is fairly standard to show that, if only ϵ -best response steps are allowed, a pseudo-polynomial time algorithm for approximate pure Nash equilibria is obtained. Combining Theorem 3 with Nash's theorem [20] we obtain Corollary 3.

Proposition 2. *Coordination polymatrix games are cardinal potential games.*

Proof of Proposition 2: Using $u_i(S)$ to denote player i 's payoff in the strategy profile is S , we show that the scaled social welfare function

$$\Phi(S) = \frac{1}{2} \sum_i u_i(S) \tag{4.1}$$

is an exact potential function of the game.

Lemma 7. Φ is an exact potential function of the game.

Proof. Let us fix a pure strategy profile S and consider the deviation of player i from strategy S_i to strategy S'_i . If j_1, j_2, \dots, j_ℓ are i 's neighbors, we have that

$$u_i(S'_i, S_{-i}) - u_i(S) = \sum_k u_{j_k}(S'_i, S_{-i}) - \sum_k u_{j_k}(S),$$

since the game on every edge is a coordination game. On the other hand, the payoffs of all the players who are not in i 's neighborhood remain unchanged. Therefore,

$$\Phi(S'_i, S_{-i}) - \Phi(S) = u_i(S'_i, S_{-i}) - u_i(S).$$

Hence, Φ is an exact potential function of the game. □

□

Now we will argue that finding a pure Nash equilibrium is PLS-complete.

Theorem 3. Finding a pure Nash equilibrium of a coordination polymatrix game is PLS-complete.

Proof. We reduce the *Max-Cut problem* with the flip neighborhood to the problem of computing a pure Nash equilibrium of a coordination polymatrix game. If the graph $G = (V, E)$ in the instance of the Max-Cut problem has n nodes, we construct a polymatrix game on the same graph $G = (V, E)$, such that every node has 2 strategies 0 and 1. For any edge $(u, v) \in E$, the payoff is $w_{u,v}$ if u and v play different strategies, otherwise the payoff is 0.

For any pure strategy profile S , we can construct a cut from S in the natural way by letting the nodes who play strategy 0 comprise one side of the cut, and those who play strategy 1 the other side. Edges that have endpoints in different groups are in the cut and we can show that $\Phi(S)$ equals the size of the cut. Indeed, for any edge (u, v) , if the edge is in the cut, u and v play different strategies, so they both receive payoff $w_{u,v}$ on this edge. So this edge contributes $w_{u,v}$ to $\Phi(S)$. If the edge is not in the cut, u and v receive payoff of 0 on this edge. In this case, the edge contributes 0 to $\Phi(S)$. So the size of the cut equals $\Phi(S)$. But $\Phi(S)$ is an exact potential function of the game, so pure Nash equilibria are in

one-to-one correspondence to the local maxima of Φ under the neighborhood defined by one player (node) flipping his strategy (side of the cut). Therefore, every pure Nash equilibrium is a local Max-Cut under the flip neighborhood. \square

Although finding an exact pure Nash equilibrium is PLS-complete, we showed that best response will lead to an ϵ -approximate Nash equilibrium in reasonable time.

Lemma 8. *Suppose that in every step of the dynamics we only allow a player to change her strategy if she can increase her payoff by at least ϵ . Then in $O(\frac{n \cdot d_{\max} \cdot u_{\max}}{\epsilon})$ steps, we will reach an ϵ -approximate pure Nash equilibrium, where u_{\max} is the magnitude of the maximum in absolute value entry in the payoff tables of the game, and d_{\max} the maximum degree.*

Proof. As showed in Lemma 7, if a player u increases her payoff by ϵ , Φ will also increase by ϵ . Since every player's payoff is at least $-d_{\max} \cdot u_{\max}$, and at most $d_{\max} \cdot u_{\max}$, Φ lies in $[-\frac{1}{2}n \cdot d_{\max} \cdot u_{\max}, \frac{1}{2}n \cdot d_{\max} \cdot u_{\max}]$. Thus, there can be at most $\frac{n \cdot d_{\max} \cdot u_{\max}}{\epsilon}$ updates to the potential function before no player can improve by more than ϵ . \square

Corollary 3. *Finding a Nash equilibrium of a coordination polymatrix game is in $\text{PLS} \cap \text{PPAD}$.*

Corollary 3 may be viewed as an indication that coordination polymatrix games are tractable, as a PPAD- or PLS-completeness result would have quite remarkable complexity theoretic implications. On the other hand, we expect the need of quite novel techniques to tackle this problem. Hence, coordination polymatrix games join an interesting family of fixed point problems that are not known to be in P, while they belong to $\text{PLS} \cap \text{PPAD}$; other important problems in this intersection are Simple Stochastic Games [8] and P -Matrix Linear Complementarity Problems [18].

Chapter 5

Combining Coordination and Zero-sum Games

We saw that, if a polymatrix game is zero-sum, we can compute an equilibrium efficiently. We also showed that, if every edge is a 2-player coordination game, the problem is in $\text{PPAD} \cap \text{PLS}$. Zero-sum and coordination games are the simplest cases of two-player games. This explains the lack of hardness results for the above models. A question often posed to us in response to these results (e.g. in [23]) is whether the combination of zero-sum and coordination games is well-behaved. What is the complexity of a polymatrix game if every edge can either be a zero-sum or a coordination game?

We eliminate the possibility of a positive result by establishing a PPAD-completeness result for this seemingly simple model. A key observation that makes our hardness result plausible is that if we allowed double edges between vertices, we would be able to simulate a general polymatrix game. Indeed, suppose that u and v are neighbors in a general polymatrix game, and the payoff matrices along the edge (u, v) are $C^{u,v}$ and $C^{v,u}$. We can define then a pair of coordination and zero-sum games as follows. The coordination game has payoff matrices $A^{u,v} = (A^{v,u})^T = (C^{u,v} + (C^{v,u})^T)/2$, and the zero-sum game has payoff matrices $B^{u,v} = -(B^{v,u})^T = (C^{u,v} - (C^{v,u})^T)/2$. Hence, $A^{u,v} + B^{u,v} = C^{u,v}$ and $A^{v,u} + B^{v,u} = C^{v,u}$. Given that general polymatrix games are PPAD-complete [10], the above decomposition shows that double edges give rise to PPAD-completeness in our model. We show next that unique edges suffice for PPAD-completeness. In fact, seemingly simple structures comprising of groups of friends who coordinate with each other while

participating in zero-sum edges against opponent groups are also PPAD-complete. These games, called *group-wise zero-sum polymatrix games*, are discussed in Section 5.3.

We proceed to describe our PPAD-completeness reduction from general polymatrix games to our model. The high level idea of our proof is to make a twin of each player, and design some gadgetry that allows us to simulate the double edges described above by unique edges. Our reduction will be equilibrium preserving. In the sequel we denote by \mathcal{G} a general polymatrix game and by \mathcal{G}^* the game output by our reduction. We start with a polymatrix game with 2 strategies per player, and call these strategies 0 and 1. Finding an exact Nash equilibrium in such a game is known to be PPAD-complete.

5.1 Gadgets

To construct the game \mathcal{G}^* , we introduce two gadgets. The first is a *copy gadget*. It is used to enforce that a player and her twin always choose the same mixed strategies. The gadget has three nodes, u_0 , u_1 and u_b , and the nodes u_0 and u_1 play zero-sum games with u_b . The games are designed to make sure that u_0 and u_1 play strategy 0 with the same probability. The payoffs on the edges (u_0, u_b) and (u_1, u_b) are defined as follows (we specify the value of M later):

- u_b 's payoff

– on edge (u_0, u_b) :	<table style="border-collapse: collapse;"> <tr> <td style="border-bottom: 1px solid black; padding: 2px 10px;"></td> <td style="padding: 2px 10px;">$u_0:0$</td> <td style="padding: 2px 10px;">$u_0:1$</td> </tr> <tr> <td style="padding: 2px 10px;">$u_b:0$</td> <td style="padding: 2px 10px;">M</td> <td style="padding: 2px 10px;">0</td> </tr> <tr> <td style="padding: 2px 10px;">$u_b:1$</td> <td style="padding: 2px 10px;">$-2M$</td> <td style="padding: 2px 10px;">$-M$</td> </tr> </table>		$u_0:0$	$u_0:1$	$u_b:0$	M	0	$u_b:1$	$-2M$	$-M$
	$u_0:0$	$u_0:1$								
$u_b:0$	M	0								
$u_b:1$	$-2M$	$-M$								

– on edge (u_1, u_b) :	<table style="border-collapse: collapse;"> <tr> <td style="border-bottom: 1px solid black; padding: 2px 10px;"></td> <td style="padding: 2px 10px;">$u_1:0$</td> <td style="padding: 2px 10px;">$u_1:1$</td> </tr> <tr> <td style="padding: 2px 10px;">$u_b:0$</td> <td style="padding: 2px 10px;">$-2M$</td> <td style="padding: 2px 10px;">$-M$</td> </tr> <tr> <td style="padding: 2px 10px;">$u_b:1$</td> <td style="padding: 2px 10px;">M</td> <td style="padding: 2px 10px;">0</td> </tr> </table>		$u_1:0$	$u_1:1$	$u_b:0$	$-2M$	$-M$	$u_b:1$	M	0
	$u_1:0$	$u_1:1$								
$u_b:0$	$-2M$	$-M$								
$u_b:1$	M	0								

The payoff of u_0 on (u_0, u_b) and of u_1 on (u_1, u_b) are defined by taking respectively the negative transpose of the first and second matrix above so that the games on these edges are zero-sum.

The second gadget is used to *simulate* in \mathcal{G}^* the game played in \mathcal{G} . For an edge (u, v) of \mathcal{G} , let us assume that the payoffs on this edge are the following:

- u 's payoff:

	$v:0$	$v:1$
$u:0$	x_1	x_2
$u:1$	x_3	x_4

- v 's payoff:

	$v:0$	$v:1$
$u:0$	y_1	y_2
$u:1$	y_3	y_4

It's easy to see that for any i , there exists a_i and b_i , such that $a_i + b_i = x_i$ and $a_i - b_i = y_i$. To simulate the game on (u, v) , we use u_0, u_1 to represent the two copies of u , and v_0, v_1 to represent the two copies of v . Coordination games are played on the edges (u_0, v_0) and (u_1, v_1) , while zero-sum games are played on the edges (u_0, v_1) and (u_1, v_0) . We only write down the payoffs for u_0, u_1 . The payoffs of v_0, v_1 are then determined, since we have already specified what edges are coordination and what edges are zero-sum games.

- u_0 's payoff

– on edge (u_0, v_0) :

	$v_0:0$	$v_0:1$
$u_0:0$	a_1	a_2
$u_0:1$	a_3	a_4

– on edge (u_0, v_1) :

	$v_1:0$	$v_1:1$
$u_0:0$	b_1	b_2
$u_0:1$	b_3	b_4

- u_1 's payoff

– on edge (u_1, v_0) :

	$v_0:0$	$v_0:1$
$u_1:0$	b_1	b_2
$u_1:1$	b_3	b_4

– on edge (u_1, v_1) :

	$v_1:0$	$v_1:1$
$u_1:0$	a_1	a_2
$u_1:1$	a_3	a_4

5.2 Construction of \mathcal{G}^*




Figure 5-1: An illustration of the PPAD-completeness reduction. Every edge (u, v) of the original polymatrix game \mathcal{G} corresponds to the structure shown at the bottom of the figure. The dashed edges correspond to coordination games, while the other edges are zero-sum games.

For every node u in \mathcal{G} , we use a copy gadget with u_0, u_1, u_b to represent u in \mathcal{G}^* . And for every edge (u, v) in \mathcal{G} , we build a simulating gadget on u_0, u_1, v_0, v_1 . The resulting game \mathcal{G}^* has either a zero-sum game or a coordination game on every edge, and there is at most one

edge between every pair of nodes. For an illustration of the construction see Figure 5-1. It is easy to see that \mathcal{G}^* can be constructed in polynomial time given \mathcal{G} . We are going to show that given a Nash equilibrium of \mathcal{G}^* , we can find a Nash equilibrium of \mathcal{G} in polynomial time.

5.3 Correctness of the Reduction

For any u_i and any pair v_0, v_1 , the absolute value of the payoff of u_i from the interaction against v_0, v_1 is at most $M_{u,v} := \max_{j,k} (|a_j| + |b_k|)$, where the a_j 's and b_k 's are obtained from the payoff tables of u and v on the edge (u, v) . Let $P = n \cdot \max_u \max_v M_{u,v}$. Then for every u_i , the payoff collected from all players other than u_b is in $[-P, P]$. We choose $M = 3P + 1$. We establish the following (proof in Appendix A.2).

Lemma 9. *In every Nash equilibrium S^* of \mathcal{G}^* , and any copy gadget u_0, u_1, u_b , the players u_0 and u_1 play strategy 0 with the same probability.*

Proof of Lemma 9: We use $\mathcal{P}_u^*(u : i, S_{-u})$ to denote the payoff for u when u plays strategy i , and the other players' strategies are fixed to S_{-u} . We also denote by x the probability with which u_0 plays 0, and by y the corresponding probability of player u_1 . For a contradiction, assume that there is a Nash equilibrium S^* in which $x \neq y$. Then

$$\begin{aligned}
& \mathcal{P}_{u_b}^*(u_b : 0, S_{-u_b}^*) \\
&= M \cdot x + (-2M) \cdot y + (-M) \cdot (1 - y) \\
&= M \cdot (x - y) - M \\
& \mathcal{P}_{u_b}^*(u_b : 1, S_{-u_b}^*) \\
&= (-2M) \cdot x + (-M) \cdot (1 - x) + M \cdot y \\
&= M \cdot (y - x) - M
\end{aligned}$$

Since u_0 and u_1 are symmetric, we assume that $x > y$ WLOG. In particular, $x - y > 0$, which implies $\mathcal{P}_{u_b}^*(u_b : 0, S_{-u_b}^*) > \mathcal{P}_{u_b}^*(u_b : 1, S_{-u_b}^*)$. Hence, u_b plays strategy 0 with probability 1. Given this, if u_0 plays strategy 0, her total payoff should be no greater than $-M + P = -2P - 1$. If u_0 plays 1, the total payoff will be at least $-P$. $-2P - 1 < -P$, thus u_0 should play strategy 1 with probability 1. In other words, $x = 0$. This is a contradiction

to $x > y$. \square

Assume that S^* is a Nash equilibrium of \mathcal{G}^* . According to Lemma 9, any pair of players u_0, u_1 use the same mixed strategy in S^* . Given S^* we construct a strategy profile S for \mathcal{G} by assigning to every node u the common mixed strategy played by u_0 and u_1 in \mathcal{G}^* . For u in \mathcal{G} , we use $\mathcal{P}_u(u : i, S_{-u})$ to denote u 's payoff when u plays strategy i and the other players play S_{-u} . Similarly, for u_j in \mathcal{G}^* , we let $\widehat{\mathcal{P}}_{u_j}^*(u_j : i, S_{-u_j}^*)$ denote the sum of payoffs that u_j collects from all players other than u_b , when u_j plays strategy i , and the other players play $S_{-u_j}^*$. We show the following lemmas (see Appendix A.2), resulting in the proof of Theorem 4.

Lemma 10. *For any Nash equilibrium S^* of \mathcal{G}^* , any pair of players u_0, u_1 of \mathcal{G}^* and the corresponding player u of \mathcal{G} , $\widehat{\mathcal{P}}_{u_0}^*(u_0 : i, S_{-u_0}^*) = \widehat{\mathcal{P}}_{u_1}^*(u_1 : i, S_{-u_1}^*) = \mathcal{P}_u(u : i, S_{-u})$.*

Lemma 11. *If S^* is a Nash equilibrium of \mathcal{G}^* , S is a Nash equilibrium of \mathcal{G} .*

Theorem 4. Finding a Nash equilibrium in polymatrix games with coordination or zero-sum games on their edges is PPAD-complete.

Theorem 4 follows from Lemma 11 and the PPAD-completeness of polymatrix games with 2 strategies per player [10]. In fact, our reduction shows a stronger result. In our reduction, players can be naturally divided into three groups. Group A includes all u_0 nodes, group B includes all u_b nodes and group C all u_1 nodes. It is easy to check that the games played inside the groups A , B and C are only coordination games, while the games played across groups are only zero-sum (recall Figure 5-1). Such games in which the players can be partitioned into groups such that all edges within a group are coordination games and all edges across different groups are zero-sum games are called *group-wise zero-sum polymatrix games*. Intuitively these games should be simpler since competition and coordination are not interleaving with each other. Nevertheless, our reduction shows that group-wise zero-sum polymatrix games are PPAD-complete, even for 3 groups of players, establishing Theorem 5.

Chapter 6

Strictly Competitive Polymatrix Games

Recall that two-player strictly competitive games are merely affine transformation of two-player zero-sum games. Given this result it is quite natural to expect that polymatrix games with strictly competitive games on their edges should be tractable. Strikingly we show that this is not the case.

Theorem 2. Finding a Nash equilibrium in polymatrix games with strictly competitive games on their edges is PPAD-complete.

The proof is based on the PPAD-completeness of polymatrix games with coordination and zero-sum games on their edges. The idea is that we can use strictly competitive games to simulate coordination games. Indeed, suppose that (A, A) is a coordination game between nodes u and v . Using two parallel edges we can simulate this game by assigning game $(2A, -A)$ on one edge and $(-A, 2A)$ on the other. Both games are strictly competitive games, but the aggregate game between u and v is the original coordination game. In our setting, we do not allow parallel edges between nodes. We go around this using our copy gadget from the previous section which only has zero-sum games.

Proof of Theorem 2: We reduce a polymatrix game \mathcal{G} with either coordination or zero-sum games on its edges to a polymatrix game \mathcal{G}^* all of whose edges are strictly competitive games. For every node u , we use a copy gadget (see Section 5) to create a pair of twin nodes u_0, u_1 representing u . By the properties of the copy gadget u_0 and u_1 use the same mixed strategy in all Nash equilibria of \mathcal{G}^* . Moreover, the copy gadget only uses zero-sum

games.

Having done this, the rest of \mathcal{G}^* is defined as follows.

- If the game between u and v in \mathcal{G} is a zero-sum game, it is trivial to simulate it in \mathcal{G}^* . We can simply let both (u_0, v_0) and (u_1, v_1) carry the same game as the one on the edge (u, v) ; clearly the games on (u_0, v_0) and (u_1, v_1) are strictly competitive. An illustration is shown in Figure 6-1.




Figure 6-1: Simulation of a zero-sum edge in \mathcal{G} (shown at the top) by a gadget comprising of only zero-sum games (shown at the bottom).

- If the game between u and v in \mathcal{G} is a coordination game (A, A) , we let the games on the edges (u_0, v_1) and (u_1, v_0) be $(2A, -A)$, and the games on the edges (u_0, v_0) and (u_1, v_1) be $(-A, 2A)$ as shown in Figure 6-2. All the games in the gadget are strictly competitive.




Figure 6-2: Simulation of a coordination edge (A, A) in \mathcal{G} . At the top we have broken (A, A) into two parallel edges. At the bottom we show the gadget in \mathcal{G}^* simulating these edges.

The rest of the proof proceeds by showing the following lemmas that are the exact analogues of the Lemmas 9, 10 and 11 of Section 5.

Lemma 12. *In every Nash equilibrium S^* of \mathcal{G}^* , and any copy gadget u_0, u_1, u_b , the players u_0 and u_1 play strategy 0 with the same probability.*

Assume that S^* is a Nash equilibrium of \mathcal{G}^* . Given S^* we construct a mixed strategy profile S for \mathcal{G} by assigning to every node u the common mixed strategy played by u_0 and u_1 in \mathcal{G}^* . For u in \mathcal{G} , we use $\mathcal{P}_u(u : i, S_{-u})$ to denote u 's payoff when u plays strategy i and the other players play S_{-u} . Similarly, for u_j in \mathcal{G}^* , we let $\widehat{\mathcal{P}}_{u_j}^*(u_j : i, S_{-u_j}^*)$ denote the sum of payoffs that u_j collects from all players other than u_b , when u_j plays strategy i , and the other players play $S_{-u_j}^*$. Then:

Lemma 13. *For any Nash equilibrium S^* of \mathcal{G}^* , any pair of players u_0, u_1 of \mathcal{G}^* and the corresponding player u of \mathcal{G} , $\widehat{\mathcal{P}}_{u_0}^*(u_0 : i, S_{-u_0}^*) = \widehat{\mathcal{P}}_{u_1}^*(u_1 : i, S_{-u_1}^*) = \mathcal{P}_u(u : i, S_{-u})$.*

Lemma 14. *If S^* is a Nash equilibrium of \mathcal{G}^* , S is a Nash equilibrium of \mathcal{G} .*

We omit the proofs of the above lemmas as they are essentially identical to the proofs of Lemmas 9, 10 and 11 of Section 5. By combining Theorem 4 and Lemma 14 we conclude the proof of Theorem 2. \square

Bibliography

- [1] Adler, I.: On the Equivalence of Linear Programming Problems and Zero-Sum Games. In: Optimization Online (2010).
- [2] Adler, I., Daskalakis, C., Papadimitriou, C. H.: A Note on Strictly Competitive Games. In: WINE (2009).
- [3] Aumann, R. J.: Game Theory. In: The New Palgrave: A Dictionary of Economics by J. Eatwell, M. Milgate, and P. Newman (eds.), London: Macmillan & Co, 460–482 (1987).
- [4] Cesa-Bianchi, N., Lugosi, G.: Prediction, learning, and games. Cambridge University Press (2006)
- [5] Bregman, L. M., Fokin, I. N.: Methods of Determining Equilibrium Situations in Zero-Sum Polymatrix Games. Optimizatsia 40(57), 70–82 (1987) (in Russian)
- [6] Bregman, L. M., Fokin, I. N.: On Separable Non-Cooperative Zero-Sum Games. Optimization 44(1), 69–84 (1998)
- [7] Chen, X., Deng, X.: Settling the complexity of 2-player Nash-equilibrium. In: FOCS (2006).
- [8] Condon, A.: The Complexity of Stochastic Games. In: Information and Computation 96(2): 203-224 (1992)
- [9] Dantzig, G. B.: Linear Programming and Extensions. Princeton University Press, (1963)
- [10] Daskalakis, C., Goldberg, P. W., Papadimitriou, C. H.: The Complexity of Computing a Nash Equilibrium. In: STOC (2006)
- [11] Daskalakis, C., Papadimitriou, C. H.: On a Network Generalization of the Minmax Theorem. In: ICALP (2009)
- [12] Daskalakis, C., Tardos, É.: private communication (2009)
- [13] Freund, Y., Schapire, R. E.: Adaptive Game Playing Using Multiplicative Weights. In: Games and Economic Behavior 29, 79–103 (1999)
- [14] Johnson, D. S., Papadimitriou, C. H. and Yannakakis, M.: How Easy is Local Search? In: Journal of Computer and System Sciences, 37(1):79100, (1988)

- [15] Kearns, M. J., Littman, M. L., Singh, S. P.: Graphical Models for Game Theory. In: UAI (2001)
- [16] Kempe, D., Kleinberg, J., Tardos, É.: Maximizing the Spread of Influence through a Social Network. In SIGKDD (2003)
- [17] Khachiyan, L. G.: A Polynomial Algorithm in Linear Programming. Soviet Mathematics Doklady, 20(1): 191-194, (1979)
- [18] Megiddo, N.: private communication (2009)
- [19] Megiddo, N. and Papadimitriou, C.H. : On Total Functions, Existence Theorems and Computational Complexity. In: Theoretical Computer Science, 81(2):317-324, (1991)
- [20] Nash, J.: Noncooperative games. Ann. Math., 54:289–295 (1951)
- [21] Papadimitriou, C. H.: On the Complexity of the Parity Argument and Other Inefficient Proofs of Existence. In: J. Comput. Syst. Sci. 48(3): 498-532 (1994)
- [22] von Neumann, J.: Zur Theorie der Gesellschaftsspiele. Math. Annalen 100, 295–320 (1928)
- [23] Workshop on Research Issues at the Interface of Computer Science and Economics, by Blume, L., Easley, D., Kleinberg, J., Tardos, E., Kalai, E. (organizers). Cornell University, September 2009

Appendix A

Omitted Proofs

A.1 Separable Zero-Sum Multiplayer Games

A.1.1 LP Formulation

Proof of Lemma 5: We show the following lemmas.

Lemma 15. *Every Nash equilibrium of the separable zero-sum multiplayer game \mathcal{GG} can be mapped to a symmetric restricted equilibrium of the lawyer game \mathcal{G} .*

Proof of Lemma 15: Let S be a Nash equilibrium of \mathcal{GG} . Denote by $S_u(i)$ the probability that u places on strategy $i \in [m_u]$ and S_u the mixed strategy of u . We construct a legitimate strategy \bar{x} by setting $\bar{x}_{u,i} = S_u(i)/n$. We claim that (\bar{x}, \bar{x}) is a symmetric restricted equilibrium. Indeed let us fix the row player's strategy to \bar{x} . For every block of the column player's strategies indexed by u , it is optimal for the column player to distribute the $1/n$ available probability mass for this block proportionally to S_u . This is because S_u is a best response for player u to the mixed strategies of the other players. \square

Lemma 16. *From any symmetric restricted equilibrium of the lawyer game \mathcal{G} , we can recover a Nash equilibrium of \mathcal{GG} in polynomial time.*

Proof of Lemma 16: Let (x, x) be a symmetric restricted equilibrium of the lawyer game. We let

$$\hat{x}_u(i) = n \cdot x_{u,i}$$

and we denote by S the strategy profile in \mathcal{GG} where every player u plays strategy $i \in [m_u]$ with probability $\hat{x}_u(i)$. We show that S is a Nash equilibrium of \mathcal{GG} .

We prove this by contradiction. If S is not a Nash equilibrium, there exists a player u who can increase her payoff by deviating from strategy S_u to some strategy S'_u . Let us then define a new legitimate strategy x' for the row player of the lawyer game. x' is the same as x , except that $x_{u:i} = S'_u(i)/n$, for all $i \in [m_u]$. It is easy to see that

$$x'^T \cdot R \cdot x - x^T \cdot R \cdot x = \frac{1}{n^2}(\mathcal{P}_u(S') - \mathcal{P}_u(S)) > 0$$

Therefore, (x, x) is not a restricted equilibrium of the lawyer game, a contradiction. \square

Combining the above we conclude the proof of Lemma 5. \square

Lemma 17 (Restricted Zero-Sum Property). *If x and y are respectively legitimate strategies for the row and column players of \mathcal{G} ,*

$$x^T \cdot R \cdot y + x^T \cdot C \cdot y = 0.$$

Proof of Lemma 17: We start with the following.

Lemma 18. *Let u be a node of \mathcal{GG} and v_1, v_2, \dots, v_k be u 's neighbors. Let y_u represent a mixed strategy for u and x_{v_i} mixed strategies for v_i , $i = 1, \dots, k$. For any fixed collection $\{x_{v_i}\}_{i=1}^k$, as we range y_u ,*

$$\sum_i x_{v_i}^T \cdot A^{v_i, u} \cdot y_u + \sum_i y_u^T \cdot A^{u, v_i} \cdot x_{v_i} \text{ remains constant.}$$

Proof of Lemma 18: Assume that the $x_{v_i}, i = 1, \dots, k$, are held fixed. As we change y_u the only payoffs that are affected are those on the edges incident to u . The sum of these payoffs is

$$\sum_i x_{v_i}^T \cdot A^{v_i, u} \cdot y_u + \sum_i y_u^T \cdot A^{u, v_i} \cdot x_{v_i}$$

Since the sum of all payoffs in the game should be 0 and the payoffs on all the other edges do not change, it must be that, as y_u varies, the quantity

$$\sum_i x_{v_i}^T \cdot A^{v_i, u} \cdot y_u + \sum_i y_u^T \cdot A^{u, v_i} \cdot x_{v_i}$$

remains constant. \square

We use Lemma 18 to establish the (restricted) zero-sum property of the lawyer game \mathcal{G} . To do this, we employ a hybrid argument. Before proceeding let us introduce some notation: If z is a legitimate strategy, then for any node $w \in \mathcal{GG}$ we let $\bar{z}_w := (z_{w:1}, z_{w:2}, \dots, z_{w:m_w})^T$.

Let y' be a legitimate strategy, such that $y'_{v:i} = y_{v:i}$ for all $v \neq u$ and $i \in [m_v]$. Assume that v_1, v_2, \dots, v_k are u 's neighbors. Then

$$\begin{aligned}
& (x^T \cdot R \cdot y + x^T \cdot C \cdot y) \\
& \quad - (x^T \cdot R \cdot y' + x^T \cdot C \cdot y') \\
&= \left(\sum_i \bar{x}_{v_i}^T \cdot A^{v_i, u} \cdot \bar{y}_u + \sum_i \bar{x}_{v_i}^T \cdot (A^{u, v_i})^T \cdot \bar{y}_u \right) \\
& \quad - \left(\sum_i \bar{x}_{v_i}^T \cdot A^{v_i, u} \cdot \bar{y}'_u + \sum_i \bar{x}_{v_i}^T \cdot (A^{u, v_i})^T \cdot \bar{y}'_u \right) \\
&= \left(\sum_i \bar{x}_{v_i}^T \cdot A^{v_i, u} \cdot \bar{y}_u + \sum_i \bar{y}'_u^T \cdot A^{u, v_i} \cdot \bar{x}_{v_i} \right) \\
& \quad - \left(\sum_i \bar{x}_{v_i}^T \cdot A^{v_i, u} \cdot \bar{y}'_u + \sum_i \bar{y}'_u^T \cdot A^{u, v_i} \cdot \bar{x}_{v_i} \right) \\
&= 0 \quad (\text{making use of Lemma 18})
\end{aligned}$$

We established that if we change strategy y on a single block u , the sum of the lawyers' payoffs remains unaltered. By doing this n times, we can change y to x without changing the sum of lawyers' payoffs. On the other hand, we know that $x^T \cdot R \cdot x$ is $1/n^2$ times the sum of all nodes' payoffs in \mathcal{GG} , if every node u plays $n \cdot \bar{x}_u$. We know that \mathcal{GG} is zero-sum and that $R = C^T$. It follows that $x^T \cdot R \cdot x = x^T \cdot C \cdot x = 0$. We conclude that

$$x^T \cdot R \cdot y + x^T \cdot C \cdot y = x^T \cdot R \cdot x + x^T \cdot C \cdot x = 0.$$

□

We conclude with a proof that a Nash equilibrium in \mathcal{GG} can be computed efficiently, and that the set of Nash equilibria is convex. This is done in two steps as follows.

Lemma 19. *Using our LP formulation we can compute a symmetric restricted equilibrium of the lawyer game \mathcal{G} in polynomial time. Moreover, the set of symmetric restricted equilibria of \mathcal{G} is convex.*

Proof of Lemma 19: We argue that a solution of the linear program will give us a symmetric restricted equilibrium of \mathcal{G} . By Nash's theorem [20], \mathcal{GG} has a Nash equilibrium S . Using S define \bar{x} as in the proof of Lemma 15. Since (\bar{x}, \bar{x}) is a restricted equilibrium of the lawyer game, $\bar{x}^T \cdot C \cdot y \leq \bar{x}^T \cdot C \cdot \bar{x} = 0$, for any legitimate strategy y for the column player.¹ Using Lemma 17 we obtain then that $\bar{x}^T \cdot R \cdot y \geq 0$, for all legitimate y . So if we hold $x := \bar{x}$ fixed in the linear program, and optimize over z, \hat{z} we would get value ≥ 0 . So the LP value is ≥ 0 . Hence, if (x', z, \hat{z}) is an optimal solution to the LP, it must be that $\frac{1}{n} \sum_u \hat{z}_u \geq 0$, which means that for any legitimate strategy y , $x'^T \cdot R \cdot y \geq 0$. Therefore, $x'^T \cdot C \cdot y \leq 0$ for any legitimate y , using Lemma 17 again. So if the row player plays x' , the payoff of the column player is at most 0 from any legitimate strategy. On the other hand, if we set $y = x'$, $x'^T \cdot C \cdot x' = 0$. Thus, x' is a (legitimate strategy) best response for the column player to the strategy x' of the row player. Since \mathcal{G} is symmetric, x' is also a (legitimate strategy) best response for the row player to the strategy x' of the column player. Thus, (x', x') is a symmetric restricted equilibrium of the lawyer game.

We show next that the optimal value of the LP is 0. Indeed, we already argued that the LP value is ≥ 0 . Let then (x', z, \hat{z}) be an optimal solution to the LP. Since x' is a legitimate strategy for \mathcal{G} , we know that $x'^T \cdot R \cdot x' = 0$ (see our argument in the proof of Lemma 17). It follows that if we hold $x = x'$ fixed in the LP and try to optimize the objective over the choices of z, \hat{z} we would get objective value $\leq x'^T \cdot R \cdot x' = 0$. But x' is an optimal choice for x . Hence the optimal value of the LP is ≤ 0 . Combining the above we get that the LP value is 0.

We showed above that if (x', z, \hat{z}) is an optimal solution of the LP, then (x', x') is a restricted equilibrium of \mathcal{G} . We show next the opposite direction, i.e. that if (x', x') is a restricted equilibrium of \mathcal{G} then (x', z, \hat{z}) is an optimal solution of the LP for some z, \hat{z} . Indeed, we argued above that for any restricted equilibrium (x', x') , $x'^T \cdot R \cdot y \geq 0$, for every legitimate strategy y . Hence, holding $x = x'$ fixed in the LP, and optimizing over z, \hat{z} , the objective value is at least 0 for the optimal choice of $z = z(x'), \hat{z} = \hat{z}(x')$. But the LP-value is 0. Hence, $(x', z(x'), \hat{z}(x'))$ is an optimal solution. But the set of optimal solutions of the LP is convex. Hence, the set $\{x' \mid \exists z, \hat{z} \text{ such that } (x', z, \hat{z}) \text{ is an optimal solution of the LP}\}$ is also convex. Hence, the set $\{(x', x') \mid \exists z, \hat{z} \text{ such that } (x', z, \hat{z}) \text{ is an optimal solution of}$

¹In the proof of Lemma 17 we show that, for any legitimate strategy x in the lawyer game, $x^T \cdot R \cdot x = x^T \cdot C \cdot x = 0$.

the LP} is also convex. But this set, as we argued above, is precisely the set of symmetric restricted equilibria of \mathcal{G} . \square

Lemma 20. *For any separable zero-sum multiplayer game \mathcal{GG} , we can compute a Nash equilibrium in polynomial time using linear programming, and the set of Nash equilibria of \mathcal{GG} is convex.*

Proof of Lemma 20: Given \mathcal{GG} , we can construct the corresponding lawyer game \mathcal{G} efficiently. By Lemma 19, we can compute a symmetric restricted equilibrium of \mathcal{G} in polynomial time, and using the mapping in Lemma 16, we can recover a Nash equilibrium of \mathcal{GG} in polynomial time. Moreover, from the proof of Lemma 19 it follows that the set $\{x' \mid (x', x') \text{ is a symmetric restricted equilibrium of } \mathcal{G}\}$ is convex. Hence, the set $\{nx' \mid (x', x') \text{ is a symmetric restricted equilibrium of } \mathcal{G}\}$ is also convex. But the latter set is by Lemma 5 precisely the set of Nash equilibria of \mathcal{GG} . \square

A.2 Polymatrix Games with Coordination and Zero-Sum Edges

Proof of Lemma 10: We first show

$$\widehat{\mathcal{P}}_{u_0}^*(u_0 : i, S_{-u_0}^*) = \widehat{\mathcal{P}}_{u_1}^*(u_1 : i, S_{-u_1}^*).$$

Since \mathcal{G}^* is a polymatrix game, it suffices to show that the sum of payoffs that u_0 collects from v_0, v_1 is the same with the payoff that u_1 collects. Since S^* is a Nash equilibrium, according to Lemma 9, we can assume that v_0 and v_1 play strategy 0 with the same probability q .

We use $u^*(u_i : j, v_0, v_1)$ to denote u_i 's payoff when playing j .

$$\begin{aligned}
u^*(u_0 : 0, v_0, v_1) &= a_1 \cdot q + a_2 \cdot (1 - q) + b_1 \cdot q + b_2 \cdot (1 - q) \\
&= (a_1 + b_1) \cdot q + (a_2 + b_2) \cdot (1 - q) \\
u^*(u_0 : 1, v_0, v_1) &= a_3 \cdot q + a_4 \cdot (1 - q) + b_3 \cdot q + b_4 \cdot (1 - q) \\
&= (a_3 + b_3) \cdot q + (a_4 + b_4) \cdot (1 - q) \\
u^*(u_1 : 0, v_0, v_1) &= a_1 \cdot q + a_2 \cdot (1 - q) + b_1 \cdot q + b_2 \cdot (1 - q) \\
&= (a_1 + b_1) \cdot q + (a_2 + b_2) \cdot (1 - q) \\
u^*(u_1 : 1, v_0, v_1) &= a_3 \cdot q + a_4 \cdot (1 - q) + b_3 \cdot q + b_4 \cdot (1 - q) \\
&= (a_3 + b_3) \cdot q + (a_4 + b_4) \cdot (1 - q)
\end{aligned}$$

So $u^*(u_0 : i, v_0, v_1) = u^*(u_1 : i, v_0, v_1)$. Thus, $\widehat{\mathcal{P}}_{u_0}^*(u_0 : i, S_{-u_0}^*) = \widehat{\mathcal{P}}_{u_1}^*(u_1 : i, S_{-u_1}^*)$.

Next we show

$$\widehat{\mathcal{P}}_{u_0}^*(u_0 : i, S_{-u_0}^*) = \mathcal{P}_u(u : i, S_{-u})$$

Since \mathcal{G} is also a polymatrix game, we can just show that the payoff that u collects from v is the same as the payoff that u_0 collects from v_0 and v_1 . By the construction of S , v plays strategy 0 with probability q . Letting $u(u : i, v)$ be the payoff for u , if u plays strategy i , we have

$$\begin{aligned}
u(u : 0, v) &= (a_1 + b_1) \cdot q + (a_2 + b_2) \cdot (1 - q) \\
u(u : 1, v) &= (a_3 + b_3) \cdot q + (a_4 + b_4) \cdot (1 - q)
\end{aligned}$$

So $u^*(u_0 : i, v_0, v_1) = u(u : i, v)$. Therefore, $\widehat{\mathcal{P}}_{u_0}^*(u_0 : i, S_{-u_0}^*) = \mathcal{P}_u(u : i, S_{-u})$. \square

Proof of Lemma 11: We only need to show that, for any player u in \mathcal{G} , playing the same

strategy that u_0, u_1 use in \mathcal{G}^* is indeed a best response for u . According to Lemma 10,

$$\widehat{\mathcal{P}}_{u_0}^*(u_0 : i, S_{-u_0}^*) = \widehat{\mathcal{P}}_{u_1}^*(u_1 : i, S_{-u_1}^*) = \mathcal{P}_u(u : i, S_{-u}).$$

Let

$$\begin{aligned} P_i &:= \widehat{\mathcal{P}}_{u_0}^*(u_0 : i, S_{-u_0}^*) \\ &= \widehat{\mathcal{P}}_{u_1}^*(u_1 : i, S_{-u_1}^*) = \mathcal{P}_u(u : i, S_{-u}). \end{aligned}$$

Also let r be the probability that u_b assigns to strategy 0 and let $u^*(u_i : j)$ be the payoff of u_i along the edge (u_i, u_b) when playing strategy j .

$$\begin{aligned} u^*(u_0 : 0) &= -M \cdot r + 2M \cdot (1 - r) = 2M - 3M \cdot r \\ u^*(u_0 : 1) &= M \cdot (1 - r) = M - M \cdot r \\ u^*(u_1 : 0) &= 2M \cdot r + (-M) \cdot (1 - r) = 3M \cdot r - M \\ u^*(u_1 : 1) &= M \cdot r \end{aligned}$$

Let p be the probability with which u_0, u_1, u play strategy 0. Since S^* is a Nash equilibrium of \mathcal{G}^* , if $p \in (0, 1)$, then we should have the following equalities:

$$\begin{aligned} u^*(u_0 : 0) + P_0 &= 2M - 3M \cdot r + P_0 \\ = u^*(u_0 : 1) + P_1 &= M - M \cdot r + P_1 \end{aligned} \quad (1)$$

$$\begin{aligned} u^*(u_1 : 0) + P_0 &= 3M \cdot r - M + P_0 \\ = u^*(u_1 : 1) + P_1 &= M \cdot r + P_1 \end{aligned} \quad (2)$$

Then

$$\begin{aligned} &2M - 3M \cdot r + P_0 + 3M \cdot r - M + P_0 \\ &= M - M \cdot r + P_1 + M \cdot r + P_1 \end{aligned}$$

$$\Rightarrow M + 2P_0 = M + 2P_1$$

$$\Rightarrow P_0 = P_1.$$

Therefore, it is a best response for u to play strategy 0 with probability p . We can show the same for the extremal case that u_0, u_1 play pure strategies ($p = 0$ or $p = 1$).

Therefore, for any u , S_u is a best response to the other players' strategy S_{-u} . So S is a Nash equilibrium of \mathcal{G} . \square