A Multiplayer Generalization of the MinMax Theorem

Yang Cai* Ozan Candogan† Constantinos Daskalakis‡ Christos Papadimitriou§

Abstract

We show that in zero-sum *polymatrix* games, a multiplayer generalization of two-person zero-sum games, Nash equilibria can be found efficiently with linear programming. We also show that the set of coarse correlated equilibria collapses to the set of Nash equilibria. In contrast, other positive properties of two-person zero-sum games are not preserved: Nash equilibrium payoffs need not be unique, and Nash equilibrium strategies need not be exchangeable or max-min.

1 Introduction

According to Robert Aumann [Aum87], two-person zero-sum games[1] are *"one of the few areas in game theory, and indeed in the social sciences, where a fairly sharp, unique prediction is made."* Indeed, in a two-person zero-sum game, max-min strategies offer a rather compelling solution: They constitute a Nash equilibrium, and this Nash equilibrium is unique modulo degeneracy. Furthermore, these mixed strategies can be easily computed with linear programming. In contrast, we now know that Nash equilibria are hard to compute in general, even for two-person non-zero-sum games [DGP06, CDT06]—and consequently for three-person zero-sum games. Von Neumann's minmax theorem [Neu28] seems to have very narrow applicability.

In this note we prove a multi-player generalization of the minmax theorem. We show that for any multi-player polymatrix game that is zero-sum the Nash equilibrium can be easily found by linear programming (and in fact by a quite direct generalization of the linear programming formulation of two-person zero-sum games). Informally, a *polymatrix game* (or *separable network game*) is defined by a graph. The nodes of the graph are the players, and the edges of the graph are two-person games. Every node has a fixed set of strategies, and chooses a strategy from this set to play in all games corresponding to adjacent edges. Given a strategy profile of all the players, the node's payoff is the sum of its payoffs in all games on edges adjacent to it. The game is *zero-sum* if, for all strategy profiles, the payoffs of all players add up to zero. This is the class of games we are considering; we present a simple method, based on linear programming, for finding a Nash equilibrium in such games (Theorem 2).

Zero-sum polymatrix games can model common situations in which nodes in a network interact pairwise and make decisions (for example, adopt one of many technologies, or choose one or more of their neighbors for preferential interaction), and which is a closed system of payoffs, in that it

*EECS, MIT; supported by NSF Award CCF-0953960 (CAREER) and CCF-1101491; `ycai@mit.edu`.

†Fuqua School of Business, Duke University; `ozan.candogan@duke.edu`.

‡EECS, MIT; supported by a Sloan Foundation fellowship, a Microsoft Reseach faculty fellowship and NSF award CCF-0953960 (CAREER) and CCF-1101491; `costis@mit.edu`.

§EECS, UC Berkeley; supported by NSF award CCF-0964033 and a Google university research award; `christos@cs.berkeley.edu`.

[1]Actually, Aumann makes the statement for (two-person) *strictly competitive* games; but these were recently shown to essentially coincide with two-person zero-sum games [ADP09].

is impossible for payoffs to flow in or out of the system. It is an intriguing class of games: since the definition involves a universal quantification over all strategy profiles, an exponential set, it is not a priori clear that there is an efficient algorithm for recognizing such games (but there is, see Section 4).

One immediate way to obtain such a game is to create a polymatrix game in which all edge-games are zero-sum, see [BF87, BF98, DP09]. But there are other ways:

Example 1. *Consider the* security game *between many evaders and many inspectors (these are the players) with many entry points (these are the strategies); each entry point is within the jurisdiction of one inspector. The network is a complete bipartite graph between evaders and inspectors. Each evader can choose any entry point, and each inspector can choose one entry point in her jurisdiction. If an evader's point of entry is inspected by some inspector, that inspector wins one unit of payoff. If the evader's point of choice is not inspected, the evader wins one unit. All other payoffs are zero. This simple polymatrix game is not zero-sum, but it is constant-sum: It is easy to see that, for any strategy profile, the total payoff equals the number of evaders. Thus it can be turned into zero-sum by, say, subtracting this amount from the payoffs of any player. But the resulting zero-sum polymatrix game has constituent games which are not zero- or constant- sum. In other words, the zero-sum nature of this game is a* global, *rather than a local, property. See Section 4 for further discussion of this point.*

2 The Main Result

We first define zero-sum polymatrix games formally.

Definition 1. *A* polymatrix game, *or* separable network game $\mathcal{G}$ *consists of the following:*

- *a finite set $V = \{1, \ldots, n\}$ of* players, *sometimes called* nodes, *and a finite set E of* edges, *which are taken to be unordered pairs $[i, j]$ of players, $i \neq j$;*
- *for each player $i \in V$, a finite set of strategies S_i;*
- *for each edge $[i, j] \in E$, a two-person game (p^{ij}, p^{ji}) where the players are i, j, the strategy sets S_i, S_j, respectively, and the payoffs $p^{ij} : S_i \times S_j \mapsto \Re$, and similarly for p^{ji};*
- *for each player $i \in V$ and strategy profile $s = (s_1, \ldots, s_n) \in \prod_{j \in V} S_j$, the payoff of player i under s is $p_i(s) = \sum_{[i,j] \in E} p^{ij}(s_i, s_j)$.*

Furthermore, $\mathcal{G}$ is zero-sum *if for all strategy profiles $s = (s_1, \ldots, s_n) \in \prod_{j \in V} S_j$, $\sum_{i \in V} p_i(s) = 0$.*

Fix a zero-sum polymatrix game $\mathcal{G}$. We shall next formulate a linear program which captures $\mathcal{G}$. The variables are the mixed strategies of the players, so we have a probability x_i^s for all $i \in V$ and $s \in S_i$. We denote the vector of all these variables by x so that x encodes a mixed strategy profile. We require that $x_i^s \geq 0$, for all i and s, and $\sum_{s \in S_i} x_i^s = 1$, for all i, writing $x \in \Delta$, if x satisfies these constraints.

For $x \in \Delta$, we write x_i for the mixed strategy of player i and $x_{-i} \in \Delta_{-i}$ for the vector of mixed strategies of all players but player i, where Δ_{-i} denotes the set of all possible x_{-i}'s. We sometimes write $(x_i,\ x_{-i})$ for x, and $(s,\ x_{-i})$ for the mixed strategy profile x such that $x_i^s = 1$, i.e. x_i corresponds to the pure strategy $s \in S_i$. Moreover, we extend $p_i(\cdot)$ to mixed strategies by taking expectations. Namely,

$$p_i(x) \equiv \sum_{[i,j] \in E} \sum_{s_i \in S_i, s_j \in S_j} p^{ij}(s_i, s_j) x_i^{s_i} x_j^{s_j}$$

represents the expected payoff of player i under mixed strategy profile x. Similarly, for $s \in S_i$, $p_j(s, x_{-i})$ represents the expected payoff of player j when player i uses pure strategy s and the other players use their mixed strategies in x_{-i}.

For each player i, player i's payoff $p_i(s, x_{-i})$ from a pure strategy s in S_i is obviously a linear function of x_{-i}. Consider the following linear program in the variables $y \in \Delta$ and $w \triangleq (w_1, ..., w_n)$:

$$\text{LP 1}: \quad \begin{array}{lll} \min\limits_{y,w} & \sum\limits_{i \in V} w_i & \\ \text{subject to} & w_i \geq p_i(s, y_{-i}), & \text{for all } i \in V, s \in S_i, \\ & y \in \Delta. & \end{array}$$

We next state our main result:

Theorem 2. *If (y, w) is an optimal solution to LP 1, then y is a Nash equilibrium of $\mathcal{G}$. Conversely, if y is a Nash equilibrium of $\mathcal{G}$, then there is a w such that (y, w) is an optimal solution to LP 1.*

We give two proofs of this theorem; the first relies on Nash's theorem, whereas, the second only employs linear programming duality.

Proof using Nash's Theorem

The constraints of LP 1 imply that at any feasible solution (y, w) we have $w_i \geq \max_{s \in S_i} p_i(s, y_{-i})$. Moreover, since $p_i(x_i, y_{-i})$ is linear in x_i, it follows that

$$\sum_{i \in V} \max_{s \in S_i} p_i(s, y_{-i}) = \max_{x \in \Delta} \sum_{i \in V} p_i(x_i, y_{-i}). \tag{1}$$

Note that zero-sum property implies that $\sum_{i \in V} p_i(y) = 0$ for any $y \in \Delta$. Using this observation together with (1) and the constraint $w_i \geq \max_{s \in S_i} p_i(s, y_{-i})$ implies that at feasible solution (y, w) we have

$$\sum_{i \in V} w_i \geq \sum_{i \in V} \max_{s \in S_i} p_i(s, y_{-i}) = \max_{x \in \Delta} \sum_{i} p_i(x_i, y_{-i}) \geq \sum_{i \in V} p_i(y_i, y_{-i}) = 0. \tag{2}$$

Hence the optimal objective of the linear program is lower bounded by zero. Nash's theorem implies that there exists a Nash equilibrium y^* such that

$$\max_{s \in S_i} p_i(s, y^*_{-i}) - p_i(y^*_i, y^*_{-i}) = 0. \tag{3}$$

Setting $w_i = \max_{s \in S_i} p_i(s, y^*_{-i})$ it can be seen that (y^*, w) is a feasible solution to LP 1, and (3) implies that all inequalities in (2) hold with equality for this solution, and the objective value zero is achieved in LP 1 so that (y^*, w) is an optimal solution. For the forward direction, consider any optimal solution (w, y) of LP 1. Since the objective value is $\sum_{i \in V} w_i = 0$ in this solution, it follows from (2) that (3) holds for this optimal solution, and hence y is a Nash equilibrium. $\square$

Proof using linear programming duality

In the above proof we use Nash's theorem to conclude that the optimal objective value of LP 1 is equal to zero. It would be surprising if the power of Nash's theorem were necessary to establish a property of a linear program. We show that it is not.

Let us rewrite the constraints of LP 1 with the help of a square matrix R with $\sum_i |S_i|$ rows and columns. The rows and columns of R are indexed by pairs $(i : s)$ and $(j : r)$ for $i, j \in V$ and

$s \in S_i$, $r \in S_j$, and $R_{(i:s),(j:r)} = p^{ij}(s,r)$ if $[i,j] \in E$, $R_{(i:s),(j:r)} = 0$ otherwise. Then $p_i(s, y_{-i})$ in LP 1 corresponds to the row $(Ry)_{(i:s)}$ of Ry indexed by $(i:s)$ for $i \in V$, $s \in S_i$, and $\sum_{i\in V} p_i(x_i, y_{-i}) = x^T R y = y^T R^T x$ for $x, y \in \Delta$. This observation suggests that LP 1 can be reformulated by replacing the constraint $w_i \geq p_i(s, y_{-i})$ with $w_i \geq (Ry)_{(i:s)}$. Thus, the dual of LP 1 (referred to as DLP 1) can be stated using the decision variables $z \in \Delta$ and $v \triangleq (v_1, ..., v_n)$ as follows:

$$\text{DLP 1}: \qquad \begin{array}{lll} \max\limits_{z,v} & \sum\limits_{j\in V} v_j & \\ \text{subject to} & v_j \leq (R^T z)_{(j:r)}, & \text{for all } j \in V, r \in S_j, \\ & z \in \Delta. & \end{array}$$

Similar to LP 1, it can be seen that a feasible solution (z, v) of DLP 1 satisfies

$$\sum_{j\in V} v_j \leq \sum_{j\in V} \min_{r\in S_j} (R^T z)_{(j,r)} = \min_{x\in\Delta} x^T R^T z \leq z^T R^T z = 0, \tag{4}$$

where the first equality follows from the linearity of $x^T R^T z$ in x, and the last one follows from the zero-sum property. So the optimal objective value of DLP 1 is bounded above by zero. Through strong duality this implies that the optimal objective value of LP 1 is bounded above by zero. Since the optimal value of LP 1 is also lower bounded by zero, it follows that LP 1 has value zero, which is what we needed to avoid the use of Nash's theorem in our previous proof of Theorem 2. $\square$

Remark: Interestingly, if (z, v) is an optimal solution to DLP 1, then z is also a Nash equilibrium. This can be seen by noting that by strong duality the optimal objective value of the dual is equal to zero, and hence (4) implies that $\sum_{j\in V} \min_{r\in S_j} (R^T z)_{(j,r)} = z^T R^T z = z^T R z = 0$ at this solution. Hence, z_j assigns positive probability only to entries of $(Rz)_{(j:r)}$ that are minimal. The definition of R implies that for any r this entry is given by $\sum_{i\in V, s\in S_i, [i,j]\in E} z_i^s p^{ij}(s,r)$, i.e., the sum of payoffs of neighbors of player j by playing with her. Since the game is zero-sum, minimizing this quantity maximizes the payoff of player j, and hence z_j is her best response to z_{-j}.

3 Properties of zero-sum polymatrix games

Thus in zero-sum polymatrix games a Nash equilibrium can be found by linear programming, just as in zero-sum two-person games. One immediate question that comes to mind is, which of the many other strong properties of zero-sum two-person games also generalize to zero-sum polymatrix games? We consider the following properties of zero-sum two-person games:

(i) Each player has a *unique payoff value* in all Nash equilibria, known as his value in the game.

(ii) Equilibrium strategies are *max-min strategies*, i.e., each player uses a strategy that maximizes her worst-case payoff (with respect to her opponent's strategies).

(iii) Equilibrium strategies are *exchangeable*, i.e., if (x_1, x_2) and (y_1, y_2) are equilibria, then so are (x_1, y_2) and (y_1, x_2).

(iv) There are no correlated equilibria (or even coarse correlated equilibria, see definition below) whose marginals with respect to the players do not constitute a Nash equilibrium.

As we shall see next, *only one* of these four properties (in particular, (iv)) generalizes to zero-sum polymatrix games.

Value of a node. Does every player in a zero-sum polymatrix game have a value, attained at all equilibria? Consider three players a, b, c. Player a has a single strategy H, whereas players b, c have two strategies H, T (for "heads" and "tails"). The polymatrix game involves two edges: an edge between players a and b, and another edge between b and c. The payoffs are as follows:

[a,b]: If player a chooses the same strategy as player b, player a receives 1 and player b receives -1, otherwise player a receives -1 and player b receives 1.

[b,c]: If player b chooses the same strategy as player c, player b receives 1 and player c receives -1, otherwise player b receives -1 and player c receives 1.

It is straightforward to check that this game is a zero-sum polymatrix game, and the following two strategy profiles are Nash equilibria:

(i) $(H, 1/2(H) + 1/2(T), H)$, i.e., player b uniformly mixes between its strategies, while players a, c choose H. The payoffs of the three players are $(0, 0, 0)$.

(ii) (H, T, H), i.e., player b chooses T, while players b, c choose H; payoffs now are $(-1, 0, 1)$.

Hence, different equilibria assign different payoffs to players in zero-sum polymatrix games.

Max-min strategies. For games with more than two players, a max-min strategy of a player is a strategy that maximizes his worst-case payoff, for any strategies of his opponents. In the game of the previous paragraph, the max-min strategy of player c is given by $1/2(H) + 1/2(T)$. However, we saw that there are Nash equilibria in which c uses a different mixed strategy. Indeed, it is not hard to check that there are no Nash equilibria in which c uses his max-min strategy.

Exchangeability. Exchangeability can be naturally generalized to multi-player games (with a set of players $V = \{1, \ldots, n\}$) as follows: If $\{x_i\}_{i \in V}$ and $\{y_i\}_{i \in V}$ are Nash equilibria, then so is the strategy profile $(x_1, \ldots, x_{i-1}, y_i, x_{i+1}, \ldots x_n)$. To disprove this property for zero-sum polymatrix games, let us consider a game with three players a, b, c, each with two strategies H, T, and three edges: $[a, b]$, $[b, c]$, and $[a, c]$. The payoffs in the games associated with these edges are the same as the payoffs of the matching-pennies game (see Figure 1). We assume that the row players associated

	H	T
H	1,-1	-1,1
T	-1,1	1,-1

Figure 1: Payoffs in a matching-pennies game.

with edges $[a, b]$, $[b, c]$, and $[a, c]$ are respectively a, b, and c. It can be seen that this is a zero-sum polymatrix game, and two Nash equilibria of this game are (i) (H, H, H), and (ii) (T, T, T). On the other hand, (T, H, H) is not an equilibrium strategy, since the third player receives a payoff of -2 at this strategy profile, but he can improve her payoff to 2 by deviating to T.

Correlated equilibria. Recall the definition of correlated equilibrium, and the more general concept of coarse correlated equilibrium:

Definition 3. *Let $\mathcal{S} = \Pi_{i \in V} S_i$ and $z \in \Delta(\mathcal{S})$ be a distribution over pure strategy profiles, where $z^{(s)}$ denotes the probability of pure strategy profile $s \in \mathcal{S}$. z is a* correlated equilibrium *iff for every player i and strategies $r, t \in S_i$,*

$$\sum_{s_{-i} \in \mathcal{S}_{-i}} p_i(r, s_{-i}) \cdot z^{(r,s_{-i})} \geq \sum_{s_{-i} \in \mathcal{S}_{-i}} p_i(t, s_{-i}) \cdot z^{(r,s_{-i})}. \tag{5}$$

z is a coarse correlated equilibrium *iff for every player i and strategy $t \in S_i$,*

$$\sum_{s \in \mathcal{S}} p_i(s) \cdot z^{(s)} \geq \sum_{s_{-i} \in \mathcal{S}_{-i}} p_i(t, s_{-i}) \cdot z_{-i}^{(s_{-i})}, \tag{6}$$

where $z_{-i}^{(s_{-i})} = \sum_{r \in S_i} z^{(r,s_{-i})}$ is the marginal probability that the pure strategy profile sampled by z for players $V \setminus \{i\}$ is s_{-i}.[2]

Theorem 4. *If z is a coarse correlated equilibrium then $\hat{x}$ is a Nash equilibrium, where, for every player i, $\hat{x}_i$ is the marginal probability distribution: $\hat{x}_i^r = \sum_{s_{-i} \in \mathcal{S}_{-i}} z^{(r,s_{-i})}$, for all $r \in S_i$.*

Proof. Since the game is polymatrix, $p_i(r, \hat{x}_{-i}) = \sum_{s_{-i} \in \mathcal{S}_{-i}} p_i(r, s_{-i}) \cdot z_{-i}^{(s_{-i})}$ for all i and $r \in S_i$. Indeed, the LHS is player i's expected payoff from strategy r when the other players use mixed strategies $\hat{x}_{-i}$, while the RHS is i's expected payoff from strategy r when the other players' strategies are jointly sampled from $z_{-i}^{(\cdot)}$. The equality follows from the fact that $\hat{x}_{-i}$ and $z_{-i}^{(\cdot)}$ have the same marginal distributions with respect to the strategy of each player in $V \setminus \{i\}$, and i's payoff only depends on these marginals. Now, let $w_i^* = \sum_{s \in \mathcal{S}} p_i(s) \cdot z^{(s)}$. Because z is a coarse correlated equilibrium, $w_i^* \geq p_i(r, \hat{x}_{-i})$ for any $r \in S_i$. On the other hand, $\sum_i w_i^* = 0$ since the game is zero-sum. These imply that $(\hat{x}, w^*)$ is an optimal solution to LP1, so that $\hat{x}$ is a Nash equilibrium by Theorem 2. □

This result has an interesting algorithmic consequence, which complements Theorem 2. The Nash equilibrium of a zero-sum polymatrix game $\mathcal{G}$ can be found not only with linear programming, but can also be arrived at in a distributed manner, as long as the players run an arbitrary *no-regret learning algorithm* [CBL06, FS99] to update their strategies in a repeated game with stage game $\mathcal{G}$. The players' average strategies can be shown to converge to a Nash equilibrium of $\mathcal{G}$ [CD11].

4 A Transformation

A special case of zero-sum polymatrix games are the *pairwise constant-sum* polymatrix games in which every edge is a two-person constant sum game, and all these constants add up to zero. Superficially, zero-sum polymatrix games appear to be more general. In this section we prove that they are not, by presenting a payoff-preserving transformation from any zero-sum polymatrix game to a pairwise constant-sum polymatrix game.

Payoff Preserving Transformation: We transform a zero-sum polymatrix game $\mathcal{G}$ to a *pairwise constant-sum polymatrix game* $\mathcal{G}'$ by modifying the payoff functions on the edges. For every edge $[i, j]$, we construct a new two player game $(\hat{p}^{ij}, \hat{p}^{ji})$ based on (p^{ij}, p^{ji}). For simplicity, we use 1 to denote the first strategy in every player's strategy set. The new payoffs are defined as follows:

$$\hat{p}^{ij}(r, s) = p^{ij}(1, 1) + \left(p^{ij}(1, s) - p^{ij}(1, 1)\right) + \left(p^{ji}(s, 1) - p^{ji}(s, r)\right). \tag{7}$$

[2]Observe that (6) follows by summing (5) over $r \in S_i$. Hence, if z is a correlated equilibrium, then z is also a coarse correlated equilibrium.

Notice that $\hat{p}^{ij}(1,1) = p^{ij}(1,1)$.

Before we argue that $(\hat{p}^{ij}, \hat{p}^{ji})$ is a constant sum game, we need to prove some useful local properties of (p^{ij}, p^{ji}).

Lemma 5. *For any edge $[i,j]$ and any $r \in S_i$, $s \in S_j$,*

$$p^{ij}(1,1) + p^{ji}(1,1) + p^{ij}(r,s) + p^{ji}(s,r) = p^{ij}(1,s) + p^{ji}(s,1) + p^{ij}(r,1) + p^{ji}(1,r).$$

Proof. Let all players except i and j fix their strategies, and $-\alpha$ represent the sum of all players' payoffs from edges that do not involve i or j as one of their endpoints. Let $P_{(w:k)}$ (w in $\{i,j\}$) be the sum of payoffs of w and her neighbors from all edges incident to w except $[i,j]$ when w plays strategy k. Since the game is zero-sum, the following are true:

- i plays strategy 1, j plays strategy s: $P_{(i:1)} + P_{(j:s)} + p^{ij}(1,s) + p^{ji}(s,1) = \alpha \quad (1)$
- i plays strategy r, j plays strategy 1: $P_{(i:r)} + P_{(j:1)} + p^{ij}(r,1) + p^{ji}(1,r) = \alpha \quad (2)$
- i plays strategy 1, j plays strategy 1: $P_{(i:1)} + P_{(j:1)} + p^{ij}(1,1) + p^{ji}(1,1) = \alpha \quad (3)$
- i plays strategy r, j plays strategy s: $P_{(i:r)} + P_{(j:s)} + p^{ij}(r,s) + p^{ji}(s,r) = \alpha \quad (4)$

Clearly, we have (1) + (2) = (3) + (4). By canceling out the common terms, we obtain the desired equality. □

Now we are ready to prove that $\mathcal{G}'$ is a pairwise constant sum game.

Lemma 6. *For every edge $[i,j]$, for all $r \in S_i, s \in S_j$, $\hat{p}^{ij}(r,s) + \hat{p}^{ji}(s,r) = C_{ij}$, where C_{ij} is an absolute constant that does not depend on r, s.*

Proof. By definition, $\hat{p}^{ji}(s,r) = p^{ji}(1,1) + \big(p^{ji}(1,r) - p^{ji}(1,1)\big) + \big(p^{ij}(r,1) - p^{ij}(r,s)\big)$.

By Lemma 5, we know

$$\big(p^{ij}(1,s) - p^{ij}(1,1)\big) + \big(p^{ji}(s,1) - p^{ji}(s,r)\big) = \big(p^{ji}(1,1) - p^{ji}(1,r)\big) + \big(p^{ij}(r,s) - p^{ij}(r,1)\big).$$

Therefore, we can rewrite $\hat{p}^{ij}(r,s)$ as

$$\hat{p}^{ij}(r,s) = p^{ij}(1,1) + \big(p^{ji}(1,1) - p^{ji}(1,r)\big) + \big(p^{ij}(r,s) - p^{ij}(r,1)\big). \tag{8}$$

Hence, $\hat{p}^{ij}(r,s) + \hat{p}^{ji}(s,r) = p^{ij}(1,1) + p^{ji}(1,1) = C_{ij}$. □

Finally, we prove that the transformation preserves the payoff of every player.

Theorem 7. *For every pure strategy profile, every player has the same payoff in games $\mathcal{G}$ and $\mathcal{G}'$.*

Proof. Since $\hat{p}^{ij}(1,1) = p^{ij}(1,1)$, if every player uses strategy 1, then every player's payoff is the same in both games.

Next we prove that, if every player has the same payoff in both games for a pure strategy profile s, then this still holds if we change some player j's strategy in s from, say, q to t. Indeed, for players that are not in j's neighborhood, their payoffs are not affected by this change in both games. Second, consider some player i in the neighborhood of j and suppose i's strategy is r in both (q, s_{-j}) and (t, s_{-j}). Then the change in i's payoff in $\mathcal{G}$ between these two strategy profiles is $p^{ij}(r,t) - p^{ij}(r,q)$. Equation (8) suggests that this equals $\hat{p}^{ij}(r,t) - \hat{p}^{ij}(r,q)$, which is exactly the change in i's payoff in $\mathcal{G}'$ between these two strategy profiles. Hence, because i's payoffs were the same in $\mathcal{G}$ and $\mathcal{G}'$ before the change in j's strategy, they should still be the same after the change.

Thus, we have argued that all players except j have the same payoff in $\mathcal{G}$ and $\mathcal{G}'$ in the new strategy profile we are considering. But both games are zero-sum, which means that j must also have the same payoff in $\mathcal{G}$ and $\mathcal{G}'$.

Using the above, we can show that every player has the same payoff in $\mathcal{G}$ and $\mathcal{G}'$ for any pure strategy profile. $\square$

An algorithm for recognizing zero-sum polymatrix games

Our main result in this paper states that Nash equilibria in zero-sum polymatrix games can be computed through linear programming, just like in two-person zero-sum games. However, it is not a priori clear that, if such a game is presented to us, we can recognize it, since the definition of a zero-sum polymatrix game involves a universal quantification over all pure strategy profiles (which scale exponentially with the number of players). We could of course apply the transformation described above and check if, in the resulting game, all edge games have constant payoffs, and if these payoffs sum to zero. But even this would not suffice, because the transformation is guaranteed to be payoff-preserving only if the original game was zero-sum.

An efficient algorithm does exist: For every player i, $s \in S_i$, and $x_{-i} \in \Delta_{-i}$, define $P_i(s, x_{-i}) = \sum_{j \in V} p_j(s, x_{-i})$ (namely the sum of all players' payoffs when player i uses pure strategy s, while players in $V \setminus \{i\}$ use mixed strategies x_{-i}). A polymatrix game is constant-sum if and only if changing a single player's strategy in a strategy profile does not affect the total sum of all players' payoffs (hence we can gradually transition between every pair of pure strategy profiles without observing any change in the total sum of players' payoffs). Thus, to check if a game is constant-sum, it suffices to verify that, for all $i \in V$, $r, s \in S_i$, and $x_{-i} \in \Delta_{-i}$, $P_i(r, x_{-i}) = P_i(s, x_{-i})$.

It follows from the above discussion that a polymatrix game is constant-sum if and only if the optimal value of the following linear program is 0 for all $i \in V$, $r, s \in S_i$.

$$\max_{x_{-i} \in \Delta_{-i}} \quad P_i(r, x_{-i}) - P_i(s, x_{-i}).^{3}$$

Given that any such linear program is solvable in polynomial time, we can efficiently determine if a polymatrix game is constant-sum (and therefore if it is zero-sum) by checking whether the optimal value of the above linear program is zero for all $i \in V$, $r, s \in S_i$.

5 Discussion

Our main result is a generalization of von Neumann's minmax theorem from two-person zero-sum games to zero-sum polymatrix games. We also showed that several other properties of two-person zero-sum games fail to generalize to polymatrix games, with one exception: Coarse correlated equilibria collapse to Nash equilibria, and no-regret play converges to Nash equilibrium.

How extensive is the class of zero-sum polymatrix games? We noted in the introduction that it trivially includes all polymatrix games with zero-sum edges, but also other games, such as the security game, for which the zero-sum property seems to be of a "global" nature. However, the results of the last section imply that any zero-sum polymatrix game can be transformed, through a nontrivial transformation, into a payoff-equivalent polymatrix game with constant-sum edges. Whether there are further generalizations of the minmax theorem to more general classes of games is an important open problem.

[3]Notice that the objective function is a linear function of x_{-i} due to the polymatrix nature of the game, and the fact that all payoffs on edges not adjacent to player i cancel out when we take the difference.

References

[ADP09] Ilan Adler, Constantinos Daskalakis, and Christos H. Papadimitriou. A Note on Strictly Competitive Games. In *the 5th Workshop on Internet and Network Economics (WINE)*, 2009.

[Aum87] Robert J. Aumann. Game Theory. *The New Palgrave: A Dictionary of Economics by J. Eatwell, M. Milgate, and P. Newman (eds.), London: Macmillan Co*, pages 460–482, 1987.

[BF87] L. M. Bregman and I. N. Fokin. Methods of Determining Equilibrium Situations in Zero-Sum Polymatrix Games. *Optimizatsia (in Russian)*, 40(57):70–82, 1987.

[BF98] L. M. Bregman and I. N. Fokin. On Separable Non-Cooperative Zero-Sum Games. *Optimization*, 44(1):69–84, 1998.

[CBL06] Nicolo Cesa-Bianchi and Gabor Lugosi. *Prediction, Learning, and Games*. Cambridge University Press, 2006.

[CD11] Yang Cai and Constantinos Daskalakis. A Minmax Theorem for Multiplayer Zero-Sum Games. In *the 22nd ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2011.

[CDT06] Xi Chen, Xiaotie Deng, and Shang-Hua Teng. Computing Nash Equilibria: Approximation and Smoothed Complexity. In *the 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 2006.

[DGP06] Constantinos Daskalakis, Paul W. Goldberg, and Christos H. Papadimitriou. The Complexity of Computing a Nash Equilibrium. In *the 38th Annual ACM Symposium on Theory of Computing (STOC)*, 2006.

[DP09] Constantinos Daskalakis and Christos H. Papadimitriou. On a Network Generalization of the Minmax Theorem. In *the 36th International Colloquium on Automata, Languages and Programming (ICALP)*, 2009.

[FS99] Yoav Freund and Robert E. Schapire. Adaptive Game Playing Using Multiplicative Weights. *Games and Economic Behavior*, 29:79–103, 1999.

[Neu28] John von Neumann. Zur Theorie der Gesellshaftsspiele. *Mathematische Annalen*, 100:295–320, 1928.