

MIT Open Access Articles

Zero-Sum Polymatrix Games: A Generalization of Minmax

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Cai, Yang; Candogan, Ozan; Daskalakis, Constantinos, et al. "Zero-Sum Polymatrix Games: A Generalization of Minmax." *Mathematics of Operations Research* 41, 2 (May 2016): 648–655. © 2016 Institute for Operations Research and the Management Sciences (INFORMS)

As Published: <http://dx.doi.org/10.1287/moor.2015.0745>

Publisher: Institute for Operations Research and the Management Sciences (INFORMS)

Persistent URL: <http://hdl.handle.net/1721.1/110830>

Version: Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

Terms of use: Creative Commons Attribution-Noncommercial-Share Alike



Zero-sum Polymatrix Games: A Generalization of Minmax

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 important 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 games¹ 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

*School of Computer Science, McGill University; Work done while the author was a student at MIT, supported by NSF Award CCF-0953960 (CAREER) and CCF-1101491; cai@cs.mcgill.edu.

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

‡EECS, MIT; supported by a Sloan Foundation fellowship, a Microsoft Research 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.

¹Actually, 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].

of their neighbors for preferential interaction), and which is a closed system of payoffs, in that it 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 several evaders and several inspectors (these are the players) with many exit points (these are the strategies); each exit point is within the jurisdiction of one inspector. The network is a complete bipartite graph between evaders and inspectors. Each evader can choose any exit point, and each inspector can choose one exit point in her jurisdiction. For every evader whose exit point is inspected, the corresponding inspector wins one unit of payoff. If the evader's exit point 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 $V = \{1, \dots, n\}$ 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 $p^{ij} : S_i \times S_j \mapsto \mathbb{R}$, and similarly for p^{ji} ;*
- *for each player $i \in V$ and strategy profile $\bar{s} = (s_1, \dots, s_n) \in \prod_{j \in V} S_j$, the payoff of player i under \bar{s} is $p_i(\bar{s}) = \sum_{[i, j] \in E} p^{ij}(s_i, s_j)$.*

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

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,

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

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 $w := (w_1, \dots, w_n)$:

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

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

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

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

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

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

$$\max_{s \in S_i} p_i(s, y_{-i}^*) - p_i(y_i^*, y_{-i}^*) = 0. \quad (3)$$

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 $\sum_{i \in V} p_i(x_i, y_{-i}) = x^T Ry = y^T R^T x$ 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 $v := (v_1, \dots, v_n)$ as follows:

$$\begin{aligned} \text{DLP 1 :} \quad & \max_{z, v} \quad \sum_{j \in V} v_j \\ & \text{subject to} \quad v_j \leq (R^T z)_{(j:r)}, \quad \text{for all } j \in V, r \in S_j, \\ & \quad \quad \quad z \in \Delta. \end{aligned}$$

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

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

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 that play against 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 her 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) . In particular, the set of equilibrium strategies of each player is convex, and the set of equilibria is the corresponding product set.
- (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 (namely, (iv)) generalizes to zero-sum polymatrix games. Moreover, Property (iii) is partially true; namely the set of equilibrium strategies of each player is convex, but it is not necessarily true anymore that the set of equilibria is the corresponding product set.

Value of a Player. 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 with different player payoffs:

- (i) (H, T, H) , i.e., player b chooses T , while players a, c choose H . The payoffs of the players are $(-1, 0, 1)$. To see that this is an equilibrium, note first that in our game player a only has a single strategy (H). Hence, for trivial reasons she cannot deviate to improve her payoff. Player c maximizes her payoff by choosing a strategy different from the one chosen by b , so she has no incentive to deviate from her strategy in this profile. Finally, given that a and c use strategy H , the payoff of b is equal to zero from both strategies, so she is best responding by playing T . Hence, (H, T, H) is an equilibrium.
- (ii) $(H, 1/2(H) + 1/2(T), H)$, i.e., player b uniformly mixes between her strategies, while players a, c choose H . The payoffs of the players are now $(0, 0, 0)$. Seeing that this profile is an equilibrium is as straightforward as in (i).

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 her worst-case payoff, for any strategies of her 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. Moreover, there are no Nash equilibria in which c uses her max-min strategy. To see this, note that when c uses the aforementioned strategy (and because a only has a single strategy, H) player b maximizes her payoff by using strategy T . On the other hand, if player b uses this strategy, player c can improve her payoff by deviating from her max-min strategy to H .

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

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

Figure 1: Payoffs in a matching-pennies game.

players associated 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 in this strategy profile, but she can improve her payoff to 2 by deviating to T . Note that this example also shows that the set of mixed strategy profiles that are equilibria cannot be expressed as a product of the sets of strategies that players use at equilibrium.²

Correlated equilibria. Recall the definition of correlated equilibrium, and the more general concept of coarse correlated equilibrium (see, e.g., [MV78, CBL06]):

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

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

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

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

where $z_{-i}^{(\bar{s}_{-i})} = \sum_{r \in S_i} z^{(r, \bar{s}_{-i})}$ is the marginal probability that the pure strategy profile sampled by z for players $V \setminus \{i\}$ is \bar{s}_{-i} .³

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: $\hat{x}_i^r = \sum_{\bar{s}_{-i} \in \mathcal{S}_{-i}} z^{(r, \bar{s}_{-i})}$, for all $r \in S_i$.

Proof. Since the game is polymatrix, $p_i(r, \hat{x}_{-i}) = \sum_{\bar{s}_{-i} \in \mathcal{S}_{-i}} p_i(r, \bar{s}_{-i}) \cdot z_{-i}^{(\bar{s}_{-i})}$ 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

²Since the set of optimal solutions of linear programs is convex, Theorem 2 implies that set of mixed strategy profiles that are equilibria is convex. However, lack of exchangeability implies that this convex set is not a product set of strategies of different players.

³Observe 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.

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 $w_i^* = \sum_{\bar{s} \in \mathcal{S}} p_i(\bar{s}) \cdot z^{(\bar{s})}$. 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. \square

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* $\hat{\mathcal{G}}$ 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:

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

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 $r \in S_i, s \in S_j$, we have*

$$p^{ij}(1, 1) + p^{ji}(1, 1) + p^{ij}(r, s) + p^{ji}(s, r) = p^{ij}(1, s) + p^{ji}(s, 1) + p^{ij}(r, 1) + p^{ji}(1, r). \quad (8)$$

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 $P_{(k:t)}$ (k in $\{i, j\}$) be the sum of payoffs of k and her neighbors from all edges incident to k except $[i, j]$, when k plays strategy t . Since the game is zero-sum, the following are true:

- i plays strategy 1, j plays strategy s : $P_{(i:1)} + P_{(j:s)} + p^{ij}(1, s) + p^{ji}(s, 1) = \alpha$ (a)
- i plays strategy r , j plays strategy 1: $P_{(i:r)} + P_{(j:1)} + p^{ij}(r, 1) + p^{ji}(1, r) = \alpha$ (b)
- i plays strategy 1, j plays strategy 1: $P_{(i:1)} + P_{(j:1)} + p^{ij}(1, 1) + p^{ji}(1, 1) = \alpha$ (c)
- i plays strategy r , j plays strategy s : $P_{(i:r)} + P_{(j:s)} + p^{ij}(r, s) + p^{ji}(s, r) = \alpha$ (d)

Clearly, we have $(a) + (b) = (c) + (d)$. By canceling out the common terms, we obtain the desired equality. \square

Next, we show that when \mathcal{G} is zero-sum, $\hat{\mathcal{G}}$ 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 $C_{ij} := p^{ij}(1, 1) + p^{ji}(1, 1)$ is an absolute constant that does not depend on r, s .

Proof. From the definition of $\hat{p}^{ij}(r, s)$ (see (7)) we have

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

Using Lemma 5, this equality can be alternatively expressed as

$$\begin{aligned} \hat{p}^{ij}(r, s) + \hat{p}^{ji}(s, r) &= p^{ij}(1, 1) + p^{ji}(1, 1) + p^{ij}(r, s) + p^{ji}(s, r) - p^{ji}(s, r) - p^{ij}(r, s) \\ &= p^{ij}(1, 1) + p^{ji}(1, 1). \end{aligned}$$

The result follows from the definition of C_{ij} . \square

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 $\hat{\mathcal{G}}$.

Proof. To prove this claim, we first use Lemma 6 to establish that $\hat{\mathcal{G}}$ is zero-sum. Consider any strategy profile \bar{s} in $\hat{\mathcal{G}}$. Observe that

$$\begin{aligned} \sum_{i \in V} \hat{p}_i(\bar{s}) &= \sum_{i \in V} \sum_{[i, j] \in E} \hat{p}^{ij}(s_i, s_j) \\ &= \sum_{[i, j] \in E} \hat{p}^{ij}(s_i, s_j) + \hat{p}^{ji}(s_j, s_i) \\ &= \sum_{[i, j] \in E} p^{ij}(1, 1) + p^{ji}(1, 1) \\ &= 0, \end{aligned}$$

where the third equality follows from Lemma 6, and the last one follows from the fact that the quantity $\sum_{[i, j] \in E} p^{ij}(1, 1) + p^{ji}(1, 1)$ equals the sum of all players' payoffs in zero-sum game \mathcal{G} when all players use their first strategy.

We next complete the proof by using the following consequence of Lemma 5 and (7):

$$\begin{aligned} \hat{p}^{ij}(r, s) &= p^{ij}(1, s) + p^{ji}(s, 1) - p^{ji}(s, r) \\ &= p^{ij}(1, 1) + p^{ji}(1, 1) + p^{ij}(r, s) - p^{ji}(1, r) - p^{ij}(r, 1). \end{aligned} \tag{9}$$

Note that this alternative representation of $\hat{p}^{ij}(r, s)$ immediately implies that for any pair of strategies $s, t \in S_j$:

$$\hat{p}^{ij}(r, s) - \hat{p}^{ij}(r, t) = p^{ij}(r, s) - p^{ij}(r, t). \tag{10}$$

Now consider an arbitrary pure strategy profile, and suppose that some player j changes her strategy from s to t . If some other player i is not a neighbor of j , then i 's payoff does not change as a result of j 's change of strategy. If player i is a neighbor of j , and plays r , then due to j 's change of strategy, the change in i 's payoff in \mathcal{G} equals $p^{ij}(r, s) - p^{ij}(r, t)$, while in $\hat{\mathcal{G}}$ it equals $\hat{p}^{ij}(r, s) - \hat{p}^{ij}(r, t)$. Thus, (10) implies that the payoff change experienced by neighbors of j is identical in both games. Hence, if every player has the same payoff in \mathcal{G} and $\hat{\mathcal{G}}$ before j changes her strategy, every player other than j still has the same payoff in the two games after j changes her strategy. Finally, since both games are zero-sum, j 's payoff is also the same in the two games after the change.

Consider the strategy profile where all players play 1. Since $\hat{p}^{ij}(1, 1) = p^{ij}(1, 1)$ (from (7)), it follows that every player has the same payoff in \mathcal{G} and $\hat{\mathcal{G}}$ in this profile. Given a target strategy profile, start from the all 1's strategy profile and successively change the strategy of every player to match the target profile (one player at a time). By the argument above, after every change, every player has the same payoff in \mathcal{G} and $\hat{\mathcal{G}}$. Thus, the claim follows. \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).

Note that Lemma 5 provides a necessary condition that zero-sum polymatrix games satisfy. Similarly, Lemma 6 indicates that the transformation of zero-sum polymatrix games provided by (7) will have a pairwise constant-sum structure that is easy to check. However, these conditions are not sufficient, i.e., not all games that satisfy the condition of Lemma 5, or whose transformations have the pairwise constant-sum structure, are zero-sum polymatrix games. For instance, consider the game given on the left in Figure 2. This game can be viewed as a polymatrix game with two players connected by an edge. Observe that the payoffs satisfy the condition provided in Lemma 5, but the game is not zero-sum. Similarly, the transformation of this game (given on the right in Figure 2) is constant-sum (in fact zero-sum), despite the fact that the original game is not zero-sum.

	A	B
A	0,0	-3,0
B	1,2	0,0

	A	B
A	0,0	-3,3
B	-2,2	-3,3

Figure 2: Consider a two-player game, where players have two strategies A, B (left), and the corresponding transformed game (right).

We next show that, even though the aforementioned simple conditions are not applicable, there exists an efficient algorithm that can be used to recognize zero-sum polymatrix games.

Theorem 8. *Let \mathcal{G} be a polymatrix game. For any player i , $s \in S_i$, and $x_{-i} \in \Delta_{-i}$, denote by $W(s, x_{-i}) := \sum_{j \in V} p_j(s, x_{-i})$ the sum of all players' payoffs when i plays s and all other players play x_{-i} . \mathcal{G} is a constant-sum game if and only if the optimal objective value of*

$$\max_{x_{-i} \in \Delta_{-i}} W(r, x_{-i}) - W(s, x_{-i}) \quad (11)$$

equals zero for all $i \in V$ and $r, s \in S_i$. Moreover, this condition can be checked in polynomial time (in the number of strategies and players).

Proof. A polymatrix game is constant-sum if and only if changing a single player's strategy in a strategy profile does not affect the sum of all players' payoffs. Equivalently, a game is constant-sum if and only if for all $i \in V$, $r, s \in S_i$, and $x_{-i} \in \Delta_{-i}$, $W(r, x_{-i}) = W(s, x_{-i})$. Observe that if the latter condition holds, then the optimal objective value of (11) equals zero, for all $i \in V$ and $r, s \in S_i$. Moreover, if the optimal objective value of (11) equals zero for all $i \in V$ and $r, s \in S_i$, then $W(r, x_{-i}) = W(s, x_{-i})$ for all $i \in V$ and $r, s \in S_i$. Thus, \mathcal{G} is a zero-sum polymatrix game if and only if (11) has objective value zero, for all $i \in V$ and $r, s \in S_i$.

Notice that the objective function of (11) is a linear function of x_{-i} since \mathcal{G} is a polymatrix game, and all payoffs on edges not adjacent to player i cancel out when we take the difference. Moreover, the constraint $x_{-i} \in \Delta_{-i}$ is linear, since Δ_{-i} is a product space of simplices. Thus, (11) is a linear optimization problem, and by solving (11) for all $i \in V$ and $r, s \in S_i$, it is possible to check in polynomial time whether \mathcal{G} is a constant-sum polymatrix game. \square

Our theorem implies that by solving the optimization problem in (11) for all $i \in V$ and $r, s \in S_i$, it is possible to check if a polymatrix game is constant-sum. Moreover, if the game is constant-sum, by evaluating W at an arbitrary strategy profile it is possible to check if the game is zero-sum.

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. Another interesting direction involves understanding the classes of games that are strategically equivalent (e.g., in the sense formalized in [MV78]) to zero-sum polymatrix games.

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.
- [MV78] H. Moulin and J.-P. Vial. Strategically zero-sum games: The class of games whose completely mixed equilibria cannot be improved upon. *International Journal of Game Theory*, 7(3-4):201–221, 1978.
- [Neu28] John von Neumann. Zur Theorie der Gesellschaftsspiele. *Mathematische Annalen*, 100:295–320, 1928.