Lecture #15 Notes Summary

Polymatrix games, star-shaped Graphical games, Stackelberg games

---

Polymatrix Games

Polymatrix games is a particular class of $N-$player games, in which the payoff of every player is obtained as a sum of individual payoffs gained against each other player. We start to define the class of Graphical Polymatrix Games (GPM games). Here, the players are associated with the nodes of a graph, and interact only with their neighbours. For a graph $G = (V, E)$ and a vertex $v \in V$, we denote by $N(v)$ the neighbourhood of v, that is, the set of vertices adjacent to v.

Definition 1 (GPM game)**.** Let $G = (V, E)$ be a simple (undirected) graph with N nodes. For $i = 1, \ldots, N$, let $n_i \geq 1$ be the number of (pure) strategies of player i. Let a pair of matrices $A_{ij} \in \mathbb{R}^{n_i \times n_j}$ and $A_{ji} \in \mathbb{R}^{n_j \times n_i}$ be given for all $e \equiv (i, j) \in E$. The graphical polymatrix game $\mathcal{G} = (G, (A_{ij}, A_{ji})_{(i,j) \in E})$ is the $N-$player game with the following payoff function:

$$\forall i \in [N], \quad \pi_i : \ \Delta_{n_1} \times \ldots \times \Delta_{n_N} \mapsto \mathbb{R}, \qquad \pi_i(\boldsymbol{p_1}, \ldots, \boldsymbol{p_N}) = \sum_{j \in N(i)} \boldsymbol{p_i}^T A_{ij}\, \boldsymbol{p_j}$$

Definition 2 (pairwise zero-sum GPM game)**.** A GMP game is called *pairwise zero-sum* if for all $e \equiv (i, j) \in E$,

$$A_{ij} = -A_{ji}^T.$$

Nash equilibriums are defined as usual, i.e., $(\boldsymbol{p_1}, \ldots, \boldsymbol{p_N})$ is a NE iff for all $i \in [N]$, $\boldsymbol{p_i} \in BR_i(\boldsymbol{p_{-i}})$.

An important result, due to Bregman & Fokin − and refined by Cai & Daskalakis − shows that the computation of a Nash equilibrium of a pairwise zero-sum GPM game can be done efficiently. For simplicity, and for the purpose of this lecture on spot-checking games, we only study the simple case where G is a star graph:

$$V = \{0, 1, \ldots, N\}, \qquad E = \{(0, i) : \ i \in [N]\}.$$

To simplify the notation in this case ($\mathcal{G}$ is a pairwise zero-sum, star-shaped polymatrix game), denote by $\boldsymbol{q}$ the strategy played by the central player, and by $\boldsymbol{p_i}$ the strategy of the i^{th} player ($i = 1, \ldots, N$). We denote the matrix of the partial game between player 0 and player i by A_i (instead of A_{0i}), so that the payoffs write:

$$\pi_0(\boldsymbol{p}, \boldsymbol{q}) = \sum_i \boldsymbol{q}^T A_i \boldsymbol{p_i}$$
$$\pi_i(\boldsymbol{p}, \boldsymbol{q}) = -\boldsymbol{q}^T A_i \boldsymbol{p_i} \qquad (\forall i \in [N]).$$

Remark 1. *Some authors call $\mathcal{G}$ a Bayesian game, because we can interprete it as follows: Player* 0 *will play a game against one of the other players ($1 \leq i \leq n$), which is chosen at random with given probabilities. Player* 0 *has to choose her strategy before knowing her opponent, and tries to mazimize her expected payoff.*

In this case, a NE can be found by solving a LP, which can be thought as a generalization of the minimax theorem.

Theorem 1 (Nash Equilibriums of a pairwise zero-sum, star-shaped GPM game)**.** *Consider the following pair of primal/dual LPs:*

$$\max_{\boldsymbol{q},\boldsymbol{u}} \quad \sum_{i\in[N]} u_i \tag{1a}$$
$$\text{s.t.} \quad \forall i \in [N],\ A_i^T \boldsymbol{q} \geq u_i \tag{1b}$$
$$\boldsymbol{q} \in \Delta_{n_0} \tag{1c}$$

$$\min_{\boldsymbol{p_1},\ldots,\boldsymbol{p_N},v} \quad v \tag{2a}$$
$$\text{s.t.} \quad \sum_{i\in[N]} A_i \boldsymbol{p_i} \leq v \tag{2b}$$
$$\forall i \in [N],\ \boldsymbol{p_i} \in \Delta_{n_i}, \tag{2c}$$

where the inequalities (1b) *and* (2b) *are componentwise, i.e.* $A_i^T\boldsymbol{q} \geq u_i$ *should be understood as* $\forall j \in [n_i],\ (A_i^T\boldsymbol{q})_j \geq u_i$. *The profile* $(\boldsymbol{q}, \boldsymbol{p_1}, \ldots, \boldsymbol{p_N})$ *is a NE of* $\mathcal{G}$ *if and only if these vectors are part of a solution of* (1)-(2).

Proof. By definition, $(\boldsymbol{q}, \boldsymbol{p_1}, \ldots, \boldsymbol{p_N})$ is a NE of $\mathcal{G}$ iff these vectors simultaneously solve the following $N+1$ optimization problems:

$$\max_{\boldsymbol{q}\in\Delta_{n_0}} \sum_i \boldsymbol{q}^T A_i \boldsymbol{p_i} \tag{P_0}$$
$$\min_{\boldsymbol{p_i}\in\Delta_{n_i}} \boldsymbol{q}^T A_i \boldsymbol{p_i} \qquad (\forall i \in [N]) \tag{$P_1,\ldots,P_N$}$$

Let $(\boldsymbol{q}, \boldsymbol{p_1}, \ldots, \boldsymbol{p_N})$ be a NE. Since $\boldsymbol{p_i}$ is a solution of (P_i), we have $\boldsymbol{q}^T A_i \boldsymbol{p_i} = \min_{j\in[n_i]}(A_i^T\boldsymbol{q})_j$ at the equilibrium. Substituting in Problem (P_0) and introducing a variable u_i such that $\min_{j\in[n_i]}(A_i^T\boldsymbol{q})_j \geq u_i$ yields the LP (1). The problems $(P_1, \ldots, P_N)$ are mutually independent, and hence they reduce to the single problem

$$\min_{(\boldsymbol{p_1},\ldots,\boldsymbol{p_N})\in\Delta_{n_1}\times\ldots\times\Delta_{n_N}} \sum_i \boldsymbol{q}^T A_i \boldsymbol{p_i}. \tag{$P_1+\ldots+P_N$}$$

Now, we use that $\boldsymbol{q}$ solves Problem (P_0), which implies $\sum_i \boldsymbol{q}^T A_i \boldsymbol{p_i} = \max_{j\in[n_0]}(\sum_i A_i\boldsymbol{p_i})_j$. Substituting in Problem $(P_1 + \ldots + P_N)$ and introducing the auxiliary variable v yields the LP (2).

Conversely, let $(\boldsymbol{q}, \boldsymbol{u})$ and $(\boldsymbol{p_1}, \ldots, \boldsymbol{p_N}, v)$ be solutions of (1)-(2). Recall the complementary slackness condition for inequality (1b): $\forall j \in [n_i], \Big((A_i^T\boldsymbol{q})_j = u_i \quad \text{OR} \quad (\boldsymbol{p_i})_j = 0\Big)$. This implies $u_i = \min_{j\in[n_i]}(A_i^T\boldsymbol{q})_j = \boldsymbol{q}^T A_i \boldsymbol{p_i}$, and hence $\boldsymbol{p_i}$ is a solution of (P_i). Similarly, the complementary slackness for (2b) says $v = \max_{j\in[n_0]}(\sum_i A_i\boldsymbol{p_i})_j = \sum_i \boldsymbol{q}^T A_i \boldsymbol{p_i}$, and so $\boldsymbol{q}$ is a solution of Problem (P_0). ☐

Similarly to what we did for congestion games, we also define the concept of Stackelberg equilibrium. Imagine that one of the player is a *leader*, who announces her strategy before playing. The other players (the *followers*) then react to the leader with a best response. Formally:

Definition 3 (Stackelberg equilibrium)**.** Let $\boldsymbol{\pi}$ be a $N-$player game in standard form, in which the first player (indexed by 1) is the leader. Define the set of partial Nash strategy profiles for the followers:

$$\mathcal{P}_{eq} := \big\{\boldsymbol{p} \in \Delta_1 \times \ldots \times \Delta_N :\ \forall i \neq 1,\ \boldsymbol{p_i} \in BR_i(\boldsymbol{p_{-i}})\big\}.$$

We say that the profile $(\boldsymbol{p_1}, \ldots, \boldsymbol{p_N})$ is a Stackelberg equilibrium of $\mathcal{G}$ if it maximizes the leader's payoff, among the set of all profiles $\boldsymbol{p} \in \mathcal{P}_{eq}$ such that the followers' strategies $\boldsymbol{p_{-1}}$ are in best response relationship to each other's action:

$$\boldsymbol{p} = (\boldsymbol{p_1}, \boldsymbol{p_{-1}}) \text{ is a Stackelberg equilibrium} \iff \boldsymbol{p} \in \underset{(\boldsymbol{p_1},\boldsymbol{p_{-1}})\in\mathcal{P}_{eq}}{\operatorname{argmax}} \pi_1(\boldsymbol{p_1}, \boldsymbol{p_{-1}})$$

$$\iff \begin{cases} (\boldsymbol{p_1}, \boldsymbol{p_{-1}}) \in \mathcal{P}_{eq} \\ \forall (\boldsymbol{p'_1}, \boldsymbol{p'_{-1}}) \in \mathcal{P}_{eq}, \quad \pi_1(\boldsymbol{p_1}, \boldsymbol{p_{-1}}) \geq \pi_1(\boldsymbol{p'_1}, \boldsymbol{p'_{-1}}) \end{cases}$$

Example: Consider the following (non-zero sum) bimatrix game:

	L	R
U	2, 1	4, 0
B	1, 0	3, 1

The strategy B for the row player is strictly dominated by U. So every NE for the row player puts all the weight on U, and the best response of the column player is to choose L. This is the unique NE, with a payoff of 2 for the row player. #1

Now, assume we have a Stackelberg game with the row player as leader. The row player could commit to the strategy B, which gives the column player an incentive to play R, and yields a payoff of 3. But with mixed strategies she can even do better. If a weight $p_U = 0.5 + \epsilon$ with $\epsilon > 0$ is assigned to strategy U, the column player prefers R, so the payoff for the row player is $3.5 - \epsilon$. The Stackelberg equilibrium is actually attained for $\epsilon = 0$. In this case, the column player could play indifferently L or R, but our definition of a Stackelberg equilibrium indicates that she breaks the tie in favor of the leader, that is, the Stackelberg equilibrium is $p_U = p_B = 0.5$, $p_L = 0$, $p_R = 1$, with a payoff of 3.5 for the row player.

In the case of a pairwise zero-sum, star-shaped polymatrix game $\mathcal{G}$, there is also a nice relationship between Nash and Stackelberg equilibriums, when the leader corresponds to the node 0 at the center of the star.

Theorem 2 (Stackelberg Equilibriums of a pairwise zero-sum, star-shaped GPM game)**.** *Let $(\boldsymbol{q}, \boldsymbol{p_1}, \ldots, \boldsymbol{p_N})$ be a NE of $\mathcal{G}$, and let $(\boldsymbol{p'_1}, \ldots, \boldsymbol{p'_N}) \in \Delta_{n_1} \times \ldots \times \Delta_{n_N}$ be arbitrary best response strategies to $\boldsymbol{q}$:*

$$\forall i \in [N], \ \min_{j\in[n_i]} (A_i^T \boldsymbol{q})_j = \boldsymbol{q}^T A_i \boldsymbol{p_i},$$

that is, the vector $\boldsymbol{p_i}$ puts only weight on the minimal coordinates of $A_i^T\boldsymbol{q}$. Then, $(\boldsymbol{q}, \boldsymbol{p'_1}, \ldots, \boldsymbol{p'_N})$ is a Stackelberg equilibrium of the stackelberg game $\mathcal{G}$ with Player 0 as the leader. In particular, every Nash equilibrium is also a Stackelberg equilibrium.

Conversely, if $(\boldsymbol{q}, \boldsymbol{p'_1}, \ldots, \boldsymbol{p'_N})$ is a Stackelberg equilibrium (with Player 0 as the leader), then $\boldsymbol{q}$ is part of a NE profile, i.e., there exists $(\boldsymbol{p_1}, \ldots, \boldsymbol{p_N}) \in \Delta_{n_1} \times \ldots \times \Delta_{n_N}$ such that $(\boldsymbol{q}, \boldsymbol{p_1}, \ldots, \boldsymbol{p_N})$ is a NE of $\mathcal{G}$.

Proof. Let $(\boldsymbol{q}, \boldsymbol{p'_1}, \ldots, \boldsymbol{p'_N})$ be as in the theorem. We have $\boldsymbol{p'_i} \in BR_i(\boldsymbol{p'_{-i}})$ for all $i \geq 1$, so $(\boldsymbol{q}, \boldsymbol{p'_1}, \ldots, \boldsymbol{p'_N})$ is a partial Nash equilibrium in $\mathcal{P}_{eq}$, in the sense of Definition 3. We know from the previous theorem that $\boldsymbol{q}$ solves the LP (1), so $\boldsymbol{q}$ maximizes $\sum_i \min_j (A_i^T\boldsymbol{q})_j$ over Δ_{n_0}. Now, let $(\bar{\boldsymbol{q}}, \bar{\boldsymbol{p}}_{\boldsymbol{1}}, \ldots, \bar{\boldsymbol{p}}_{\boldsymbol{N}})$ be another profile in $\mathcal{P}_{eq}$. We have $\min_{j\in[n_i]}(A_i^T\bar{\boldsymbol{q}})_j = \bar{\boldsymbol{q}}^T A_i \bar{\boldsymbol{p}}_{\boldsymbol{i}}$ for all i, and hence

$$\pi_0(\bar{\boldsymbol{p}}, \bar{\boldsymbol{q}}) = \sum_i \min_{j\in[n_i]} (A_i^T \bar{\boldsymbol{q}})_j \leq \sum_i \min_{j\in[n_i]} (A_i^T \boldsymbol{q})_j = \pi_0(\boldsymbol{p'}, \boldsymbol{q}).$$

This shows that $(\boldsymbol{q}, \boldsymbol{p'})$ is a Stackelberg equilibrium of $\mathcal{G}$ (with Player 0 as the leader).

Conversely, if $(\boldsymbol{q}, \boldsymbol{p'})$ is a Stackelberg equilibrium of $\mathcal{G}$, then $\boldsymbol{q}$ must maximize $\sum_i \min_j (A_i^T \boldsymbol{q})_j$ over Δ_{n_0}, and so it is a solution of the primal LP (1). By Theorem 1, any solution $\boldsymbol{p}$ of the dual LP (2) yields a Nash equilibrium $(\boldsymbol{q}, \boldsymbol{p})$ of $\mathcal{G}$. Such a solution always exists, because (2) is feasible and bounded. $\square$

We point out that this theorem is true only if the leader is the player at the center of the star (index 0). For a counter example, see Exercise 6 of Worksheet #15.

Let us now turn to the case where the game is not zero sum. We keep the star pattern of the graph G, and for all i we set $A_i := A_{0i}$, $B_i := -A_{i0}^T$, so that the payoff reads:

$$\begin{aligned} \pi_0(\boldsymbol{p}, \boldsymbol{q}) &= \sum_i \boldsymbol{q}^T A_i \boldsymbol{p_i} \\ \pi_i(\boldsymbol{p}, \boldsymbol{q}) &= -\boldsymbol{q}^T B_i \boldsymbol{p_i} \qquad (\forall i \in [N]). \end{aligned} \tag{3}$$

In this case, both Stackelberg and Nash equilibrium are hard to compute. However, the next theorem shows that a Stackelberg equilibrium with pure follower strategies always exist:

Theorem 3 (Stackelberg equilibrium with pure follower strategies)**.** *Consider the star-shaped GPM game $\mathcal{G}$ with payoffs described in Eq.* (3) *and Player 0 as leader. There exists a Stackelberg equilibrium profile $(\boldsymbol{q}, \boldsymbol{p_1}, \ldots, \boldsymbol{p_N})$ in which every follower's strategy $\boldsymbol{p_i}$ is pure, that is, $\boldsymbol{p_i}$ has a single nonzero component (which must be* 1*).*

Proof. First note that the existence part of this theorem is not evident (actually a Stackelberg profile of mixed strategies exists for every polymatrix game, but we won't prove this in the lecture). But we can write an optimization problem with a *supremum*, such that every profile attaining the maximum is a Stackelberg equilibrium:

$$\begin{aligned} \sup_{\boldsymbol{p}, \boldsymbol{q}} \quad & \sum_i \boldsymbol{q}^T A_i \boldsymbol{p_i} \\ \text{s.t.} \quad & \forall i \in [N],\ \boldsymbol{p_i} \text{ minimizes } \boldsymbol{q}^T B_i \boldsymbol{p_i} \ \text{ over } \Delta_{n_i} \\ & \boldsymbol{q} \in \Delta_{n_0} \end{aligned} \tag{4}$$

This kind of optimization problem is called a *bilevel programming* problem, because some secondary variables (here, the $\boldsymbol{p_i}$) are constrained to solve an auxiliary optimization problem that depends on the primary variable (here, $\boldsymbol{q}$). Let $(\boldsymbol{q}, \boldsymbol{p})$ be a feasible profile for this problem (so we have $(\boldsymbol{q}, \boldsymbol{p}) \in \mathcal{P}_{eq}$). For all i, define $I_i(\boldsymbol{q}) := \arg\min_{j \in [n_i]} (B_i^T \boldsymbol{q})_j$. The vector $\boldsymbol{p_i}$ puts weight only on coordinates in $I_i(\boldsymbol{q})$ (because $\boldsymbol{p_i}$ is a best response to $\boldsymbol{q}$). Now choose a coordinate $j_i \in \arg\max_{j \in I_i(\boldsymbol{q})} (A_i^T \boldsymbol{q})_j$, and let $\boldsymbol{p'_i}$ be the unit vector with a 1 on the j_i^{th} coordinate. By definition, $(\boldsymbol{q}, \boldsymbol{p'})$ is feasible for Problem (4), and $\sum_i \boldsymbol{q}^T A_i \boldsymbol{p_i} \leq \sum_i \boldsymbol{q}^T A_i \boldsymbol{p'_i}$. This shows that we can restrain our search to pure strategies $\boldsymbol{p_i}$ of the followers:

$$\begin{aligned} \sup_{j_1, \ldots, j_N, \boldsymbol{q}} \quad & \sum_i (A_i^T \boldsymbol{q})_{j_i} \\ \text{s.t.} \quad & \forall i \in [N],\ j_i \in [n_i] \\ & \forall i \in [N], \forall j \in [n_i], \quad (B_i^T \boldsymbol{q})_j \geq (B_i^T \boldsymbol{q})_{j_i} \\ & \boldsymbol{q} \in \Delta_{n_0}, \end{aligned} \tag{5}$$

This problem is in fact the maximum of a finite number of LPs. Each LP is bounded (some might be infeasible), which guarantees that the supremum is attained, and proves the existence of a Stackelberg equilibrium (with pure strategies for the followers). $\square$

As a byproduct of this proof, we obtain a method to compute a Stackelberg equilibrium of the game $\mathcal{G}$, based on a sequence of $\prod_{i=1}^{N} n_i$ linear programming problems.

Corollary 4. *Consider the star-shaped graphical game $\mathcal{G}$ with payoffs described in Eq.* (3) *and Player 0 as leader. For all $\boldsymbol{j} = (j_1, \ldots, j_N) \in [n_1] \times \ldots \times [n_N]$, let $u(\boldsymbol{j}) \in \mathbb{R} \cup \{-\infty\}$ denote the value of the following LP:*

$$\begin{aligned} \max_{\boldsymbol{q}} \quad & \sum_i (A_i^T \boldsymbol{q})_{j_i} \\ \text{s.t.} \quad & \forall i \in [N], \forall j \in [n_i], \quad (B_i^T \boldsymbol{q})_j \geq (B_i^T \boldsymbol{q})_{j_i} \\ & \boldsymbol{q} \in \Delta_{n_0}, \end{aligned} \tag{$P_{\boldsymbol{j}}$}$$

If $\boldsymbol{j}^*$ maximizes $u(\boldsymbol{j})$ over $[n_1] \times \ldots \times [n_N]$, then the profile $(\boldsymbol{q}(\boldsymbol{j}^*), \boldsymbol{p}(\boldsymbol{j}^*))$ is a Stackelberg equilibrium, where $\boldsymbol{q}(\boldsymbol{j}^*)$ is a solution of $(P_{\boldsymbol{j}^*})$ and $\boldsymbol{p_i}(\boldsymbol{j}^*)$ is the unit vector with a 1 on the j_i^{*th} coordinate.