Finding Equilibria for Two-Player Games Using Labeled Prototypes

Definition of Bi-Matrix Games

Bi-Matrix Games:

- A two-player game represented by two payoff matrices A and B.
- Player 1 chooses a strategy from the set $M = \{1, 2, \ldots, m\}$.
- Player 2 chooses a strategy from the set $N = \{m+1, m+2, \ldots, m+n\}$.
- Payoff matrices:
  - $A \in \mathbb{R}^{m \times n}$: Payoffs for Player 1.
  - $B \in \mathbb{R}^{m \times n}$: Payoffs for Player 2.
- If Player 1 chooses $i \in M$ and Player 2 chooses $j \in N$, the payoffs are A_{ij} for Player 1 and B_{ij} for Player 2.

Supports in Bi-Matrix Games

Support:

- The *support* of a mixed strategy is the set of pure strategies played with non-zero probability.
- For Player 1, a mixed strategy is a probability vector $x \in \mathbb{R}^m$ such that $x \geq 0$ and $\sum_{i=1}^{m} x_i = 1$.
- For Player 2, a mixed strategy is a probability vector $y \in \mathbb{R}^n$ such that $y \geq 0$ and $\sum_{j=1}^{n} y_j = 1$.
- A pure strategy is a mixed strategy where only one component is non-zero.

Best Response Condition

Best Response:

- Player 1's payoff when using strategy x against y is given by:

$$u_1(x, y) = x^T Ay.$$

- Player 1's strategy x is a *best response* to Player 2's strategy y if:

$$x_i > 0 \implies (Ay)_i = \max_{k \in M}(Ay)_k.$$

- Similarly, Player 2's strategy y is a *best response* to Player 1's strategy x if:

$$y_j > 0 \implies (x^T B)_j = \max_{l \in N}(x^T B)_l.$$

Nash Equilibrium

Nash Equilibrium:

- A pair of mixed strategies (x^*, y^*) is a *Nash equilibrium* if:

$$
\begin{aligned}
x^T A y^* \geq (x')^T A y^* \quad & \text{for all } x' \in X, \\
(x^*)^T B y \geq (x^*)^T B y' \quad & \text{for all } y' \in Y.
\end{aligned}
$$

- Intuitively, neither player can improve their payoff by unilaterally deviating from their strategy.

In Terms of Best Response:

- A Nash equilibrium is a pair of strategies where each strategy is a best response to the other.

Polyhedron Definitions (Part 1)

Polyhedron:

- A polyhedron P in $\mathbb{R}^d$ is a set $\{z \in \mathbb{R}^d \mid Cz \leq q\}$ for some matrix C and vector q.
- It is called *full-dimensional* if it has dimension d.
- It is called a *polytope* if it is bounded.

Faces of a Polyhedron:

- A face of P is a set $\{z \in P \mid cz = q_0\}$ for some $c \in \mathbb{R}^d$, $q_0 \in \mathbb{R}$, where the inequality $cz \leq q_0$ holds for all $z \in P$.
- A vertex of P is the unique element of a zero-dimensional face.
- An edge of P is a one-dimensional face.

Polyhedron Definitions (Part 2)

Facets and Properties of Faces:

- A facet of a d-dimensional polyhedron P is a face of dimension $d - 1$.
- Any nonempty face F of P can be obtained by turning some of the inequalities defining P into equalities, called *binding inequalities*.
- $F = \{z \in P \mid c_i z = q_i, i \in I\}$, where $c_i z \leq q_i$ for $i \in I$ are some rows in $Cz \leq q$.

Simple Polyhedron:

- A d-dimensional polyhedron P is *simple* if no point belongs to more than d facets of P.
- This holds if there are no special dependencies between the facet-defining inequalities.

Best Response Polyhedron

Best Response Polyhedron:

- The best response polyhedron of a player is the set of that player's mixed strategies together with the *upper envelope* of expected payoffs (and any larger payoffs) to the other player.
- The best response polyhedron P for Player 1 is defined analogously.
- Generally:

$$P = \{(x, v) \in \mathbb{R}^M \times \mathbb{R} \mid x \geq 0, 1^T x = 1, B^T x \leq 1v\},$$
$$Q = \{(y, u) \in \mathbb{R}^N \times \mathbb{R} \mid Ay \leq 1u, y \geq 0, 1^T y = 1\}.$$

Example

Example:

- Consider the following payoff matrices for a two-player game:

$$A = \begin{bmatrix} 3 & 3 \\ 2 & 5 \\ 0 & 6 \end{bmatrix}, \quad B = \begin{bmatrix} 1 & 0 \\ 0 & 2 \\ 4 & 3 \end{bmatrix}$$

- Player 1 chooses rows, and Player 2 chooses columns.
- Payoff for Player 1 is given by matrix A.
- Payoff for Player 2 is given by matrix B.

Example

Best Response Polyhedron for Player 1:

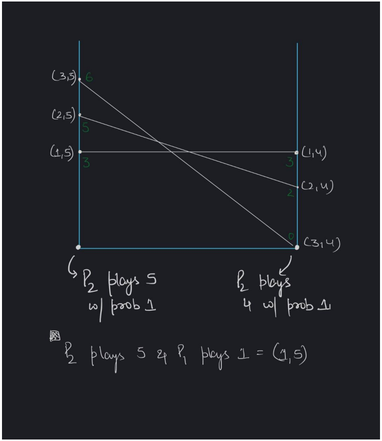

(a)

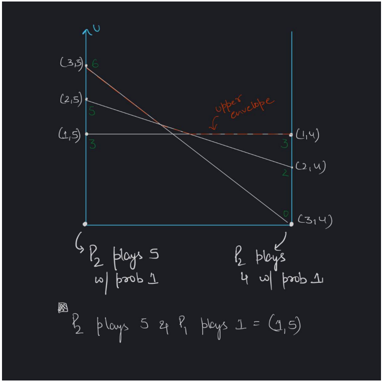

(b)

Example

Best Response Polyhedron for Player 1:

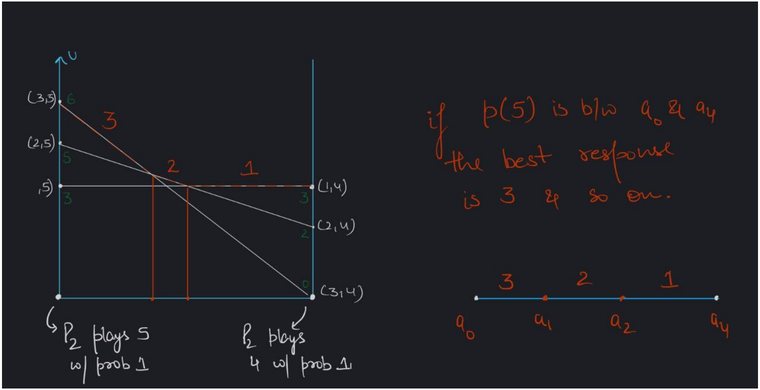

(c)

Example

Best Response Polyhedron for Player 2:

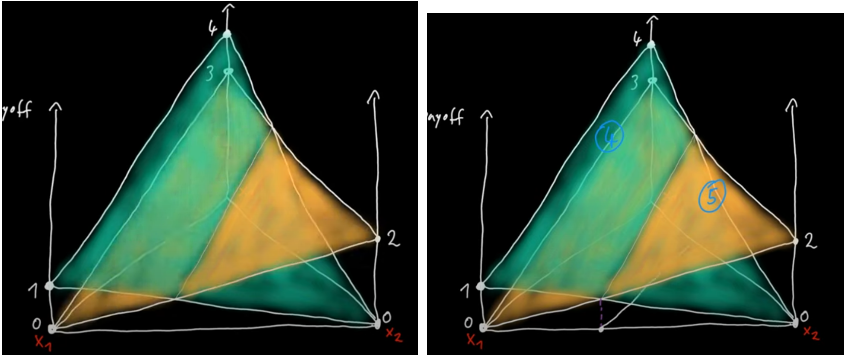

(d) (e)

Example

Best Response Polyhedron for Player 2:

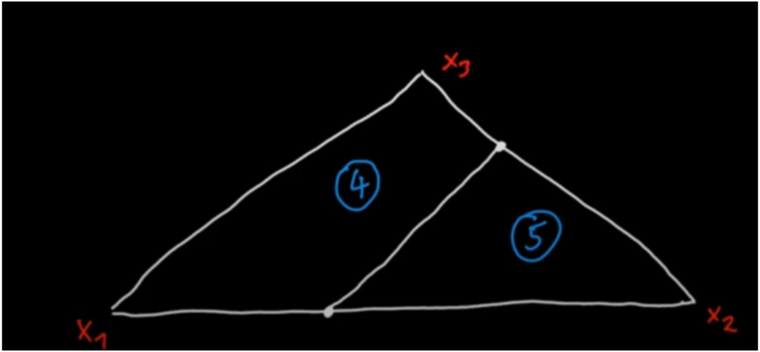

(f)

Labels in Best Response Polyhedra

Labels:

- A point $(y, u) \in Q$ has a label $k \in M \cup N$ if:
  - For $k = i \in M$ (Player 1's pure strategies): The i-th inequality $\sum_{j \in N} a_{ij} y_j = u$ is binding, meaning i is a best response to y with payoff u.
  - For $k = j \in N$ (Player 2's pure strategies): The inequality $y_j = 0$ is binding, meaning j is not played.
- Similarly, a point $(x, v) \in P$ has a label $k \in M \cup N$ if:
  - For $k = i \in M$: The inequality $x_i = 0$ is binding, meaning i is not played.
  - For $k = j \in N$: The inequality $\sum_{i \in M} b_{ij} x_i = v$ is binding, meaning j is a best response to x with payoff v.
- A pair $((x, v), (y, u)) \in P \times Q$ is *completely labeled* if every $k \in M \cup N$ appears as a label of either (x, v) or (y, u).

Why Labels?

Purpose of Labels:

- Ensure the **best response condition**:
  - Played strategies ($x_i > 0$ or $y_j > 0$) yield the maximum payoff.
  - Unplayed strategies ($x_i = 0$ or $y_j = 0$) are explicitly excluded.
- Guarantee a **complete representation** of all pure strategies:
  - Every pure strategy is either a best response or unused.
  - No strategy is ignored.

Complete Labeling:

- Every label $k \in M \cup N$ must appear in either P or Q:
  - Ensures all pure strategies are evaluated.
  - Captures the mutual best-response structure of a Nash equilibrium.

What Goes Wrong Without Labels:

- Violation of Nash equilibrium:
  - Strategies with positive probability may not be optimal (not best responses).
  - Some strategies may be "ignored," leading to incomplete or invalid equilibria.

Adding Labels to our example

In the two examples we have, we will add points where a pure strategy is played with probability zero.

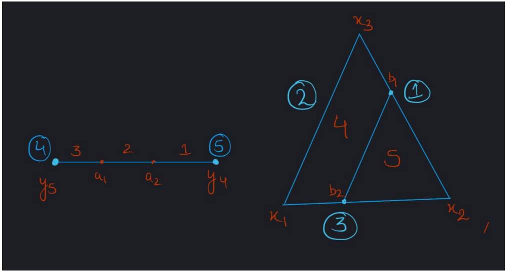

In the first best-response-polyhedron (BRP) of P_1 we have marked 4 and 5 because that is where moves 4 and 5 have zero probability of being played by player 2.
Similarly, in the BRP of P_2, we have marked the line x_1x_3 with 2 since the probability of 2 being played by P1 is 0 on that line.

Adding Labels to our example

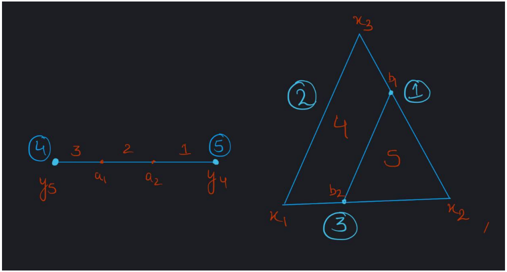

Consider (a_1, b_1)
On a_1 we can play both 3 and 2.
On b_1 we can play both 4 and 5.
The only thing thats missing is 1 which also has a zero probability on b_1.
$\text{Label}(a_1, b_1) = M \cup N$ hence that pair of point is a $nash - equilibrium$.

Simplifying Best Response Polytopes

The Problem: The initial best response polyhedra can be complicated:

- They are often **unbounded**.
- They contain payoff variables (u and v) in their coordinates.

This makes it harder to analyze and solve for equilibrium strategies.

Our Approach: We simplify this by:

- Assuming matrices M and N^T are **non-negative** and have no zero columns (we can adjust payoffs to ensure this).
- **Normalizing** the payoffs by dividing each inequality $N^T x \leq v$ by v, which reduces the payoff scale and simplifies the model.

Normalized Best Response Polyhedra: After normalization, the best response polytopes for both players become:

- **P** for Player 1: $P = \left\{x \in \mathbb{R}^m : x \geq 0, N^T x \leq 1\right\}$
- **Q** for Player 2: $Q = \{y \in \mathbb{R}^n : y \geq 0, My \leq 1\}$

Now What?

1. **Structure of Nondegenerate Games:**
   - A bimatrix game is **nondegenerate** if each mixed strategy with support of size k has at most k pure best responses.
   - This ensures that strategies have a well-defined and limited number of best responses, avoiding ambiguities.
2. **Role of Vertices in P and Q:**
   - The polytopes P (for Player 1) and Q (for Player 2) are defined by linear constraints derived from the game's payoffs.
   - Each point in P or Q represents a mixed strategy, and their **vertices** represent pure or extreme strategies.
   - In a nondegenerate game:
     - Each point in P has at most m labels, where m is the dimension of P.
     - Each point in Q has at most n labels, where n is the dimension of Q.

Now What

- Recall that a bimatrix game is nondegenerate if there are at most k pure best responses to every mixed strategy with support of size k.
- In these games NE correspond to pairs of completely labeled vertices.
    - Since G is nondegenerate, each point of P has at most m labels. This is because if x has support of size k, then x has at most $m - k$ labels in A_1 and so if x had more than m labels, then x would have more than k best responses in A_2. Analogously, each point of Q has at most n labels.
- Thus, P and Q are both simple polytopes (each point of P or Q contained in more than m or n facets has more than m or n labels).
- Since $\dim(P) = m$ and $\dim(Q) = n$, only vertices of P and Q can have m and n labels.
- Since we need total $\|m\| + \|n\|$ labels, we only need to check vertices for the Nash Equilibrium!!