

Lesson 33. LP Duality and Game Theory

Today...

- LP duality and two-player zero-sum game theory

Game theory

- Game theory is the mathematical study of strategic interactions, in which an individual's success depends on his/her own choice as well as the choices of others
- We'll look at one type of game, and use LP duality to give us some insight about behavior in these games

Two-player zero-sum games

- Two players make decisions simultaneously
- Payoff depends on joint decisions
- Zero-sum: whatever one person wins, the other person loses
- Examples:
 - Rock-paper-scissors
 - Advertisers competing for market share (gains/losses over existing market share)

Payoff matrices

- 2 players
 - player R (for “row”)
 - player C (for “column”)
- Player R chooses among m rows (**actions**)
- Player C chooses among n columns
- Example: rock-paper-scissors, $m = 3$, $n = 3$

	Rock	Paper	Scissors
Rock	0	-1	1
Paper	1	0	-1
Scissors	-1	1	0

- This is the **payoff matrix** for player R
- Zero-sum: Player C receives the negative

- Another example: $m = 2, n = 3$

	1	2	3
1	-2	1	2
2	2	-1	0

- Player R chooses row 2, Player C chooses column 1
- What is the payoff of each player?

Pure and mixed strategies

- **Pure strategy:** pick one row (or column) over and over again
- **Mixed strategy:** each player assigns probabilities to each of his/her strategies
- For example:

	1	2	3
1	-2	1	2
2	2	-1	0
3	1	0	-2

- Suppose player R plays all three actions with equal probability
 - Row 1 with probability 1/3
 - Row 2 with probability 1/3
 - Row 3 with probability 1/3
- For example:

	1	2	3	Prob.
1	-2	1	2	1/3
2	2	-1	0	1/3
3	1	0	-2	1/3
Expected payoffs				

- Suppose player R plays all three actions with equal probability
- ⇒ Can compute **expected payoffs**:
- If player C plays
 - * column 1:
 - * column 2:
 - * column 3:

Who has the advantage?

	1	2	3
1	-2	1	2
2	2	-1	0
3	1	0	-2

- Can we find “optimal” (mixed) strategies for two-player zero-sum games?
- What can player R guarantee in return, regardless of what C chooses?

Player R and payoff lower bounds

- Suppose Player R plays all three actions with equal probability
- With this mixed strategy, R can guarantee a payoff of at least:
- This is a lower bound on the payoff R gets when playing (1/3, 1/3, 1/3)

Player C and payoff upper bounds

	1	2	3	Expected payoff (for R)
1	-2	1	2	
2	2	-1	0	
3	1	0	-2	
Prob.	1/3	1/3	1/3	

- Player C’s payoff = -(Player R’s payoff)
- Player C wants to limit Player R’s payoff
- Suppose Player C plays all three actions with equal probability
- With this mixed strategy, C can guarantee that R gets a payoff of at most:
- This is an upper bound on the payoff R gets when C plays (1/3, 1/3, 1/3)

Let’s optimize: Player R’s problem

- Want to decide mixed strategy that maximizes guaranteed payoff

⇒ Decision variables:

$$x_i = \text{prob. of choosing action } i$$

$$\text{for } i \in \{1, 2, 3\}$$

	1	2	3	Prob
1	-2	1	2	x_1
2	2	-1	0	x_2
3	1	0	-2	x_3

- Optimization model:

- Player R's problem: maximin
- Convert Player R's problem to LP:

Player C's problem

- Want to decide mixed strategy that limits Player R's payoff

⇒ Decision variables:

$y_i = \text{prob. of choosing action } i$
for $i \in \{1, 2, 3\}$

	1	2	3
1	-2	1	2
2	2	-1	0
3	1	0	-2
Prob.	y_1	y_2	y_3

- Optimization model:

- Player C's problem: minimax
- Convert Player C's problem to LP:

Optimal mixed strategy for Player R

	1	2	3	Prob.
1	-2	1	2	7/18
2	2	-1	0	5/18
3	1	0	-2	1/3
Expected payoff	1/9	1/9	1/9	

- Solve Player R's LP

⇒ Optimal mixed strategy for R guarantees that R can get at least

$$\min\{1/9, 1/9, 1/9\} = 1/9$$

- “Maximin” payoff = 1/9

Optimal mixed strategy for Player C

	1	2	3	Expected payoff (for R)
1	-2	1	2	1/9
2	2	-1	0	1/9
3	1	0	-2	1/9
Prob.	1/3	5/9	1/9	

- Solve Player C's LP

⇒ Optimal mixed strategy for C guarantees that C can limit R's payoff to at most

$$\max\{1/9, 1/9, 1/9\} = 1/9$$

- “Minimax” payoff = 1/9
- “Maximin” payoff = “Minimax” payoff - **not** a coincidence

Fundamental Theorem of 2-Player 0-Sum Games

- For any 2-player 0-sum game, let

$p_R(\mathbf{x})$ = lower bound on R's payoff if R plays probability vector \mathbf{x}

$p_C(\mathbf{y})$ = upper bound on R's payoff if C plays probability vector \mathbf{y}

- Let

\mathbf{x}^* = maximizer of p_R

\mathbf{y}^* = minimizer of p_C

- Then,

$$p_R(\mathbf{x}^*) = p_C(\mathbf{y}^*)$$

- That is, maximin payoff = minimax payoff

- Why is this remarkable?
 - Think back to example
 - Imagine you are Player R, and you have to announce in advance what your mixed strategy is
 - Intuitively, this seems like a bad idea
 - But, if you play the optimal maximin strategy, you are guaranteed an expected payoff of $1/9$
 - And, Player C cannot do anything to prevent this
 - Announcing the strategy beforehand does not cost you in this case
- Why is this true?
 - Turns out Player R's LP and Player C's LP form a primal-dual pair
 - Theorem follows immediately from strong duality for LP
 - For example, after some manipulation, it is easy to see that in our game, Player R's LP and Player C's LP are duals of each other

Player R's LP:

$$\begin{array}{ll}
 \max & z \\
 \text{s.t.} & 2x_1 - 2x_2 - x_3 + z \leq 0 \\
 & -x_1 + x_2 + z \leq 0 \\
 & -2x_1 + 2x_3 + z \leq 0 \\
 & x_1 + x_2 + x_3 = 1 \\
 & x_1, x_2, x_3 \geq 0
 \end{array}$$

Player C's LP:

$$\begin{array}{ll}
 \min & w \\
 \text{s.t.} & 2y_1 - y_2 - 2y_3 + w \geq 0 \\
 & -2y_1 + y_2 + w \geq 0 \\
 & -y_1 + 2y_3 + w \geq 0 \\
 & y_1 + y_2 + y_3 = 1 \\
 & y_1, y_2, y_3 \geq 0
 \end{array}$$

- Homework: verify that these LPs are duals of each other