CS 170: Zero-Sum Games — Lecture Notes

Course: CS 170 — Efficient Algorithms and Intractable Problems, Spring 2026 **Instructors:** Lijie Chen, Umesh V. Vazirani

---

Part 1: What Is a Zero-Sum Game?

You and a friend are about to play a game. You each make a choice — simultaneously, without seeing each other — and then money changes hands based on the outcome. Your gain is your friend's loss, and vice versa. This is a **zero-sum game**: the total payoff is always zero, because whatever one player wins, the other loses.

Zero-sum games show up everywhere: poker, penalty kicks in soccer, pricing wars between companies, even the design of algorithms that must work against adversaries. In this lecture, we'll see that LP duality — the theory we developed in the previous lectures — gives us a complete solution to these games.

1.1 Matching Pennies

Here's the simplest interesting game. You and your friend each secretly choose Heads or Tails. You reveal simultaneously. If the coins match, your friend pays you $1. If they differ, you pay your friend $1.

What should you do? If you always play Heads, your friend catches on and always plays Tails — you lose every round. Switch to Tails, and your friend switches to Heads. You're stuck in a cat-and-mouse loop: no matter what you commit to, your friend can exploit it.

1.2 Rock-Paper-Scissors

Now consider the game everyone knows. Rock beats Scissors, Scissors beats Paper, Paper beats Rock. If you win you get $1, if you lose you pay $1, ties pay nothing.

The same cat-and-mouse problem appears, but now with three options: always play Rock? Your friend plays Paper. Switch to Paper? Your friend plays Scissors. No fixed choice is safe. The game is also perfectly symmetric — neither player has a structural advantage.

1.3 Who Goes First Matters — But Only If You're Deterministic

There's a crucial observation lurking in these examples. Imagine the game isn't played simultaneously — instead, one player has to commit first, and the other gets to see the commitment before responding.

In Matching Pennies, if you go first and announce "I'm playing Heads," your friend sees this and plays Tails — you lose. If you announce "I'm playing Tails," your friend plays Heads — you lose again. Going first is a disaster: *whatever* deterministic choice you make, your friend exploits it. You can guarantee at most $-\$1$. But if your friend goes first, the tables turn: you see their choice and pick the matching coin, guaranteeing $+\$1$.

The same thing happens in Rock-Paper-Scissors: the player who commits to a fixed choice first always loses $1, the player who sees it and responds always wins $1.

Now here's where it gets interesting. What if you're allowed to *randomize*? Suppose you go first, but instead of announcing a fixed choice, you announce: "I will play Heads with probability 1/2 and Tails with probability 1/2." Your friend knows your *strategy* — the probabilities — but doesn't know which coin will actually come up. It's as if you flip a fair coin to decide, and your friend has to choose before seeing the flip. No matter what your friend picks, your expected payoff is 0.

And if your friend goes first with the same 50-50 announcement? You face the same situation — your expected payoff is again 0.

So with deterministic strategies, going first costs you $1 in Matching Pennies. But with randomization, going first costs you nothing — both players can guarantee the same expected payoff of $0 regardless of the order. The remarkable punchline of this lecture is that **this is always true**: for *any* zero-sum game, if both players use the right randomized strategy, the order of play doesn't matter. There is a single number — the **value** of the game — that both players can guarantee.

Part 2: Mixed Strategies and the Minimax Theorem

2.1 The Payoff Matrix

To reason about games precisely, we write down a **payoff matrix** G. We call the two players **Row** and **Column**. Row picks a row i, Column picks a column j (simultaneously), and the entry G_{ij} is what Column pays Row. Positive entries are good for Row, negative entries are good for Column.

For Matching Pennies (rows/columns are Heads, Tails):

$$G = \begin{pmatrix} 1 & -1 \\ -1 & 1 \end{pmatrix}$$

For Rock-Paper-Scissors (rows/columns are Rock, Paper, Scissors):

$$G = \begin{pmatrix} 0 & -1 & 1 \\ 1 & 0 & -1 \\ -1 & 1 & 0 \end{pmatrix}$$

More generally, a two-player zero-sum game is defined by an $m \times n$ payoff matrix G, where m is the number of strategies for Row and n for Column.

2.2 Mixed Strategies

As we saw, any fixed deterministic choice — called a **pure strategy** — can be exploited. The way out is **randomization**: instead of committing to a single choice, a player picks a **probability distribution** over their options. This is called a **mixed strategy**.

For Row, a mixed strategy is a list of probabilities $(x_1, \dots, x_m)$ — one for each row — that are nonneg and sum to 1. In Matching Pennies, "play Heads with probability 1/2, Tails with probability 1/2" is the mixed strategy $(1/2, 1/2)$. In Rock-Paper-Scissors, "play each option with probability 1/3" is $(1/3, 1/3, 1/3)$. Column's mixed strategy $(y_1, \dots, y_n)$ works the same way.

When Row uses mixed strategy $\mathbf{x}$ and Column uses mixed strategy $\mathbf{y}$, the **expected payoff** to Row — averaging over both players' randomness — is $\mathbf{x}^\top G\,\mathbf{y}$, which is just the weighted average $\sum_{i,j} x_i\,G_{ij}\,y_j$.

2.3 What Can Each Player Guarantee?

In Part 1, we saw that the order of play matters enormously for deterministic strategies. Let's revisit this with mixed strategies using the language of maximin and minimax.

Imagine Row goes first: Row publicly announces a mixed strategy $\mathbf{x}$ (the *probabilities*, not the actual move — Column knows the distribution but not the realization). Column, knowing $\mathbf{x}$, picks the best response. The best Row can guarantee in this scenario is:

$$\max_{\mathbf{x}}\ \min_{\mathbf{y}}\ \mathbf{x}^\top G\,\mathbf{y}$$

This is the **maximin**: Row's guaranteed payoff even in the worst case where Column knows Row's strategy.

Now flip it: Column goes first, announces a mixed strategy $\mathbf{y}$, and Row responds optimally:

$$\min_{\mathbf{y}}\ \max_{\mathbf{x}}\ \mathbf{x}^\top G\,\mathbf{y}$$

This is the **minimax**: Column's guaranteed payment even in the worst case where Row knows Column's strategy.

Since going first is a disadvantage, the maximin should be $\leq$ the minimax — and indeed it always is. But is there a gap?

2.4 The Minimax Theorem

With deterministic strategies, the gap can be huge (in Matching Pennies: -1 vs. $+1$). But the stunning fact is that **with mixed strategies, the gap is always zero**:

> **Von Neumann's Minimax Theorem (1928).** For any payoff matrix G:
>
> $$\max_{\mathbf{x}}\ \min_{\mathbf{y}}\ \mathbf{x}^\top G\,\mathbf{y} \;=\; \min_{\mathbf{y}}\ \max_{\mathbf{x}}\ \mathbf{x}^\top G\,\mathbf{y}$$
>
> where $\mathbf{x}$ and $\mathbf{y}$ range over probability distributions (mixed strategies).

What does this really say? It says that **if both players play optimally, the order doesn't matter**. There is a single number V — the **value** of the game — such that:

- Row has a mixed strategy that guarantees an expected payoff of at least V, even if Column knows Row's strategy and responds optimally.
- Column has a mixed strategy that guarantees an expected payment of at most V, even if Row knows Column's strategy and responds optimally.

Neither player can do better. It doesn't matter who "goes first" — both players can announce their strategy openly, and the expected outcome is the same. For Matching Pennies: $V = 0$. For Rock-Paper-Scissors: $V = 0$.

How do we prove this? Through LP duality — which we develop in the next section.

Part 3: Solving Games with Linear Programming

3.1 Row's LP

The minimax theorem is not just a theoretical statement — it gives us an efficient algorithm for finding optimal strategies. The key observation is that each player's optimization problem is a **linear program**.

Row wants to choose a mixed strategy $\mathbf{x}$ to maximize the worst-case expected payoff. A useful observation: Column's best response to any fixed $\mathbf{x}$ is always a **pure strategy** (a single column j), because Column is minimizing a linear function over a simplex, and the minimum of a linear function over a polytope is always at a vertex. So Row's problem is:

$$\max_{\mathbf{x}} \min_{j} \sum_i x_i G_{ij}$$

To express this as an LP, introduce a variable V representing the worst-case payoff:

$$\begin{gathered} \max\ V \\ \sum_i x_i G_{ij} \geq V \quad \text{for } j = 1, \dots, n \\ \sum_i x_i = 1 \\ x_i \geq 0 \quad \text{for } i = 1, \dots, m \end{gathered}$$

The constraints say: "against every pure column strategy j, Row's expected payoff is at least V." Row maximizes this guaranteed floor.

3.2 Column's LP

By the same logic, Column wants to minimize the worst-case expected payment. Row's best response to any $\mathbf{y}$ is the pure strategy i maximizing $\sum_j G_{ij} y_j$, so Column's problem is:

$$\begin{gathered} \min\ W \\ \sum_j G_{ij} y_j \leq W \quad \text{for } i = 1, \dots, m \\ \sum_j y_j = 1 \\ y_j \geq 0 \quad \text{for } j = 1, \dots, n \end{gathered}$$

The constraints say: "against every pure row strategy i, Column's expected payment is at most W." Column minimizes this guaranteed ceiling.

3.3 Deriving Column's LP as the Dual of Row's LP

We now show that Column's LP arises naturally as the **dual** of Row's LP. The motivation is the same as in the LP I lecture: we want to certify an **upper bound** on Row's optimal value V^* by taking clever linear combinations of Row's constraints.

The multiplier idea. Recall how we derived the dual of the chocolate LP: we assigned nonneg multipliers to the primal's inequality constraints, added them up, and tried to produce an inequality bounding the objective from above. Let's do the same thing here — but Row's LP also has an equality constraint, which requires a small twist.

Row's LP has n inequality constraints and one equality constraint:

$$\sum_i x_i\, G_{ij} \geq V \quad \text{for } j = 1, \dots, n \qquad \text{and} \qquad \sum_i x_i = 1$$

Assign a nonneg multiplier $y_j \geq 0$ to each inequality constraint (nonneg so that multiplying preserves the $\geq$ direction). Multiply and sum:

$$\sum_j y_j \left(\sum_i x_i\, G_{ij} \right) \geq V \sum_j y_j$$

Now here is the key step. If we choose the multipliers so that $\sum_j y_j = 1$ — i.e., $\mathbf{y}$ is itself a probability distribution — the right-hand side simplifies to just V:

$$\sum_i x_i \left(\sum_j G_{ij}\, y_j \right) \geq V$$

The left-hand side is $\sum_i x_i \cdot$ (Row i's payoff against Column's mix $\mathbf{y}$). Since $x_i \geq 0$ and $\sum_i x_i = 1$ (using the equality constraint!), this is a convex combination, which is at most the largest term:

$$V \leq \sum_i x_i \left(\sum_j G_{ij}\, y_j \right) \leq \max_i \sum_j G_{ij}\, y_j$$

Call this upper bound $W = \max_i \sum_j G_{ij}\, y_j$. We have shown: **for any probability vector $\mathbf{y}$ (i.e., $y_j \geq 0$, $\sum_j y_j = 1$), the quantity $W = \max_i \sum_j G_{ij}\, y_j$ is an upper bound on V^*.** This is **weak duality** — any feasible solution to Column's problem bounds Row's optimum from above.

Tightening the bound. Column wants the tightest possible upper bound, so Column should minimize W over all probability vectors $\mathbf{y}$. Expressing $W \geq \sum_j G_{ij}\, y_j$ for each row i as explicit constraints, we arrive at:

$$\min\ W$$
$$\sum_j G_{ij}\, y_j \leq W \quad \text{for } i = 1, \dots, m$$
$$\sum_j y_j = 1, \quad y_j \geq 0$$

This is **Column's LP** — derived purely from the desire to certify an upper bound on Row's game value!

Notice the beautiful interplay: the inequality constraints of Row's LP gave rise to the dual variables y_j (Column's mixed strategy), and the equality constraint $\sum x_i = 1$ gave rise to the normalization $\sum y_j = 1$ (Column plays a probability distribution). The entire dual has a natural game-theoretic meaning.

Strong duality and the Minimax Theorem. Weak duality tells us $V^* \leq W^*$: Row's guarantee is at most Column's guarantee. The remarkable fact — **strong duality** — is that equality always holds: $V^* = W^*$. We state this without re-proving it (it follows from the simplex algorithm, just as in the basic LP case).

This is exactly the **Minimax Theorem**: the best payoff Row can guarantee equals the least payoff Column can be forced into. LP duality, which itself generalized the max-flow min-cut theorem, now proves one of the most famous results in all of game theory. ■

3.4 Worked Example: Rock-Paper-Scissors

Let us solve Rock-Paper-Scissors from Section 1.3. Recall the payoff matrix:

$$G = \begin{pmatrix} 0 & -1 & 1 \\ 1 & 0 & -1 \\ -1 & 1 & 0 \end{pmatrix}$$

Row's LP has variables x_1 (Rock), x_2 (Paper), x_3 (Scissors), and V:

$$\begin{gathered} \max\ V \\ 0 \cdot x_1 + x_2 - x_3 \geq V \quad \text{(Column plays Rock)} \\ -x_1 + 0 \cdot x_2 + x_3 \geq V \quad \text{(Column plays Paper)} \\ x_1 - x_2 + 0 \cdot x_3 \geq V \quad \text{(Column plays Scissors)} \\ x_1 + x_2 + x_3 = 1, \quad x_1, x_2, x_3 \geq 0 \end{gathered}$$

By the symmetry of the game, we guess $x_1 = x_2 = x_3 = 1/3$. Let's verify the constraints:

- Against Rock: $0(1/3) + 1(1/3) + (-1)(1/3) = 0$
- Against Paper: $(-1)(1/3) + 0(1/3) + 1(1/3) = 0$
- Against Scissors: $1(1/3) + (-1)(1/3) + 0(1/3) = 0$

All three payoffs equal 0, so $V = 0$ is feasible. Can Row do better? No — by the same symmetry, Column can also play $(1/3, 1/3, 1/3)$ and guarantee a payment of at most 0. By weak duality, $V^* \leq W^* = 0$. So the game value is $V^* = 0$, and the unique optimal strategy for both players is the uniform distribution.

Part 4: Summary

Concept	Key Idea
Zero-sum game	Two-player game with opposing interests; specified by payoff matrix G
Mixed strategy	Probability distribution over pure strategies
Expected payoff	$\mathbf{x}^\top G\,\mathbf{y}$ under mixed strategies $\mathbf{x}, \mathbf{y}$
Minimax Theorem	Mixed maximin = mixed minimax; every game has a well-defined value V
Row's LP	max V subject to $\sum_i x_i G_{ij} \geq V$, $\sum_i x_i = 1$, $x_i \geq 0$
Column's LP	min W subject to $\sum_j G_{ij} y_j \leq W$, $\sum_j y_j = 1$, $y_j \geq 0$
Proof via LP duality	Row's and Column's LPs are dual; strong duality $\Rightarrow$ minimax theorem

Previous lecture: LP duality, the dual LP, weak and strong duality, max-flow min-cut as a special case.

Part 5: Further Reading

- **Dasgupta, Papadimitriou, Vazirani — *Algorithms***, Chapter 7, Section 7.5. The primary reference for this lecture. Covers zero-sum games, the minimax theorem, and the connection to LP duality.
- **Tim Roughgarden — *Twenty Lectures on Algorithmic Game Theory***, Lecture 5. A clear and accessible treatment of zero-sum games and the minimax theorem.
- **John von Neumann and Oskar Morgenstern — *Theory of Games and Economic Behavior*** (1944). The foundational work that launched game theory.