Lecture 2

Lecturer: Constantinos Daskalakis *Scribe: Constantinos Daskalakis*

In this lecture, we focus on two-player zero-sum games. Our goal is to show that linear programming (LP) duality implies the existence of an efficiently computable Nash equilibrium in such games. Before describing this result, we introduce some definitions and terminology.

1 Preliminaries

Two-Player (Normal-Form) Games. A *two-player normal-form game* is specified via a pair (R, C) of $m \times n$ *payoff matrices*. The two players of the game, called the *row player* and the *column player*, have respectively m and n *pure strategies*. As the players' names imply, the pure strategies of the row player are in one-to-one correspondence with the rows of the payoff matrices, while the strategies of the column player correspond to columns. So if the row player plays strategy i and the column player strategy j, then their respective payoffs are given by $R[i,j]$ and $C[i,j]$. Players may also randomize over their strategies, leading to *mixed*—as opposed to pure—strategies. We represent the simplex of mixed strategies available to the row player by Δ_m and those available to the column player by Δ_n. Thus, if $\mathbf{x} \in \Delta_m$, then $\mathbf{x}$ is an m-dimensional vector $\langle x_1, x_2, \ldots, x_m \rangle$ such that $0 \le x_i \le 1$ for each $i \in [m] := \{1, \ldots, m\}$ and $\sum_{i=1}^m x_i = 1$; similarly for $\mathbf{y} \in \Delta_n$. For a pair of mixed strategies $\mathbf{x} \in \Delta_m$ and $\mathbf{y} \in \Delta_n$, the expected payoff of the row and column player are respectively $\mathbf{x}^T R \mathbf{y}$ and $\mathbf{x}^T C \mathbf{y}$.

Nash Equilibrium. A pair of strategies $(\mathbf{x}, \mathbf{y})$ is said to be a *Nash equilibrium* iff neither player can increase her expected payoff by unilaterally deviating from her strategy. Mathematically,

$$\begin{aligned} \mathbf{x}^T R \mathbf{y} &\ge \mathbf{x'}^T R \mathbf{y},\ \forall \mathbf{x}' \in \Delta_m; \\ \mathbf{x}^T C \mathbf{y} &\ge \mathbf{x}^T C \mathbf{y}',\ \forall \mathbf{y}' \in \Delta_n. \end{aligned}$$

Since the expected payoff is linear in the players' mixed strategies, the support of a player's strategy in a Nash equilibrium should include only those pure strategies that maximize her payoff given the other player's mixed strategy. Conversely, any strategy whose support satisfies the above condition must be a Nash equilibrium. This leads to a pair of alternative conditions defining a Nash equilibrium:

$$\begin{aligned} &\forall i \quad \text{s.t. } x_i > 0, \quad i \in \arg\max_k \{\mathbf{e_k}^T R \mathbf{y}\}; \\ &\forall j \quad \text{s.t. } y_j > 0, \quad j \in \arg\max_\ell \{\mathbf{x}^T C \mathbf{e}_\ell\}, \end{aligned}$$

where $\mathbf{e_k} \in \Delta_m, \mathbf{e}_\ell \in \Delta_n$ represent the vectors of appropriate dimensions with a single 1 at the coordinate k and ℓ respectively; these correspond to the pure strategies k and ℓ respectively for the row and column players.

Another pair of equivalent conditions asserts that a player cannot improve her payoff by unilaterally switching to a pure strategy, i.e.

$$\begin{aligned} \mathbf{x}^T R \mathbf{y} &\ge \mathbf{e_i}^T R \mathbf{y},\ \forall i \in [m] \\ \mathbf{x}^T C \mathbf{y} &\ge \mathbf{x}^T C \mathbf{e_j},\ \forall j \in [n], \end{aligned}$$

where $\mathbf{e_i}$ and $\mathbf{e_j}$ respectively represent the row player's pure strategy i and the column player's pure strategy j.

Examples. We give a few classical examples of two-player games. Each is described by a table whose (i, j)-th entry has a pair of values, the first corresponding to $R[i, j]$ and the second to $C[i, j]$.

Rock Paper Scissors (RPS)

This is the familiar school-yard game with the same name.

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

It can be checked that his game has a unique Nash equilibrium, namely when both players randomize uniformly over their strategies: $\mathbf{x} = \mathbf{y} = \langle \frac{1}{3}, \frac{1}{3}, \frac{1}{3} \rangle$.

Prisoner's Dilemma

Two suspects are arrested by the police. The police have insufficient evidence for a conviction, and, having separated both prisoners, visit each of them to offer the same deal. If one testifies for the prosecution against the other and the other remains silent, the betrayer goes free and the silent accomplice receives the full 10-year sentence. If both remain silent, both prisoners are sentenced to only six months in jail for a minor charge. If each betrays the other, each receives a five-year sentence. Each prisoner must choose to betray the other or to remain silent. Each one is assured that the other would not know about the betrayal before the end of the investigation. How should the prisoners act?

	Silent	Betray
Silent	$-\frac{1}{2}, -\frac{1}{2}$	-10, 0
Betray	0, -10	-5, -5

This game has a pure Nash equilibrium where both players "Betray". Observe that this is not the *social optimum*, that is the pair of strategies optimizing the sum of the players' payoffs, since both players stand to gain if they simultaneously switched to "Silent". This exposes a key feature of Nash equilibria—that they may fail to optimize social welfare and are stable only under unilateral deviations by the players.

Battle of the Sexes

A pair of friends with different interests need to decide where to spend their afternoon. They make simultaneous suggestions to each other and will only go somewhere if their suggestions match. One prefers football and the other prefers theater. Their payoffs are as follows.

	Theater!	Football fine
Theater fine	1, 5	0, 0
Football!	0, 0	5, 1

This game has multiple Nash equilibria. In particular, both the squares with non-zero payoffs represent Nash equilibria. In addition, $\mathbf{x} = \langle \frac{1}{6}, \frac{5}{6} \rangle$ and $\mathbf{y} = \langle \frac{5}{6}, \frac{1}{6} \rangle$ is a Nash equilibrium. Thus, this game has 3 Nash equilibria. In fact, as we will establish later this semester, *every "non-degenerate game" must have an odd number of Nash equilibria!*

Two-player Zero-sum Games. A two-player game is said to be *zero-sum* if $R + C = 0$, i.e. $R[i, j] + C[i, j] = 0$ for all $i \in [m]$ and $j \in [n]$. For example, the RPS game described above is zero-sum. We give another less symmetrical game below.

Presidential Elections (adopted from [4])

Two candidates for presidency have two possible strategies on which to focus their campaigns. Depending on their decisions they may win the favor of a number of voters. The table below describes their gains in millions of voters.

	Morality	Tax-cuts
Economy	3, -3	-1, 1
Society	-2, 2	1, -1

Clearly, this is also a zero-sum game. Suppose now that the row player announces a strategy of $\langle \frac{1}{2}, \frac{1}{2} \rangle$ in advance. This is presumably a suboptimal way to play the game, since it exposes the row player's strategy to the column player who can now react optimally. Given the row player's announced strategy, the column player will choose the strategy that maximizes her payoff given the row player's declared strategy; in this case, she will choose "Tax-cuts". More generally, if the row player announces strategy $\langle x_1, x_2 \rangle$, then the column player has the following expected payoffs:

$$\begin{aligned} \mathbb{E}[\text{``Morality''}] &= -3x_1 + 2x_2 \\ \mathbb{E}[\text{``Tax} - \text{cuts''}] &= x_1 - x_2. \end{aligned}$$

Hence, the payoff that the row player will get after reacting optimally to the row player's announced strategy is $\max(-3x_1 + 2x_2, x_1 - x_2)$. Since this is a zero-sum game, the resulting row player's payoff is $-\max(-3x_1 + 2x_2, x_1 - x_2) \equiv \min(3x_1 - 2x_2, -x_1 + x_2)$. So if the row player is forced to announce his strategy in advance he will choose

$$(x_1, x_2) \in \arg\max_{(x_1,x_2)} \min(3x_1 - 2x_2, -x_1 + x_2).$$

In other words, the row player's optimal announced strategy (x_1, x_2) can be computed via the following linear program, whose value will be equal to the row player's payoff after the column player's optimal response to (x_1, x_2).

$$\max z$$

$$\begin{aligned} \text{s.t.} \quad 3x_1 - 2x_2 &\geq z \\ -x_1 + x_2 &\geq z \\ x_1 + x_2 &= 1 \\ x_1, x_2 &\geq 0. \end{aligned}$$

Solving this LP yields $x_1 = \frac{3}{7}, x_2 = \frac{4}{7}$ and $z = \frac{1}{7}$. Conversely, if the column player were forced to announce her strategy in advance, she would aim to solve the following LP.

$$\max w$$

$$\begin{aligned} \text{s.t.} \quad -3y_1 + y_2 &\geq w \\ 2y_1 - y_2 &\geq w \\ y_1 + y_2 &= 1 \\ y_1, y_2 &\geq 0, \end{aligned}$$

which yields the following solution: $y_1 = \frac{2}{7}, y_2 = \frac{5}{7}$ and $w = -\frac{1}{7}$.

Here is a remarkable observation. If the row player plays $\mathbf{x} = \langle \frac{3}{7}, \frac{4}{7} \rangle$ and the column player plays $\mathbf{y} = \langle \frac{2}{7}, \frac{5}{7} \rangle$ then

- The payoff of the row player is $1/7$ and the payoff of the column payer is $-1/7$; in particular, the sum of their payoffs is zero.

- Given strategy $\mathbf{x} = \langle \frac{3}{7}, \frac{4}{7} \rangle$ for the row player, the column player cannot obtain a payoff greater than $-\frac{1}{7}$ since the row player gets a payoff of at least $\frac{1}{7}$ irrespective of the column player's strategy (by the definition of the first LP) and the game is zero-sum;
- Similarly, given $\mathbf{y} = \langle \frac{2}{7}, \frac{5}{7} \rangle$ for the column player, the row player cannot obtain a payoff greater than $\frac{1}{7}$ since the column player gets a payoff of at least $-\frac{1}{7}$ irrespective of the row player's strategy.

Hence, the pair of strategies $\mathbf{x} = \langle \frac{3}{7}, \frac{4}{7} \rangle$ and $\mathbf{y} = \langle \frac{2}{7}, \frac{5}{7} \rangle$ is a Nash equilibrium! This were somewhat unexpected since the two LPs were formulated independently by the two players without any care/guessing about what their opponent might do. In particular, it seemed a priori suboptimal/too pessimistic for a player to announce her strategy and let the other player respond. Moreover, there was no reason to expect that the pessimistic strategies for the two players are in fact best responses to each other. As we show below, this coincidence happens for a deep reason, it is true in all two-player zero-sum games, and is a ramification of the strong LP duality.

2 Generalization

2.1 LP Formulation

Let us now study general two-player zero-sum games. Recall that a two-player game is zero-sum iff the payoff matrices (R, C) satisfy $R + C = 0$. As in the previous section, let us write the LP for the row player if he is forced to announce his strategy in advance.

$$\begin{aligned} & \max \ z \\ \text{s.t.} \quad & \boldsymbol{x}^T R \geq z \mathbf{1}^T \\ & \boldsymbol{x}^T \mathbf{1} = 1 \\ & \forall i, \ x_i \geq 0. \end{aligned} \qquad LP(1)$$

Notice that the maximum of z is actually the same as

$$\max_{\boldsymbol{x}} \min_{\boldsymbol{y}} \boldsymbol{x}^T R \boldsymbol{y}.$$

Now let us consider the dual of $LP(1)$:

$$\begin{aligned} & \min \ z' \\ \text{s.t.} \quad & -\boldsymbol{y}^T R^T + z' \mathbf{1}^T \geq \mathbf{0} \\ & \boldsymbol{y}^T \mathbf{1} = 1 \\ & \forall j, \ y_j \geq 0. \end{aligned} \qquad LP(2)$$

Let $z'' = -z'$. Using the fact that $C = -R$, we can change $LP(2)$ into another equivalent LP.

$$\begin{aligned} & \max \ z'' \\ \text{s.t.} \quad & C \boldsymbol{y} \geq z'' \mathbf{1} \\ & \boldsymbol{y}^T \mathbf{1} = 1 \\ & \forall j \ y_j \geq 0. \end{aligned} \qquad LP(3)$$

It is not hard to see that the maximum of z'' is equal to

$$\max_{\boldsymbol{y}} \min_{\boldsymbol{x}} \boldsymbol{x}^T C \boldsymbol{y} = -\min_{\boldsymbol{y}} \max_{\boldsymbol{x}} \boldsymbol{x}^T R \boldsymbol{y}.$$

Here is the interesting fact: $LP(3)$ is nothing but the linear program that the column player would solve if she were forced to announce her strategy in advance! Using the strong duality of linear programming we are going to show next that the solutions to $LP(1)$ and $LP(3)$ constitute a Nash equilibrium of the game.

2.2 LP Strong Duality

Since $LP(2)$ is the dual of $LP(1)$, LP strong duality implies the following: If $(\boldsymbol{x}, z)$ is optimal for $LP(1)$ and $(\boldsymbol{y}, z')$ is optimal for $LP(2)$, then $z = z'$. On the other hand, given our construction of $LP(3)$ it follows that, if $(\boldsymbol{y}, z')$ is optimal for $LP(2)$, then $(\boldsymbol{y}, -z')$ is optimal for $LP(3)$. Thus, if $(\boldsymbol{x}, z)$ is optimal for $LP(1)$ and $(\boldsymbol{y}, z'')$ is optimal for $LP(3)$, then $z = -z''$.

We show next that the solutions of these LPs actually give us a Nash equilibrium.

Theorem 1. *If $(\boldsymbol{x}, z)$ is optimal for $LP(1)$, and $(\boldsymbol{y}, z'')$ is optimal for $LP(3)$, then $(\boldsymbol{x}, \boldsymbol{y})$ is a Nash equilibrium of (R, C). Moreover, the payoffs of the row/column player in this Nash equilibrium are z and $z'' = -z$ respectively.*

Proof: Since $(\boldsymbol{x}, z)$ is feasible for $LP(1)$, and $(\boldsymbol{y}, z'')$ is feasible for $LP(3)$:

$$\boldsymbol{x}^T R \geq z\mathbf{1} \Rightarrow \boldsymbol{x}^T R\boldsymbol{y} \geq z \tag{1}$$

$$\begin{aligned} C\boldsymbol{y} \geq z''\mathbf{1} &\Rightarrow \boldsymbol{x'}^T C\boldsymbol{y} \geq z'', \forall \boldsymbol{x'} \\ &\Rightarrow \boldsymbol{x'}^T R\boldsymbol{y} \leq -z'', \forall \boldsymbol{x'}. \end{aligned} \tag{2}$$

From equation (2) we know that, if the column player is playing $\boldsymbol{y}$, the row player's payoff is at most $-z'' \equiv z$ (since we argued above that $z'' = -z$ from strong LP duality). Moreover, equation (1) implies that if the row player plays $\boldsymbol{x}$ against $\boldsymbol{y}$, her payoff will be at least z; in fact, given (2) it is exactly z. Thus, $\boldsymbol{x}$ is a best response for the row player against strategy $\boldsymbol{y}$ of the column player, giving payoff of z to the row player.

Using a similar argument, we can show that $\boldsymbol{y}$ is a best response for the column player against strategy $\boldsymbol{x}$ of the row player, giving payoff z'' to the column player. Hence, $(\boldsymbol{x}, \boldsymbol{y})$ is a Nash equilibrium in which the players' payoffs are z and $z'' = -z$ respectively. $\square$

We give a couple of corollaries of the above theorem.

Corollary 1. *There exists a Nash equilibrium in every two-player zero-sum game.*

Corollary 2 (The Minmax Theorem)**.**

$$\max_{\boldsymbol{x}} \min_{\boldsymbol{y}} \boldsymbol{x}^T R\boldsymbol{y} = \min_{\boldsymbol{y}} \max_{\boldsymbol{x}} \boldsymbol{x}^T R\boldsymbol{y}.$$

Next, we argue that every Nash equilibrium of the two-player zero-sum game can be found by our approach.

Theorem 2. *If $(\boldsymbol{x}, \boldsymbol{y})$ is a Nash equilibrium of (R, C), then $(\boldsymbol{x}, \boldsymbol{x}^T R\boldsymbol{y})$ is an optimal solution of $LP(1)$, and $(\boldsymbol{y}, -\boldsymbol{x}^T C\boldsymbol{y})$ is an optimal solution of $LP(2)$.*

Proof: Let $z = \boldsymbol{x}^T R\boldsymbol{y} = -\boldsymbol{x}^T C\boldsymbol{y}$. Since $(\boldsymbol{x}, \boldsymbol{y})$ is a Nash equilibrium.

$$\begin{array}{rll} & \boldsymbol{x}^T C\boldsymbol{y} & \geq \boldsymbol{x}^T C\boldsymbol{e_j} \quad \forall j \\ \Rightarrow & -\boldsymbol{x}^T R\boldsymbol{y} & \geq -\boldsymbol{x}^T R\boldsymbol{e_j} \quad \forall j \\ \Rightarrow & \boldsymbol{x}^T R & \geq z\mathbf{1}. \end{array}$$

So we have showed that $(\boldsymbol{x}, z)$ is a feasible solution for $LP(1)$; similarly we can argue that $(\boldsymbol{y}, z)$ is a feasible solution for $LP(2)$. Since $LP(2)$ is the dual of $LP(1)$, by weak duality, we know that the value achieved by any feasible solution of $LP(2)$ upper bounds the value achieved by any feasible solution of $LP(1)$. But, these values are equal here. So, $(\boldsymbol{x}, z)$ and $(\boldsymbol{y}, z)$ must be optimal solutions for $LP(1)$ and $LP(2)$ respectively. $\square$

Since all Nash equilibria define a pair of optimal solutions to $LP(1)$ and $LP(2)$ and all pairs of optimal solutions to $LP(1)$ and $LP(2)$ define a Nash equilibrium, we obtain the following important corollaries of Theorems 1 and 2.

Corollary 3 (Convexity of Optimal Strategies)**.** *The set*

$$\{\boldsymbol{x} \mid \textit{there exists } \boldsymbol{y} \textit{ such that } (\boldsymbol{x}, \boldsymbol{y}) \textit{ is a Nash equilibrium of } (R, C)\}$$

of equilibrium strategies for the row player in zero-sum game is convex. Ditto for the column player.

Corollary 4 (Anything Goes)**.** *If $\boldsymbol{x}$ is any equilibrium strategy for the row player in a zero-sum game (R, C) and $\boldsymbol{y}$ is any equilibrium strategy for the column player, then $(\boldsymbol{x}, \boldsymbol{y})$ is a Nash equilibrium of the game.*

Corollary 5 (Convexity of Equilibrium Set)**.** *The equilibrium set*

$$\{(\boldsymbol{x}, \boldsymbol{y}) \mid (\boldsymbol{x}, \boldsymbol{y}) \ \textit{ is a Nash equilibrium of } (R, C)\}$$

of a zero-sum game is convex.

Corollary 6 (Uniqueness of Payoffs)**.** *The payoff of the row player is equal in all Nash equilibria of a zero-sum game. Ditto for the column player.*

Inspired by this last corollary, we define the *value of a zero-sum game*, to be the payoff of the row player in any (all) Nash equilibria of the game.

Definition 1 (Value of a Zero-Sum Game)**.** *If (R, C) is zero-sum game, then the* value of the game *is the unique payoff of the row player in all Nash equilibria of the game.*

3 Homework

We conclude with a homework problem inspired by our discussion of zero-sum games.

Adjacent Zero-Sum Games (2 points): Players 1 and 2 are playing a zero-sum game with payoff matrices $(A^{1,2}, A^{2,1})$; at the same time players 1 and 3 are playing another zero-sum game with payoff matrices $(A^{1,3}, A^{3,1})$. Instead of using different strategies in the two games, player 1 must use the same strategy in both games.

1. Show that there exists a Nash equilibrium in this game.
2. Give an algorithm that can compute a Nash equilibrium efficiently.

4 Historical Remarks

The study of games from a mathematical perspective goes back to the work of mathematician Émile Borel. The existence of a Nash equilibrium in two-player zero-sum games was established by John von Neumann in 1928, in a paper [6] that arguably gave birth to the field of Game Theory. Two decades after von Neumann's result it became understood in a discussion between von Neumann and Danzig—the father of mathematical programming—that the existence of an equilibrium in zero-sum games is implied by Strong Linear Programming duality (see, e.g., [2], [3]). Hence, as was established another three decades later by Khachiyan [5], finding an equilibrium in zero-sum games is computationally tractable. Dantzig also proposed a reduction from a generic Linear Programming problem to the computation of a min-max strategy in a zero-sum game, noting that the reduction does not necessarily work. Despite the lack of a complete proof, it has since been often quoted that the Min-Max theorem and the Strong Duality Theorem of linear programming are equivalent. In reality, it was only recently that Ilan Adler (himself a student of Danzig) provided a complete reduction from Linear Programming to Zero-sum Games [1].

References

[1] Ilan Adler. On the Equivalence of Linear Programming Problems and Zero-Sum Games. *Optimization Online,* 2010.

[2] Donald J. Albers and Constance Reid. An Interview with George B. Dantzig: The Father of Linear Programming. *The College Mathematics Journal,* 17(4):293314, 1986.

[3] George B. Dantzig. *Linear Programming and Extensions.* Princeton University Press, 1963.

[4] Sanjoy Dasgupta, Christos Papadimitriou and Umesh Vazirani. *Algorithms.* McGraw-Hill, 2006.

[5] Leonid G. Khachiyan. A Polynomial Algorithm in Linear Programming. *Soviet Mathematics Doklady,* 20(1):191194, 1979.

[6] John von Neumann. Zur Theorie der Gesellshaftsspiele. *Mathematische Annalen,* 100:295320, 1928.