

1 Overview

In our previous lecture we explored the concept of duality which is the cornerstone of optimization theory. The goal of today's lecture is to explore the consequences of duality in two fields where optimization is ubiquitous: Game Theory and Learning Theory, and explore the interesting connections between these areas.

2 A Primer in Game Theory

We will focus on two-player games which constitute a nice introduction to the key concepts of Game Theory. A two-player game is defined by:

- a set of two players
- set of strategies $M = \{1, \dots, m\}$ and $N = \{1, \dots, n\}$ for player 1 and player 2 respectively.
- payoff matrices $A \in \mathbb{R}^{m \times n}$ and $B \in \mathbb{R}^{m \times n}$ for player 1 and player 2 respectively.

The signification of the payoff matrices is the following: when player 1 plays action $i \in \{1, \dots, m\}$ and player 2 plays action $j \in \{1, \dots, n\}$, A_{ij} is the gain of player 1 and B_{ij} is the gain of player 2.

Example: battle of the sexes. A famous example of a two-player game is *battle of the sexes*. This game models a situation where a couple is trying to agree on a program for the evening. The sets of strategies are the same for the husband and the wife: watch a football game or go to the opera. The couple would prefer to do the same activity, but the husband would prefer to watch the football game while the wife would prefer to go to the opera. The payoff matrices of the two players is depicted in Figure 1.

When analyzing a game, one would like to know which combination of strategies are likely to be chosen by the players. Assuming that the players are rational, *i.e.* utility maximizers, the notion of *Nash equilibria* naturally emerges.

Definition. A strategy profile $(i, j) \in M \times N$ is called a Nash equilibrium if the strategy played by each player is optimal given the strategy chosen by the other player:

- $A_{ij} \geq A_{kj}$ for all $k \in M$.
- $B_{ij} \geq B_{ik}$ for all $k \in N$.

	Opera	Football
Opera	(3, 2)	(1, 1)
Football	(0, 0)	(2, 3)

Figure 1: Battle of the sexes game. The rows represent the strategies of the wife and the columns represent the strategies of the husband. In this compact representation, we combine both payoff matrices into one table. The upper left cell of this table should be read like this: *when the wife and the husband both agree to go to the opera, the gain of the wife is 3 and the gain of the husband is 2.*

In the battle of sexes game introduced above, there are two Nash equilibria: (Opera, Opera) and (Football, Football). (Opera, Football), for example, is not a Nash equilibrium since the wife could increase her gain from 1 to 2 by choosing Football instead of Opera.

Example: Prisoner’s dilemma. Another famous example of two-player games is the *prisoners’ dilemma*. In this problem, two prisoners are in solitary confinement. The police is convinced that both prisoners are guilty but does not have enough evidence to convict them to a long prison sentence. Thus, the police interrogates the prisoners separately. Each prisoner can either stay silent (not reveal anything to the police) or betray the other. If only one of the prisoners betrays the other, then he will be set free and the other will be convicted to a long prison sentence. However, if they betray each other, both prisoners will serve a long prison sentence. Possible payoffs modeling this situation are depicted in the matrices below.

	Silent	Betray
Silent	(-1, -1)	(-10, 0)
Betray	(0, -10)	(-8, -8)

A close inspection of this table shows that there is only one Nash equilibrium: (Betray, Betray). Indeed, in all other cases, a silent prisoner can always increase his gain (or reduce his loss) by choosing to betray instead of staying silent.

In this case, the Nash equilibrium is not optimal for the prisoners: they could be both better off by staying silent. This suboptimality is a consequence of the nature of the game: both prisoners choose their action without knowing what the other prisoner does.

3 Mixed strategies, zero-sum games and the minimax theorem

When studying games, a more general notion of strategies can also be defined: *mixed strategy*. A mixed strategy is an assignment of a probability to each (pure) strategy so that. A mixed strategy allows the player to choose a pure strategy at random.

Formally, a mixed strategy for player 1 over the set of pure strategies $M = \{1, \dots, m\}$ is a probability vector $\mathbf{x} \in \Delta_M = \{u \in [0, 1]^m : \sum_{i=1}^m u_i = 1\}$. A mixed strategy for player 2 is defined in a similar way. Note that a pure strategy can be seen as a degenerate case of mixed strategy where a single probability is set to 1 and all the others are set to 0.

The notion of gain and Nash equilibrium can be extended to this generalized notion of strategy. If player 1 and player 2 have mixed strategies $x \in \Delta_M$ and $y \in \Delta_N$ respectively, their gains are defined as the expected gain when they both play selecting pure strategies at random according to x and y . Formally, the gain of player 1 will be:

$$\sum_{(i,j) \in M \times N} x_i y_j A_{ij} = \mathbf{x}^\top \mathbf{A} \mathbf{y}$$

and similarly for player 2 by replacing A with B .

The notion of Nash equilibria is also extended using this expected gain.

Definition. A profile of mixed strategies $\mathbf{x} \in \Delta_M$ and $\mathbf{y} \in \Delta_N$ is a mixed Nash equilibrium iff:

- $\mathbf{x}^\top \mathbf{A} \mathbf{y} \geq \mathbf{u}^\top \mathbf{A} \mathbf{y}$ for all $\mathbf{u} \in \Delta_M$.
- $\mathbf{x}^\top \mathbf{B} \mathbf{y} \geq \mathbf{x}^\top \mathbf{B} \mathbf{v}$ for all $\mathbf{v} \in \Delta_N$.

Example: Penalty-kick game. The *penalty kick* game models the situation in soccer where a player is about to kick a penalty: usually the goalie stands at the middle of the goal line and the player choose to aim either to the left or to the right of the goal. Because the player stands very close to the goal, the goalie has to decide on a direction in advance and starts diving in this direction even before the player hits the ball. If the goalie dives in the direction chosen by the player, he saves the shot, otherwise he fails. Possible payoffs modeling this situation could be:

	Left	Right
Left	(+1, -1)	(-1, +1)
Right	(-1, +1)	(+1, -1)

where the rows represent the dive direction chosen by the goalie and the columns represent the kick direction chosen by the player. Note that this game has no pure Nash equilibria: in any situation the player with gain -1 can increase his gain by changing his strategy. Intuitively, this simply expresses that there is always a better strategy *in retrospect* for the losing player. However, this game has one mixed Nash equilibrium: $(1/2, 1/2)$, $(1/2, 1/2)$. When both players choose a direction uniformly at random, this leads to an expected gain of 0 for both of them. Any other choice of probabilities would lead to an unbalancedness either to the left or to the right which could be exploited by the other player.

The fact that the penalty kick game has a mixed Nash equilibrium is not surprising. It is in fact a consequence of a general theorem. Remark that in the penalty kick game, the gain of the kicker is exactly the opposite of the gain of the goalie. More generally we have:

Definition. A two-player game is a zero-sum game iff $B = -A$ with A and B the payoffs matrices of the players.

Zero-sum game are sometimes called *strictly competitive* as the gain of a player is exactly the loss of the other(s) player(s): cooperation is useless in these games. We will soon prove the existence of equilibria of zero-sum games. But before doing so, we will prove the minimax theorem.

This theorem is a simple consequence of the following theorem which is itself a simple consequence of the strong duality theorem.

3.1 Von Neumann's minimax theorem

Theorem 1. For any matrix $A \in \mathbb{R}^{|M| \times |N|}$:

$$\max_{\mathbf{x} \in \Delta_M} \min_{\mathbf{y} \in \Delta_N} \mathbf{x}^\top \mathbf{A} \mathbf{y} = \min_{\mathbf{y} \in \Delta_N} \max_{\mathbf{x} \in \Delta_M} \mathbf{x}^\top \mathbf{A} \mathbf{y}$$

Proof. First observe that for a given $\mathbf{x} \in \Delta_M$:

$$\min_{\mathbf{y} \in \Delta_N} \mathbf{x}^\top \mathbf{A} \mathbf{y} = \min_{j \in N} (\mathbf{x}^\top \mathbf{A})_j$$

hence, the left-hand side of the minimax theorem can be written as the following LP:

$$\begin{aligned} \max \quad & t \\ \text{s.t.} \quad & t \leq (\mathbf{x}^\top \mathbf{A})_j, \quad j \in N \\ & \mathbf{x} \geq 0 \\ & \sum_{i \in M} x_i = 1 \end{aligned}$$

Similarly, the right-hand side can be written as the following LP:

$$\begin{aligned} \min \quad & s \\ \text{s.t.} \quad & s \geq (\mathbf{A} \mathbf{y})_i, \quad i \in M \\ & \mathbf{y} \geq 0 \\ & \sum_{j \in N} y_j = 1 \end{aligned}$$

and both LPs are dual of each other. The claim now follows from the strong duality theorem. \square

Let us now see how to prove the existence of mixed Nash equilibria in zero-sum two-player games:

Theorem 2. For every two-player zero-sum game, there exists a mixed Nash equilibrium.

Proof. Let us consider a zero-sum two-player game with A the pay-off matrix for player A . Let us write $v = \max_{\mathbf{x} \in \Delta_M} \min_{\mathbf{y} \in \Delta_N} \mathbf{x}^\top \mathbf{A} \mathbf{y} = \min_{\mathbf{y} \in \Delta_N} \max_{\mathbf{x} \in \Delta_M} \mathbf{x}^\top \mathbf{A} \mathbf{y}$ and consider $\mathbf{x}^* \in \arg \max_{\mathbf{x} \in \Delta_M} \min_{\mathbf{y} \in \Delta_N} \mathbf{x}^\top \mathbf{A} \mathbf{y}$ and $\mathbf{y}^* \in \arg \min_{\mathbf{y} \in \Delta_N} \max_{\mathbf{x} \in \Delta_M} \mathbf{x}^\top \mathbf{A} \mathbf{y}$. Then we claim that $(\mathbf{x}^*, \mathbf{y}^*)$ is a mixed Nash equilibrium. Indeed, it follows from the definition of \mathbf{x}^* and \mathbf{y}^* that:

- $\mathbf{x}^\top \mathbf{A} \mathbf{y}^* \leq v \leq \mathbf{x}^{*\top} \mathbf{A} \mathbf{y}^*$ for all $\mathbf{x} \in \Delta_M$.
- $\mathbf{x}^{*\top} \mathbf{A} \mathbf{y} \geq v \geq \mathbf{x}^{*\top} \mathbf{A} \mathbf{y}^*$, *i.e.* $\mathbf{x}^{*\top} \mathbf{B} \mathbf{y} \leq \mathbf{x}^{*\top} \mathbf{B} \mathbf{y}^*$ for all $\mathbf{y} \in \Delta_N$.

which concludes the proof. □