16 Non-Cooperative Games

The second part of the course will be concerned with situations in which multiple self-interested entities, or *agents*, operate in the same environment. *Game theory* provides mathematical models, so-called games, for studying these types of situations. We focus for now on *non-cooperative games*, where agents independently optimize different objectives and outcomes must be self-enforcing. In a later lecture we also consider *cooperative games*, which focus on conditions under which cooperation among subsets of the agents can be sustained.

16.1 Games and Solutions

The central object of study in non-cooperative game theory are normal-form games. A *normal-form game* is a tuple $\Gamma = (N, (A_i)_{i \in N}, (p_i)_{i \in N})$ where N is a finite set of *players*, and for each player $i \in N$, A_i is a non-empty and finite set of *actions* available to i and $p_i : (\times_{i \in N} A_i) \to \mathbb{R}$ is a function mapping each action profile, i.e., each combination of actions, to a real-valued *payoff* for i. Unless noted otherwise, the results we consider are invariant under positive affine transformations, and payoffs and their relative intensities will not be comparable across players.

More complicated games where players move sequentially and base their decisions on their and others' earlier moves can also be represented as normal-form games, by encoding every possible course of action in the former by an action of the latter. It should be noted, however, that this generally leads to a large increase in the number of actions.

We henceforth restrict our attention to two-player games, but note that most concepts and results extend in a straightforward way to games with more than two players. A two-player game with m actions for player 1 and n actions for player 2 can be represented by a pair of matrices $P, Q \in \mathbb{R}^{m \times n}$, where p_{ij} and q_{ij} are the payoffs of players 1 and 2 when player 1 plays action i and player 2 plays action j. Two-player games are therefore sometimes referred to as bimatrix games, and players 1 and 2 as the row and column player, respectively.

Assume that players can choose their actions randomly and denote the set of possible *strategies* of the two players by X and Y, respectively, i.e., $X = \{x \in \mathbb{R}^m_{\geqslant 0} : \sum_{i=1}^m x_i = 1\}$ and $Y = \{y \in \mathbb{R}^n_{\geqslant 0} : \sum_{i=1}^n y_i = 1\}$. A *pure strategy* is a strategy that chooses some action with probability one, and we make no distinction between pure strategies and the corresponding actions. A profile $(x, y) \in X \times Y$ of strategies induces a lottery over outcomes, and we write $p(x, y) = x^T P y$ and $q(x, y) = x^T Q y$ for the expected payoff of the two players in this lottery.

Consider for example the well-known prisoner's dilemma, involving two suspects

	S	T
S	(2, 2)	(0, 3)
T	(3, 0)	(1, 1)

Figure 16.1: Representation of the prisoner's dilemma as a normal-form game. The matrices P and Q are displayed as a single matrix with entries (p_{ij}, q_{ij}), and players 1 and 2 respectively choose a row and a column of this matrix. Action S corresponds to remaining silent, action T to testifying.

	C	D
C	(2, 2)	(1, 3)
D	(3, 1)	(0, 0)

Figure 16.2: The game of chicken, where players can chicken out or dare

accused of a crime who are being interrogated separately. If both remain silent, they walk free after spending a few weeks in pretrial detention. If one of them testifies against the other and the other remains silent, the former is released immediately while the latter is sentenced to ten years in jail. If both suspects testify, each of them receives a five-year sentence. A representation of this situation as a two-player normal-form game is shown in Figure 16.1.

It is easy to see what the players in this game should do, because action T yields a strictly larger payoff than action S for *every* action of the respective other player. More generally, for two strategies $x, x' \in X$ of the row player, x is said to *(strictly) dominate* x' if for every strategy $y \in Y$ of the column player, $p(x, y) > p(x', y)$. Dominance for the column player is defined analogously. Strategy profile (T, T) in the prisoner's dilemma is what is called a *dominant strategy equilibrium*, a profile of strategies that dominate every other strategy of the respective player. The source of the dilemma is that outcome resulting from (T, T) is strictly worse for both players than the outcome resulting from (S, S). More generally, an outcome that is weakly preferred to another outcome by all players, and strictly preferred by at least one player is said to *Pareto dominate* that outcome. An outcome that is Pareto dominated is clearly undesirable.

In the absence of dominant strategies, it is less obvious how players should proceed. Consider for example the game of chicken illustrated in Figure 16.2. This game has its origins in a situation where two cars drive towards each other on a collision course. Unless one of the drivers yields, both may die in a crash. If one of them yields while the other goes straight, however, the former will be called a "chicken," or coward. It is easily verified that this game does not have any dominant strategies.

The most cautious choice in a situation like this would be to ignore that the other player is self-interested and choose a strategy that maximizes the payoff in the worst case, taken over all of the other player's strategies. A strategy of this type is known as a *maximin strategy*, and the payoff thus obtained as the player's *security level*. It is easy to see that it suffices to maximize the minimum payoff over all *pure* strategies of the other player, i.e., to choose x such that $\min_{j\in\{1,\ldots,n\}}\sum_{i=1}^{m} x_i p_{ij}$ is as large as possible. Maximization of this minimum can be achieved by maximizing a lower bound that holds for all $j = 1, \ldots, n$, so a maximin strategy and the security level for the row player can be found as a solution of the following linear program with variables $v \in \mathbb{R}$ and $x \in \mathbb{R}^m$:

$$\begin{aligned} \text{maximize} \quad & v \\ \text{subject to} \quad & \sum_{i=1}^{m} x_i p_{ij} \geqslant v \quad \text{for } j = 1, \ldots, n \\ & \sum_{i=1}^{m} x_i = 1 \\ & x \geqslant 0. \end{aligned} \tag{16.1}$$

The unique maximin strategy in the game of chicken is to yield, for a security level of 1. This also illustrates that a maximin strategy need not be optimal: assuming that the other player yields, the best response is in fact to go straight. Formally, strategy $x \in X$ of the row player is a *best response* to strategy $y \in Y$ of the column player if for all $x' \in X$, $p(x, y) \geqslant p(x', y)$. The concept of a best response for the column player is defined analogously.

16.2 The Minimax Theorem

In the special case when the interests of the two players are diametrically opposed, maximin strategies turn out to be optimal in a very strong sense. A two-player game is called *zero-sum game* if $q_{ij} = -p_{ij}$ for all $i = 1, \ldots, m$ and $j = 1, \ldots, n$. A game of this type is sometimes called a matrix game, because it can be represented just by the matrix P containing the payoffs of the row player. Assuming invariance of utilities under positive affine transformations, results for zero-sum games in fact apply to the larger class of *constant-sum* games, in which the payoffs of the two players always sum up to the same constant. For games with more than two players, these properties are far less interesting, as one can always add an extra player who "absorbs" the payoffs of the others.

Theorem 16.1 (von Neumann, 1928). *Let* $P \in \mathbb{R}^{m\times n}$, $X = \{x \in \mathbb{R}^m_{\geqslant 0} : \sum_{i=1}^{m} x_i = 1\}$, $Y = \{y \in \mathbb{R}^n_{\geqslant 0} : \sum_{i=1}^{n} y_i = 1\}$. *Then,*

$$\max_{x\in X}\min_{y\in Y} p(x, y) = \min_{y\in Y}\max_{x\in X} p(x, y).$$

Proof. Again consider the linear program (16.1), and recall that the optimal solution of this linear program is equal to $\max_{x\in X}\min_{y\in Y} p(x, y)$. By adding a slack variable

$z \in \mathbb{R}^n$ with $z \geqslant 0$ we obtain the Lagrangian

$$\begin{aligned} L(v, x, z, w, y) &= v + \sum_{j=1}^{n} y_j \Big(\sum_{i=1}^{m} x_i p_{ij} - z_j - v \Big) - w \Big(\sum_{i=1}^{m} x_i - 1 \Big) \\ &= \Big(1 - \sum_{j=1}^{n} y_j \Big) v + \sum_{i=1}^{m} \Big(\sum_{j=1}^{n} p_{ij} y_j - w \Big) x_i - \sum_{j=1}^{n} y_j z_j + w, \end{aligned}$$

where $w \in \mathbb{R}$ and $y \in \mathbb{R}^n$. The Lagrangian has a finite maximum for $v \in \mathbb{R}$ and $x \in \mathbb{R}^m$ with $x \geqslant 0$ if and only if $\sum_{j=1}^{n} y_j = 1$, $\sum_{j=1}^{n} p_{ij} y_j \leqslant w$ for $i = 1, \ldots, m$, and $y \geqslant 0$. The dual of (16.1) is therefore

$$\begin{aligned} \text{minimize} \quad & w \\ \text{subject to} \quad & \sum_{j=1}^{n} p_{ij} y_j \leqslant w \quad \text{for } i = 1, \ldots, m \\ & \sum_{j=1}^{n} y_j = 1 \\ & y \geqslant 0. \end{aligned}$$

It is easy to see that the optimal solution of the dual is $\min_{y \in Y} \max_{x \in X} p(x, y)$, and the theorem follows from strong duality. □

The number $\max_{x \in X} \min_{y \in Y} p(x, y) = \min_{y \in Y} \max_{x \in X} p(x, y)$ is also called the *value* of the matrix game with payoff matrix P.