Game theory, lecture 6

2 person general sum games

General sum games don't have as clean or satisfying a theory as zero sum games, but new phenomena arise. We have noncooperative and cooperative games. (In a 0-sum game there is no reason to cooperate.) The general set up is as follows:

We have 2 sets X and Y (move spaces for players I and II respectively) and two functions

$$U_1 : X \times Y \to \mathbb{R} ,$$

$$U_2 : X \times Y \to \mathbb{R} ,$$

representing the players' payoff if the moves made are x and y. We easily see that

$$U_1(x, y) = -U_2(x, y) \qquad \Leftrightarrow \qquad \text{0-sum game.}$$

A finite two-person game can be represented by a bimatrix, where the elements are pairs of numbers, for example:

$$\begin{pmatrix} (1,4) & (2,3) \\ (5,6) & (4,-1) \end{pmatrix}$$

We can also think of this as being 2 matrices A and B, one being the payoff matrix for player I (i.e. the first coordinate), and the other being the payoff matrix for player II.

How much can player I guarantee himself? The answer depends only on the matrix A:

$$v(A) = \max_{p_i} \min_j \sum_i p_i a_{i,j} \ .$$

This is called the safety level of player I. Similarly player II has a safety level which is

$$v(B^T) = \max_{q_j} \min_i \sum_j q_j b_{i,j} \ .$$

Playing according to the safety level (i.e. playing optimally for your payoff matrix) is not neccesserily good.

Example

$$\begin{pmatrix} (2,0) & (1,-100) \\ (1,0) & (2,-100) \end{pmatrix}$$

In this game player I could play $(\frac{1}{2}, \frac{1}{2})$ to be guaranteed $\frac{3}{2}$\$, but we could say that the "correct" way to play is to play 1 and get 2\$ (since player II probably will play 1). But if you feel your opponent is "out to get you", you should only play your optimal strategy for your matrix.

One replacement for an optimal strategy in a 0 sum game involves the notion of a strategic equilibrium, otherwise known as Nash equilibrium.

Definition A pair (i,j) is called a pure Nash equilibrium (PNE) if $a_{i,j} \geq a_{i',j}$ for all i' (in other words, i is a best response to j for player I), and $b_{i,j} \geq b_{i,j'}$ for all j' (j is a best response to i for player II).

Hence a pure Nash equilibrium is a pair of pure strategies such that there is no incentive for either player to unilaterally deviate. For 0-sum games these are simply saddle points. In particular PNE do not necessarily exist.

Example Example of a game with no PNE:

$$\begin{pmatrix} (1,-1) & (0,0) \\ (0,0) & (1,-1) \end{pmatrix}$$

Unlike 0-sum games, the payoff at different PNE may differ.

Example Example of a game which has different payoffs at different PNE:

$$\begin{pmatrix} (0,1) & (6,9) \\ (5,4) & (0,0) \end{pmatrix}$$

Definition A Nash Equilibrium (NE) for a 2-person general sum game (A,B) is a pair of probability vectors (p,q) on X and Y respectively such that

$$\sum_{i,j} p_i a_{i,j} q_j \geq \sum_{i,j} \tilde{p}_i a_{i,j} q_j \quad \forall \tilde{p}$$

$$\sum_{i,j} p_i b_{i,j} q_j \geq \sum_{i,j} p_i b_{i,j} \tilde{q}_j \quad \forall \tilde{q}$$

i.e. no player has an incentive to unilaterally deviate. In particular, being a NE is a property of the pair, i.e. "I plays a NE" is meaningless.

Theorem A Nash equilibrium always exists.

We will state and prove this for a k-person game.

Corollary Von Neumann's minimax theorem for 0-sum games.

Proof Given a 0-sum game, let (p, q) be a NE. Let

$$V := \sum_{i,j} a_{i,j} p_i q_j \ .$$

We need to show that

$$V_I \geq V \geq V_{II} \ .$$

(By symmetry, we only need to show $V_I \geq V$). By definition we have $V = pAq$, and $pAq \geq pAq'$ for all q' since (p, q) is a NE. In particular, $V_I \geq V$.

$\square$

Nash equilibrium for k-person games

Consider a game with k players. Each player has moves available to them. Player i's moves are given by a finite set S_i. Hence everyone's moves together is simply an element in $\prod_{i=1}^k S_i$. Let f represent the payoffs. So

$$f : \prod_{i=1}^{k} S_i \to \mathbb{R}^k \ ,$$

where

$$f_i(s), \ s \in \prod S_i \ ,$$

is the payoff to player i if the moves made are s. We would like to have a PNE meaning an element $\tilde{s} = (\tilde{s}_1, ..., \tilde{s}_k) \in S$ such that

$$f_i(\tilde{s}_1, ..., \tilde{s}_i, ..., \tilde{s}_k) \geq f_i(\tilde{s}_1, ..., \tilde{s}_{i-1}, s_i, \tilde{s}_{i+1}, ..., \tilde{s}_k) \quad \forall i \ \forall s_i \in S_i$$

But, of course, this might not exist. Denote $\mathcal{P}(S_i)$ to be the set of probability vectors on S_i. Let $\mathcal{P} = \prod_{i=1}^k \mathcal{P}(S_i)$, and extend f to a function $f : \mathcal{P} \to \mathbb{R}^k$

according to

$$f_i(\pi_1, ..., \pi_k) = \sum_{(s_1,...s_k)\in S} f_i(s_1, ..., s_k) \prod_{1=1}^{k} \pi_i(s_i) \ , \qquad i = 1, ...k.$$

Hence f_i is the expected payoff to player i, when the players use mixed strategies $\pi_1, ..., \pi_k$.

Definition A NE for a k-person game with finite move spaces $S_1, ...S_k$ and payoff functions $f_i : \prod S_i \to \mathbb{R}^k$ is an element $(\tilde{\pi}_1, ..., \tilde{\pi}_k) \in \mathcal{P}$ such that for all i:

$$f_i(\tilde{\pi}_1, ..., \tilde{\pi}_i, ..., \tilde{\pi}_k) \geq f_i(\tilde{\pi}_1, ..., \tilde{\pi}_{i-1}, \pi_i, \tilde{\pi}_{i+1}, ..., \tilde{\pi}_k) \quad \forall \pi_i \in \mathcal{P}(S_i)$$

i.e. no player has an incentive to unilaterally deviate.

Nash's theorem A Nash equilibrium always exists.

Nash proved this first by using the Kakatani fixed point theorem. Later he realised that you could prove it using the Brouwer fixed point theorem, which is weaker.

Definition If $f : X \to X$ and $f(x_0) = x_0$, x_0 is called a fixed point of f.

Brouwer's fixed point theorem Let C be a compact convex set in $\mathbb{R}^n$. Let $f : C \to C$ be a continuous function. Then f has a fixed point.

We now prove Nash's theorem from Brouwer's fixed point theorem, and then prove the Brouwer's fixed point theorem.

Proof of Nash's theorem Let $|S_i| = l_i$ be the number of moves player i has, and let

$$\mathcal{P} = \prod_{i=1}^{k} \mathcal{P}(S_i) \ .$$

Observe that $\mathcal{P}$ can be thought of as a subset of $\mathbb{R}^{\sum_{i=1}^{k} l_i}$. It is then easy to check that $\mathcal{P}$ is compact and convex.

We now define $g : \mathcal{P} \to \mathcal{P}$. Given $(\pi_1, ..., \pi_k) \in \mathcal{P}$, for every $i \in \{1, ..., k\}$ and $j \in \{1, ..., l_i\}$, let

$$\varphi_{i,j} = \max\{0, f_i(\pi_1, ..., \pi_{i-1}, \delta_j, \pi_{i+1}, ..., \pi_k) - f_i(\pi_1, ..., \pi_{i-1}, \pi_i, \pi_{i+1}, ..., \pi_k)\} \ ,$$

i.e. $\varphi_{i,j}$ is the improvement (if any) for i by changing unilaterally to δ_j from π_i. A key fact is that for all i some $\varphi_{i,j} = 0$ (check this!). Since f is continuous the $\varphi_{i,j}$ must also be. We define $g(\pi_1, ..., \pi_k)$ to be $(\tilde{\pi}_1, ..., \tilde{\pi}_k)$, where

$$\tilde{\pi}_i = \frac{\pi_i + \sum_{j=1}^{l_i} \varphi_{i,j}\delta_j}{1 + \sum_{j=1}^{l_i} \varphi_{i,j}}$$

i.e. $\tilde{\pi}_i$ is π_i modified by weighing a little heavier those pure strategies that are good, and then renormalizing. Now we have a well-defined continuous mapping g.

Next time we finish the proof by showing that $(\pi_i, ..., \pi_k)$ is a NE if and only if it is a fixed point for g. This would imply (by the Brouwer fixed point theorem) that there exists a NE.