Computational Game Theory **Spring Semester, 2005/06**

Lecture 6: April 25, 2006

Lecturer: Yishay Mansour *Scribe: Lior Gavish, Andrey Stolyarenko, Asaph Arnon*

Partially based on scribe by Nataly Sharkov and Itamar Nabriski from 2003/4

6.1 Outline

This lecture shall deal with the existence of Nash Equilibria in general (i.e. non-zero-sum) games. We start with a proof of the existence theorem. This proof uses a fixed-point theorem known as Brouwer's Lemma, which we shall prove in the following section. We then move to an algorithm for finding Nash Equilibria in two-players general sum games.

6.2 Existence Theorem

Theorem 6.1 *Every finite game has a (mixed-strategy) Nash Equilibrium.*

This section shall outline a proof of this theorem. We begin with a definition of the model, proceed with a statement of Brouwer's Lemma and conclude with the proof.

6.2.1 Model and Notations

Recall that a finite strategic game consists of the following:

- A finite set of players, namely $N = \{1, \ldots, n\}$.
- For every player i, a set of actions $A_i = \{a_{i1}, \ldots, a_{im}\}$.
- The set $A = \otimes_{i=1}^{n} A_i$ of joint actions.
- For every player i, a utility function $u_i : A \rightarrow \Re$.

A mixed strategy for player i is a random variable over his actions. The set of such strategies is denoted $\triangle(A_i)$. Letting every player have his own mixed strategy (independent of the others) leads to the set of joint mixed strategies, denoted $\triangle(A) = \otimes_{i=1}^{n} \triangle(A_i)$.

Every joint mixed strategy $p \in \triangle(A)$ consists of n vectors $\vec{p_1}, \ldots, \vec{p_n}$, where $\vec{p_i}$ defines the distribution played by player i. Taking the expectation over the given distribution, we define the utility for player i by

$$u_i(p) = E_{a \sim p}\left[u_i(a)\right] = \sum_{a \in A} p(a) u_i(a) = \sum_{a \in A} \left(\prod_{i=1}^{n} \vec{p_i}(a_i) \right) u_i(a)$$

We can now define a Nash Equilibrium (NE) as a joint strategy where no player profits from unilaterally changing his strategy:

Definition A joint mixed strategy $p^* \in \triangle(A)$ is NE, if for every player $1 \le i \le n$ it holds that

$$\forall q_i \in \triangle(A_i) \quad u_i(p^*) \ge u_i(p^*_{-i}, q_i)$$

or equivalently

$$\forall a_i \in A_i \quad u_i(p^*) \ge u_i(p^*_{-i}, a_i)$$

6.2.2 Brouwer's Lemma

The following lemma shall aid us in proving the existence theorem:

Lemma 6.2 *Let B be a compact (i.e. closed and bounded) set. Further assume that B is convex. If $f : B \to B$ is a continuous map, then there must exist $x \in B$ such that $f(x) = x$.*

We will sketch the proof later. First, let us explore some examples:

B	**f(x)**	**Fixed Points**
$[0,1]$	x^2	$0, 1$
$[0,1]$	$1-x$	$\frac{1}{2}$
$[0,1]^2$	(x^2, y^2)	$\{0,1\} \times \{0,1\}$
unit ball (in polar coord.)	$(\frac{r}{2}, 2\theta)$	$(0, \theta)$ for all θ

6.2.3 Proof of Existence of Nash Equilibrium

We now turn to the proof of the existence theorem. For $1 \le i \le n$, $j \in A_i$, $p \in \triangle(A)$ we define

$$g_{ij}(p) = \max\{u_i(p_{-i}, a_{ij}) - u_i(p), 0\}$$

to be the gain for player i from switching to the deterministic action a_{ij}, when p is the joint strategy (if this switch is indeed profitable). We can now define a continuous map between mixed strategies $y : \triangle(A) \to \triangle(A)$ by

$$y_{ij}(p) = \frac{p_{ij} + g_{ij}(p)}{1 + \sum_{j=1}^{m} g_{ij}(p)}.$$

Observe that:

- For every player i and action a_{ij}, the mapping $y_{ij}(p)$ is continuous (w.r.t. p). This is due to the fact that $u_i(p)$ is obviously continuous, making $g_{ij}(p)$ and consequently $y_{ij}(p)$ continuous.

- For every player i, the vector $(y_{ij}(p))_{j=1}^m$ is a distribution, i.e. it is in $\triangle(A_i)$. This is due to the fact that the denominator of $y_{ij}(p)$ is a normalizing constant for any given i.

Therfore y fullfils the conditions of Brouwer's Lemma. Using the lemma, we conclude that there is a fixed point p for y. This point satisfies

$$p_{ij} = \frac{p_{ij} + g_{ij}(p)}{1 + \sum_{j=1}^m g_{ij}(p)}.$$

This is possible only in one of the following cases. Either $g_{ij}(p) = 0$ for every i and j, in which case we have an equilibrium (since no one can profit from changing his strategy). Otherwise, assume there is a player i s.t. $\sum_{j=1}^m g_{ij}(p) > 0$. Then,

$$p_{ij} \left(1 + \sum_{j=1}^m g_{ij}(p) \right) = p_{ij} + g_{ij}(p)$$

or

$$p_{ij} \left(\sum_{j=1}^m g_{ij}(p) \right) = g_{ij}(p).$$

This means that $g_{ij}(p) = 0$ iff $p_{ij} = 0$, and therefore $p_{ij} > 0 \Rightarrow g_{ij}(p) > 0$. However, this is **impossible** by the definition of $g_{ij}(p)$. Recall that $u_i(p)$ is a mean with weights p_{ij}. Therefore, it cannot be that player i can profit from every pure action in $\vec{p_i}$'s support (with respect to the mean).

We are therefore left with the former possibility, i.e. $g_{ij}(p) = 0$ for all i and j, implying a NE. Q.E.D.

6.3 Brouwer's Lemma

Let us restate the lemma:

Lemma 6.3 *(Brouwer) Let $f : B \to B$ be a continuous function from a non-empty, compact, convex set $B \subset \Re^n$ to itself. Then there is $x^* \in S$ such that $x^* = f(x^*)$ (i.e. x^* is a fixed point of f).*

We shall first show that the conditions are necessary, and then outline a proof in 1D and in 2D. The proof of the general N-D case is similar to the 2D case (and is omitted).

6.3.1 Necessity of Conditions

To demonstrate that the conditions are necessary, we show a few examples:

When B is not bounded: Consider $f(x) = x + 1$ for $x \in \Re$. Then, there is obviously no fixed point.

When B is not closed: Consider $f(x) = x/2$ for $x \in (0, 1]$. Then, although $x = 0$ is a fixed point, it is not in the domain.

When B is not convex: Consider a circle in 2D with a hole in its center (i.e. a ring). Let f rotate the ring by some angle. Then, there is obviously no fixed point.

6.3.2 Proof of Brouwer's Lemma for 1D

Let $B = [0, 1]$ and $f : B \longrightarrow B$ be a continuous function. We shall show that there exists a fixed point, i.e. there is a x_0 in $[0, 1]$ such that $f(x_0) = x_0$. There are 2 possibilities:

1. If $f(0) = 0$ or $f(1) = 1$ then we are done.
2. If $f(0) \neq 0$ and $f(1) \neq 1$. Then define:

   $$F(x) = f(x) - x$$

   In this case:

   $$F(0) = f(0) - 0 = f(0) > 0$$
   $$F(1) = f(1) - 1 < 0$$

   Thus, we have $F : [0, 1] \longrightarrow \Re$, where $F(0) \cdot F(1) < 0$. Since $f(\cdot)$ is continuous, $F(\cdot)$ is continuous as well. By the *Intermediate Value Theorem*, there exists $x_0 \in [0, 1]$ such that $F(x_0) = 0$. By definition of $F(\cdot)$:

   $$0 = F(x_0) = f(x_0) - x_0$$

   And thus:

   $$f(x_0) = x_0$$

6.3.3 Proof of Brouwer's Lemma for 2D

Let $B = [0, 1]^2$ and $f : B \longrightarrow B$ be a continuous function $f(x, y) = (x', y')$. The function can be split into two components:

$$x' = f_1(x, y) \ ; \ y' = f_2(x, y)$$

From the one-dimensional case we know that for any given value of y there is at least one value of x such that:

$$x_0 = f_1(x_0, y)$$

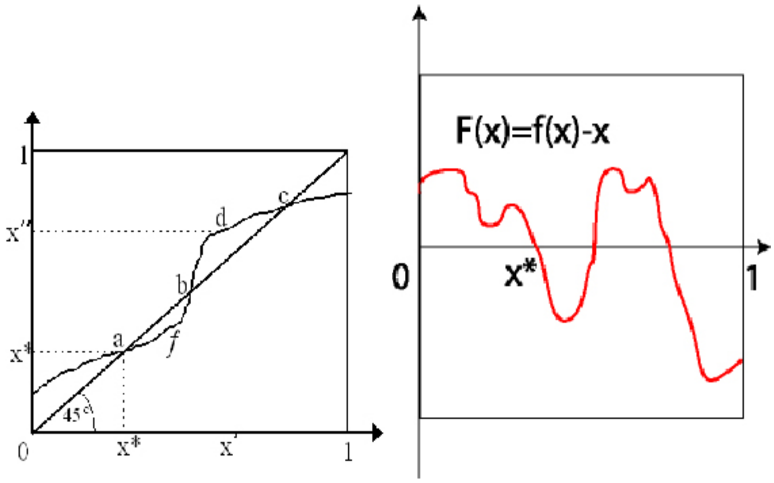

Figure 6.1: A one dimensional fixed point (left) and the function $F(\cdot)$ (right)

Let us (falsely) assume that there is always one such value x_0 for any given y. The existence of an overall fixed point will follow immediately. Let $F(y)$ denote the fixed value of x for any given y. F is a continuous function of y (this can be easily proven from f's continuity), so the composite function:

$$y' = f_2(F(y), y)$$

is a continuous function of y, and hence we can invoke the one-dimensional case again to assert the existence of a value y_0 such that:

$$y_0 = f_2(F(y_0), y_0)$$

It is now clear that $(F(y_0), y_0)$ is a fixed point.

Analogously, we could let $G(x)$ define the fixed value of y for any given x. On a two-dimensional plot, the function $y = G(x)$ must pass continuously from the left edge to the right edge of the cell, and the function $x = F(y)$ must pass continuously from the top edge to the bottom edge. Obviously, they must intersect at some point - the fixed point - as shown in Figure 6.2.

However, the assumption that there is always one value x_0 such that $x_0 = f_1(x_0, y)$ for any given y is not true. The number of such fixed values may change as we vary y. Thus, it is not straightforward to define continuous functions F and G passing between opposite edge pairs. In general, for any given value of y, the function $y' = f_1(x, y)$ could have multiple fixed points, as shown in Figure 6.3.

Furthermore, as y is varied, the curve of $f_1(x, y)$ can shift and change shape in such a way that some of the fixed points "disappear" and others "appear" elsewhere. Notice that the single-valued functions f_1 and f_2 (when restricted to the interior of the interval) cross the diagonal (i.e. $y = x$) an odd number of times. If the curve is tangent to the diagonal, but does not cross it, we will consider the fixed point as 2 fixed points with the same value. The edges of the domain can be treated separately (and

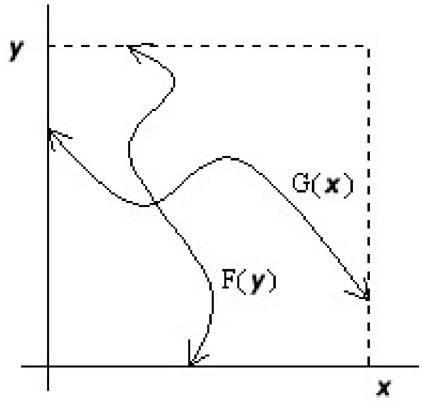

Figure 6.2: The fixed point functions $G(x)$ and $F(y)$ must intersect at some point.

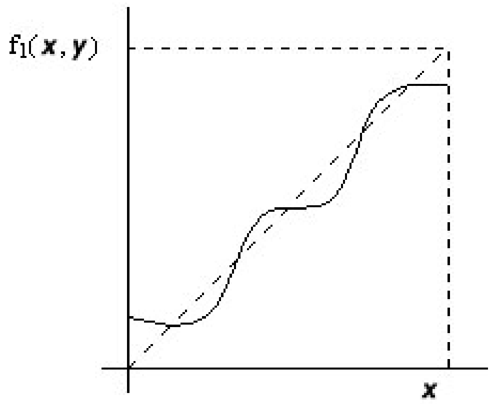

Figure 6.3: The function $f_1(x, y)$ can have multiple fixed points.

we shall not discuss them here). A plot of the multi-valued function $F(y)$, mapping y to the fixed points of $f_1(x, y)$ for an illustrative case is shown in Figure 6.4.

We see that for each value of y, there is an odd number of fixed points of $f_1(x, y)$ (when double counting tangent fixed points). This implies that there is a continuous path extending all the way from the left edge to the right edge (the path AB in Figure 6.4). This is due to the fact that a closed loop contributes an even number of crossings at each y. Only paths that start at one edge and end at the other have odd parity. It follows that there must be curves that extend continuously from one edge to the other[1]. Similarly, the corresponding function $G(x)$ has a contiguous path from the top edge to the bottom edge. As a result, the intersection between the multi-valued functions F and G is guaranteed. This intersection is f's fixed point.

[1]Notice that there might be a continuous segment of fixed points for a given y. We shall not discuss this case here.

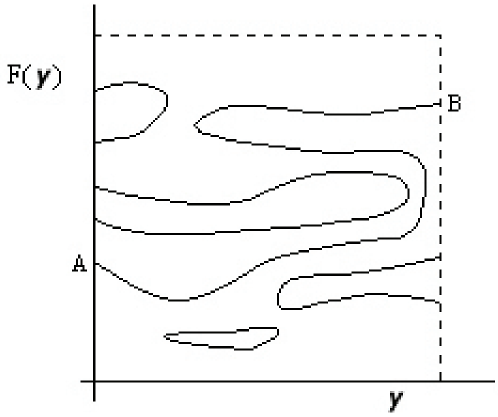

Figure 6.4: The fixed points function F of x as a function of y. AB is the continuous path from the left side to the right.

6.4 Computing Nash Equilibria

We shall first examine an example of a general game, and then describe how to generally compute a NE for a 2-player game.

6.4.1 Simple Example

Consider the following general sum game. Let U and V be the payoff matrices of players 1 and 2 respectively. Player 1 has two actions while Player 2 has four actions. Hence, the payoff matrices are 2x4 (note that V is slightly different from the one shown in class).

$$U = \begin{bmatrix} 1 & 5 & 7 & 3 \\ 2 & 3 & 4 & 3 \end{bmatrix}$$

$$V = \begin{bmatrix} 2 & 1 & 1 & 3 \\ 4 & 5 & 6 & 0 \end{bmatrix}$$

We assume the strategy for Player 1 is $(1-p, p)$, i.e., the bottom action has probability p while the upper has probability $1-p$. Player 2's strategy is a vector of length 4. The utility for Player 2 from playing action $j \in [1,4]$ is $V_{1j}(1-p) + V_{2j}p$, as depicted in Figure 6.5.

Player 2's best response (BR) as a function of p is given by the upper envelope of the graph. We examine the following points as candidates for Nash Equilibrium:

- The extreme point $p = 0$ (Player 1 chooses upper strategy). Player 2's BR is to play $j = 4$. Therefore, the payoff vector for player 1 is $(3,3)$, and thus he is indifferent, so this is an equilibrium.

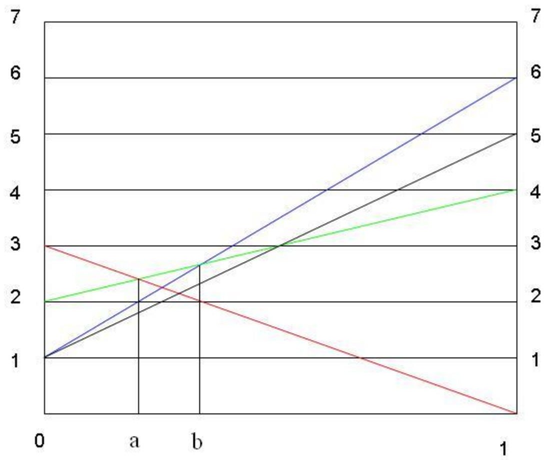

Figure 6.5: Player 2's utility as a function of p. The lines are derived from coloumn vectors of V.

- The extreme point $p = 1$ (Player 1 chooses bottom strategy). Player 2's BR is to play $j = 3$. Therefore, the payoff vector for player 1 is $(7, 4)$, and thus player 1 prefers to play $p = 0$, so this is not an equilibrium.
- The left segment (red, $p \in (0, a)$). Player 2 playes $j = 4$, and so player 1 is indifferent, because his payoff vector is $(3, 3)$. Therefore, we have an equilibrium.
- The middle segment (green, $p \in (a, b)$). Player 2 plays $j = 1$, and so Player 1's payoff vector is $(1, 2)$. His BR is $p = 1$, but $p \in (a, b)$, so we have no equilibrium.
- The right segment (blue, $p \in (b, 1)$). Player 2 plays $j = 3$, and so player 1's payoff vector is $(7, 4)$. His BR is $p = 0$, but $p \in (b, 1)$, so this is not an equilibrium.
- The intersection points. NE in intersection points occurs whenever Player 1's BR 'flips'.

  For Player 1, the first point ($p = a$) is a transition bewteen payoffs $(3, 3)$ and $(1, 2)$. So, he should play $p = 1$ (since $1 \leq 2$, and $3 \leq 3$). However, recall $p = a$, so there is no equilibrium.

  The second point ($p = b$) is a transition between $(1, 2)$ and $(7, 4)$. There is a 'flip' in player 1's BR, and consequently a NE. We shall now derive the NE. We start by calculating p as the intertesection point:

$$2p + 4(1 - p) = p + 6(1 - p)$$

$$p = 2(1 - p)$$

$$\Rightarrow p = \frac{2}{3}$$

We now calculate q, the probability of Player 2 choosing the green (i.e. first) strategy. Since this is a NE, Player 1 should be indifferent, so:

$$1 \cdot q + 7(1-q) = 2q + 4(1-q)$$

$$3(1-q) = q$$

$$\Rightarrow q = \frac{3}{4}$$

To summarize, a NE occurs when Player 1 plays the second action with probability $p = \frac{2}{3}$, and Player 2 plays strategy 1 with probability $q = \frac{3}{4}$ and strategy 3 with probability $1 - q = \frac{1}{4}$.

6.4.2 Algorithm for Finding Nash Equilibria

Consider a general two-players game, where U is the payoff matrix for Player 1 and V is the payoff matrix for Player 2. According to the existence theorem, a Nash Equilibrium (p, q) does exist.

Suppose we know the support of p and q. Let p have support $S_1 = \{i : p_i \neq 0\}$ and q have support $S_2 = \{j : q_j \neq 0\}$. How can the Nash Equilibrium be computed? Requiring that Player 1 plays his best response, we have:

$$\forall i \in S_1 \quad \forall l \in A_1 \quad e_i U q^T \geq e_l U q^T,$$

where e_i is the i-th row of the identity matrix. It is clear that if $i, l \in S_1$, then $e_i U q^T = e_l U q^T$. Similarly, for Player 2:

$$\forall j \in S_2 \quad \forall l \in A_2 \quad pVe_j^T \geq pVe_l^T$$

Of course, since p and q are distributions, and since S_1 and S_2 are their supports, we must also require

$$\sum_i p_i = 1$$

$$\forall i \in S_1 \quad p_i > 0$$

$$\forall i \notin S_1 \quad p_i = 0$$

and an analogous set of constraints on q. The inequalities with the constraints can be solved by Linear Programming algorithms. However, introducing variables u and v for the payoffs of the players in the equilibrium, we obtain the following equations

$$\forall i \in S_1 \quad e_i U q^T = u$$

$$\forall j \in S_2 \quad pUe_j^T = v$$

$$\forall i \notin S_1 \quad p_i = 0$$

$$\forall j \notin S_2 \quad q_j = 0$$

$$\sum_i p_i = 1$$

$$\sum_j q_j = 1$$

These equations can be solved more efficiently than the original inequalities.

Notice there are $2(n+1)$ equations with $2(n+1)$ variables. When the equations are non-degenerate, there is a unique solution, which must be tested for validity (i.e. $p_i \geq 0$). Otherwise, there might be an infinite number of Nash Equilibriums.

The algorithm for computing a NE is now straightforward. For all possible supports S_1, S_2 (since we do not know them in advance), we solve the equations. Whenever we find a valid solution, we can output it, as it is a NE.

Unfortunately, this algorithm has worst-case exponential running time, since there are $2^n \cdot 2^n = 4^n$ possible supports (the support could be any subset of the actions for each player). We cannot, however, expect to do much better at finding all Nash Equilibria. Consider, for instance, a game where the payoff for both players is the identity matrix, i.e. $U_{ij} = V_{ij} = \delta_{ij}$. Then, when both players play a uniform distribution on the same support (for any support), we have a NE. Therefore, there are 2^n Nash Equilibria in this game.

6.5 Approximate Nash Equilibrium

Definition A joint mixed strategy p^* is ϵ-Nash if for every player i and every mixed strategy $a \in \triangle(A_i)$ it holds that

$$u_i(p^*_{-i}, a) \leq u_i(p^*) + \epsilon$$

Note that if $\epsilon = 0$ then ϵ-Nash Equilibrium becomes a general Nash Equilibrium. This definition is useful for cases where changing strategy has a cost of ϵ. Therefore, a player will not bother changing his strategy for less than ϵ.