

Existence of Mixed Nash Equilibria

Instructor: Thomas Kesselheim

Recall that last lecture we introduced normal-form games. We defined them as cost-minimization games but often it is more natural that player want to maximize their payoff or utility rather than minimizing their cost.

Definition 4.1. A (normal-form, payoff-maximization) game is a triple $(\mathcal{N}, (S_i)_{i \in \mathcal{N}}, (u_i)_{i \in \mathcal{N}})$. Here, \mathcal{N} is the set of players, $|\mathcal{N}| = n$, often $\mathcal{N} = \{1, \dots, n\}$. For each player $i \in \mathcal{N}$, S_i is the set of (pure) strategies of player i . The set $S = \prod_{i \in \mathcal{N}} S_i$ is called the set of states or strategy profiles. For each $i \in \mathcal{N}$, $u_i: S \rightarrow \mathbb{R}$ is the payoff function of player i . In state $s \in S$, player i has a payoff of $u_i(s)$.

These two perspectives are entirely equivalent. We can derive a cost-minimization game by setting $c_i(s) = -u_i(s)$ and reuse the definitions. So, in particular, again mixed Nash equilibria are defined by requiring each player to be playing a best response.

Recall this lemma from last lecture, it will be useful and important today.

Lemma 4.2. A mixed strategy σ_i is a best-response strategy against σ_{-i} if and only if $u_i(\sigma_i, \sigma_{-i}) \geq u_i(s'_i, \sigma_{-i})$ for all pure strategies $s'_i \in S_i$.

1 Nash's Theorem

We saw that in this more general class of games, pure Nash equilibria do not necessarily exist. As we will prove today, mixed Nash equilibria always exist if the number of players and the number of strategies is finite.


Theorem 4.3 (Nash's Theorem). Every finite normal-form game has a mixed Nash equilibrium.

Nash's theorem is usually proved via Brouwer's fixed point theorem.

Theorem 4.4 (Brouwer's Fixed Point Theorem). Every continuous function $f: D \rightarrow D$ mapping a compact and convex nonempty subset $D \subseteq \mathbb{R}^m$ to itself has a fixed point $x^* \in D$ with $f(x^*) = x^*$.

As a reminder, these are the definitions of the terms used in Brouwer's fixed point theorem. Here, $\|\cdot\|$ denotes an arbitrary norm, for example, $\|x\| = \max_i |x_i|$.

- A set $D \subseteq \mathbb{R}^m$ is *convex* if for any $x, y \in D$ and any $\lambda \in [0, 1]$ we have $\lambda x + (1 - \lambda)y \in D$.



- A set $D \subseteq \mathbb{R}^m$ is *compact* if and only if it is closed and bounded.

- A set $D \subseteq \mathbb{R}^m$ is *bounded* if and only if there is some bound $r \geq 0$ such that $\|x\| \leq r$ for all $x \in D$.
- A set $D \subseteq \mathbb{R}^m$ is *closed* if it contains all its limit points. That is, consider any convergent sequence $(x_n)_{n \in \mathbb{N}}$ within D , i.e., $\lim_{n \rightarrow \infty} x_n$ exists and $x_n \in D$ for all $n \in \mathbb{N}$. Then $\lim_{n \rightarrow \infty} x_n \in D$.

$[0, 1]$ is closed and bounded

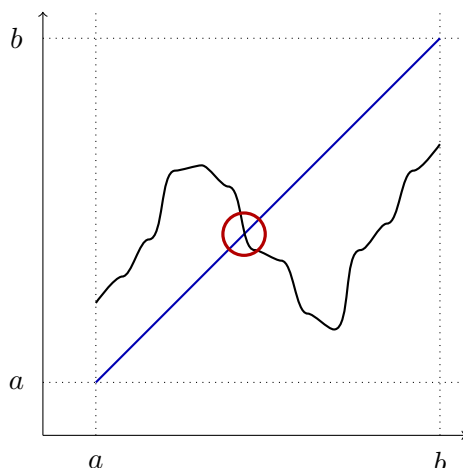
$(0, 1)$ is not closed but bounded

$[0, \infty)$ is closed and unbounded

- A function $f: D \rightarrow \mathbb{R}^m$ is *continuous at a point* $x \in D$ if for all $\epsilon > 0$, there exists $\delta > 0$, such that for all $y \in D$: If $\|x - y\| < \delta$ then $\|f(x) - f(y)\| < \epsilon$.
 f is called *continuous* if it is continuous at every point $x \in D$.

Equivalent formulation of Brouwer's fixed point theorem in one dimension:

For all $a, b \in \mathbb{R}$, $a < b$, every continuous function $f: [a, b] \rightarrow [a, b]$ has a fixed point.



Proof of Theorem 4.3. Consider a finite normal form game. Without loss of generality we assume it to be a payoff-maximization game. Let $\mathcal{N} = \{1, \dots, n\}$, $S_i = \{1, \dots, m_i\}$. So the set of mixed states X can be considered a subset of \mathbb{R}^m with $m = \sum_{i=1}^n m_i$.

Exercise: Show that X is convex and compact.

We will define a function $f: X \rightarrow X$ that transforms a mixed strategy profile into another mixed strategy profile. The fixed points of f are shown to be the mixed Nash equilibria of the game.

For mixed state x and for $i \in \mathcal{N}$ and $j \in S_i$, let

$$\phi_{i,j}(x) = \max\{0, u_i(j, x_{-i}) - u_i(x)\} .$$

So, $\phi_{i,j}(x)$ is the amount by which player i 's payoff would increase when unilaterally moving from x to j if this quantity is positive, otherwise it is 0. Observe that u_i is a continuous function as it is defined as $u_i(x) = \sum_{s_1=1}^{m_1} \dots \sum_{s_n=1}^{m_n} x_{1,s_1} \dots x_{n,s_n} u_i(s)$. Therefore $\phi_{i,j}$ is also continuous.

Observe that by Lemma 4.2 a mixed state x is a Nash equilibrium if and only if $\phi_{i,j}(x) = 0$ for all $i = 1, \dots, n$, $j = 1, \dots, m_i$.

Define $f: X \rightarrow X$ with $f(x) = x' = (x'_{1,1}, \dots, x'_{n,m_n})$ by

$$x'_{i,j} = \frac{x_{i,j} + \phi_{i,j}(x)}{1 + \sum_{k=1}^{m_i} \phi_{i,k}(x)}$$

for all $i = 1, \dots, n$ and $j = 1, \dots, m_i$.

Observe that $x^* \in X$. That means, $f: X \rightarrow X$ is well defined. Furthermore, f is continuous because each $\phi_{i,j}$ is. Therefore, by Theorem 4.4, f has a fixed point, i.e., there is a point $x^* \in X$ such that $f(x^*) = x^*$.

We only need to show that every fixed point x^* of f is a mixed Nash equilibrium. So, in other words, we need to show that $f(x^*) = x^*$ implies that $\phi_{i,j}(x^*) = 0$ for all $i = 1, \dots, n$, $j = 1, \dots, m_i$.

Fix some $i \in \mathcal{N}$. Once we have shown that $\phi_{i,j}(x^*) = 0$ for $j = 1, \dots, m_i$, we are done.

Let j' be chosen such that $u_i(j', x_{-i}^*)$ is minimized among the j' such that $x_{i,j'}^* > 0$. As $u_i(x^*)$ is defined to be $\sum_{j=1}^{m_i} x_{i,j}^* \cdot u_i(j, x_{-i}^*)$, we have $u_i(x^*) = \sum_{j=1}^{m_i} x_{i,j}^* \cdot u_i(j, x_{-i}^*) \geq \sum_{j=1}^{m_i} x_{i,j}^* \cdot u_i(j', x_{-i}^*) = u_i(j', x_{-i}^*)$. Therefore $\phi_{i,j'}(x^*) = \max\{0, u_i(j', x_{-i}^*) - u_i(x^*)\} = 0$.

We now use the fact that x^* is a fixed point. Therefore, we have

$$x_{i,j'}^* = \frac{x_{i,j'}^* + \phi_{i,j'}(x^*)}{1 + \sum_{k=1}^{m_i} \phi_{i,k}(x^*)} = \frac{x_{i,j'}^*}{1 + \sum_{k=1}^{m_i} \phi_{i,k}(x^*)} .$$

As $x_{i,j'}^* > 0$, we also have

$$1 = \frac{1}{1 + \sum_{k=1}^{m_i} \phi_{i,k}(x^*)} ,$$

and so

$$\sum_{k=1}^{m_i} \phi_{i,k}(x^*) = 0 .$$

Since $\phi_{i,k}(x^*) \geq 0$ for all k , we have to have $\phi_{i,k}(x^*) = 0$ for all k . This completes the proof. \square

2 Computing Nash Equilibria in Bimatrix Games

While we now know that mixed Nash equilibria always exist, we have no idea how to compute them. There are infinitely many mixed strategies, so even exhaustive search is not an option.

In the remainder of today's lecture, we will see how to reduce the search space to a finite one and get an idea how a more sophisticated algorithm work. To this end, we will consider a bimatrix utility-maximization game. Let $A \in \mathbb{R}^{m_1 \times m_2}$ be the row player's utility matrix, and $B \in \mathbb{R}^{m_1 \times m_2}$ be the column player's utility matrix. Mixed strategies correspond to real-valued vectors. We call $x \in \mathbb{R}^{m_1}$ the row player's mixed strategy and $y \in \mathbb{R}^{m_2}$ the column player's mixed strategy. Note that by this notation, we can simply write the players' utilities as matrix products, namely $u_{\text{row}}(x, y) = x^T A y$, $u_{\text{col}}(x, y) = x^T B y$.

Without loss of generality, we can assume that all entries in A and B are positive. This is because adding some value c to all entries in A or B shifts all utilities the same way.

Furthermore, for a vector x let its $\text{supp}(x)$ denote the set of entries in which it is strictly positive (support), i.e., $\text{supp}(x) = \{i \mid x_i > 0\}$. Given this new notation, Lemma 4.2 can be re-written as follows.

Lemma 4.5. x is best response to y if and only if for all $i \in \text{supp}(x)$: $(Ay)_i = \max_k (Ay)_k$
 y is best response to x if and only if for all $i \in \text{supp}(y)$: $(x^T B)_i = \max_k (x^T B)_k$

Using this lemma, we can state: a pair of arbitrary real-valued vectors $(x, y) \in \mathbb{R}^{m_1} \times \mathbb{R}^{m_2}$ is

a mixed Nash equilibrium if and only if there are numbers $u, v > 0$ such that

$$\begin{aligned} \sum_i x_i &= 1 & (Ay)_i &\leq u & \text{for all } i & & (Ay)_i &= u & \text{for all } i \in \text{supp}(x) \\ \sum_i y_i &= 1 & (x^T B)_i &\leq v & \text{for all } i & & (x^T B)_i &= v & \text{for all } i \in \text{supp}(y) \\ x, y &\geq 0 \end{aligned}$$

We can substitute \tilde{x} for x/v , \tilde{y} for y/u . To find a mixed Nash equilibrium, it is sufficient to find real vectors \tilde{x}, \tilde{y} such that there are $u, v > 0$ for which

$$\begin{aligned} \sum_i \tilde{x}_i &= 1/v & (A\tilde{y})_i &\leq 1 & \text{for all } i & & (A\tilde{y})_i &= 1 & \text{for all } i \in \text{supp}(\tilde{x}) \\ \sum_i \tilde{y}_i &= 1/u & (\tilde{x}^T B)_i &\leq 1 & \text{for all } i & & (\tilde{x}^T B)_i &= 1 & \text{for all } i \in \text{supp}(\tilde{y}) \\ \tilde{x}, \tilde{y} &\geq 0 \end{aligned}$$

Note that $u, v > 0$ only appear in the conditions $\sum_i \tilde{x}_i = 1/v$ and $\sum_i \tilde{y}_i = 1/u$. As vectors are non-negative, they exist if and only if $\tilde{x} \neq 0$ and $\tilde{y} \neq 0$.

This, once again, simplifies our task: We have to find vectors $\tilde{x}, \tilde{y} \neq 0$ for which

$$\begin{aligned} (A\tilde{y})_i &\leq 1 & \text{for all } i & & (A\tilde{y})_i &= 1 & \text{for all } i \in \text{supp}(\tilde{x}) \\ (\tilde{x}^T B)_i &\leq 1 & \text{for all } i & & (\tilde{x}^T B)_i &= 1 & \text{for all } i \in \text{supp}(\tilde{y}) \\ \tilde{x}, \tilde{y} &\geq 0 \end{aligned}$$

This already gives us a naive algorithm that runs in finite time: Try out any non-empty set for $\text{supp}(\tilde{x})$ and $\text{supp}(\tilde{y})$. Note that we now have $|\text{supp}(\tilde{x})| + |\text{supp}(\tilde{y})|$ variables to set and the same number of linear equations. So, usually, there will be a unique solution. Find this solution and check whether it also fulfills the inequalities. We already followed this approach last time when we computed the mixed Nash equilibria of the 2×2 Inspection Game. We guessed that both supports comprise both strategies and solved the resulting equalities.

3 Bonus: Lemke–Howson Algorithm

A smarter algorithm is the Lemke–Howson algorithm. It is a lot like the simplex algorithm for linear programming. The idea is to simplify notation by combining the two matrices A and B into one matrix C by setting

$$C = \begin{pmatrix} 0 & B^T \\ A & 0 \end{pmatrix}$$

Then, the problem becomes to find $\tilde{z} \neq 0$ such that

$$(C\tilde{z})_i \leq 1 \quad \text{and} \quad \tilde{z}_i \geq 0 \quad \text{and} \quad (\tilde{z}_i = 0 \text{ or } (C\tilde{z})_i = 1) \quad \text{for all } i$$


The conditions $(C\tilde{z})_i \leq 1$ and $\tilde{z}_i \geq 0$ define a polytope in $\mathbb{R}^{m_1+m_2}$. Each condition $(C\tilde{z})_i = 1$ and $\tilde{z}_i = 0$ corresponds to a hyperplane in this space. The hyperplanes are the boundaries of the polytope. Overall, we have $2(m_1 + m_2)$ such hyperplanes. In each vertex of the polytope (in the non-degenerate case), $m_1 + m_2$ hyperplanes intersect because this is the dimension of the space. We have to find the point at which *the right hyperplanes* intersect, namely for each i we either want $\tilde{z}_i = 0$ or $(C\tilde{z})_i = 1$.

It is easy to find a vertex fulfilling this condition for all but one i . The Lemke–Howson algorithm starts at such a vertex and moves to one for which again all or all but one of the conditions are fulfilled. If one does things the right way, no cycles can occur and therefore eventually one has to reach a vertex at which all conditions are fulfilled.

Example 4.6. *Although matrices C derived from bimatrix games have the off-diagonal block structure, it is instructive to consider the following 2×2 matrix*

$$C = \begin{pmatrix} 1 & 2 \\ 2 & 1 \end{pmatrix} .$$

In this example, the polygon looks as follows.



There are three vertices, $(0, 0)$ and $(\frac{1}{3}, \frac{1}{3})$ have the property that for each i we have $\tilde{z}_i = 0$ or $(C\tilde{z})_i = 1$. For the other two either $\tilde{z}_1 = 0$ and $(C\tilde{z})_1 = 1$, or $\tilde{z}_2 = 0$ and $(C\tilde{z})_2 = 1$.

Recommended Literature

- Philip D. Straffin. Game Theory and Strategy, The Mathematical Association of America, fifth printing, 2004. (For basic concepts)
- J. Nash. Non-Cooperative Games. The Annals of Mathematics 54(2):286-295. (Nash's original paper)