Algorithmic Game Theory — Summer 2015, Week 1

Game Theory Basics

ETH Zürich — Paul Dütting

These notes cover the material presented in the first exercise session. We introduce basic notions from Game Theory, such as the concept of a (normal form) game and elementary equilibrium concepts.

1 Normal Form Game

Definition 1.1. *A* (normal form, cost minimization) game *is a triple* $(\mathcal{N}, (S_i)_{i\in N}, (c_i)_{i\in N})$ *where*

- $\mathcal{N}$ *is the set of players,* $n = |\mathcal{N}|$*,*
- S_i *is the set of* (pure) strategies *of player* i*,*
- $S = \prod_{i\in\mathcal{N}} S_i$ *is the set of states,*
- $c_i \colon S \to \mathbb{R}$ *is the cost function of player* $i \in \mathcal{N}$*. In state* $s \in S$*, player* i *has a cost of* $c_i(s)$*.*

We denote by $s_{-i} = (s_1, ..., s_{i-1}, s_{i+1}, ..., s_n)$ a state s without the strategy s_i. This notation allows us to concisely define a unilateral deviation of a player. For $i \in \mathcal{N}$, let $s \in S$ and $s'_i \in S_i$, then $(s'_i, s_{-i}) = (s_1, \ldots, s_{i-1}, s'_i, s_{i+1}, \ldots, s_n)$.

2 Dominant Strategy Equilibrium

Definition 1.2. *A pure strategy* s_i *is called a* dominant strategy *for player* $i \in \mathcal{N}$ *if* $c_i(s_i, s_{-i}) \leq c_i(s'_i, s_{-i})$ *for every* $s'_i \in S_i$ *and every* s_{-i}*.*

Definition 1.3. *A state* $s \in S$ *is called a* dominant strategy equilibrium *if for every player* $i \in \mathcal{N}$ *strategy* $s_i \in S_i$ *is a dominant strategy.*

A dominant strategy equilibrium provides a very strong guarantee for each player: What they do is the best they can no matter what the other players do.

Example 1.4 (Prisoner's Dilemma)**.** *Two criminals are interrogated separately. Each of them has two possible strategies: (C)onfess, remain (S)ilent. Confessing yields a smaller sentence if the other one is silent. If both confess, the sentence is larger for both (4 years) compared to when they both remain silent (2 years).*

	S	C
S	2, 2	5, 1
C	1, 5	4, 4

- *If both players remain (S)ilent, the total cost is smallest.*

- *If both players (C)onfess, the cost is larger for both of them.*
- *Still, for each player confessing is always the preference!*

In other words, for both players strategy C is a dominant strategy and the state (C, C) in which both players choose strategy C is a dominant strategy equilibrium.

Not every game has a dominant strategy equilibrium.

Example 1.5 (Battle of the Sexes). *Suppose Angelina and Brad go to the movies. Angelina prefers watching movie A, Brad prefers watching movie B. However, both prefer watching a movie together to watching movies separately.*

	A	*B*
A	*1* *2*	*6* *6*
B	*5* *5*	*2* *1*

There is no dominant strategy of either of the two player: In state (A,A) the preference for both is A. In state (B,B) the preference for both is B.

3 Pure Nash Equilibrium

Definition 1.6. *A strategy s_i is called a* best response *for player $i \in \mathcal{N}$ against a collection of strategies s_{-i} if $c_i(s_i, s_{-i}) \leq c_i(s'_i, s_{-i})$ for all $s'_i \in S_i$.*

Note: A strategy s_i is a dominant strategy if and only if s_i is a best response for all s_{-i}.

Definition 1.7. *A state $s \in S$ is called a* pure Nash equilibrium *if s_i is a best response against the other strategies s_{-i} for every player $i \in \mathcal{N}$.*

So, a pure Nash equilibrium is stable against unilateral deviation. No player can reduce his cost by only changing his only strategy.

Clearly, every dominant strategy equilibrium is a Nash equilibrium.

Example 1.8 (Prisoner's Dilemma). *Recall the game from Example 1.4. We can find its unique pure Nash equilibrium (C, C) by marking best responses with boxes.*

	S	*C*
S	*2* *2*	$\boxed{1}$ *5*
C	*5* $\boxed{1}$	$\boxed{4}$ $\boxed{4}$

Pure nash equilibria need not be unique.

Example 1.9 (Battle of the Sexes)**.** *Recall the game from Example 1.4. We can find its pure Nash equilibria (A, B) and (B, A) by marking best responses with boxes.*

	A	B
A	$\boxed{2}$, $\boxed{1}$	6, 6
B	5, 5	$\boxed{1}$, $\boxed{2}$

Not every game has a pure Nash equilibrium.

Example 1.10 (Inspection Game)**.** *Consider a game between a user of the Polybahn and an inspector from ZVV. The user can either (C)omply and buy a ticket or (D)efect and do not buy one. The inspector can decide to be (L)azy and inspect or he can decide to (I)nspect. The inspector prefers to inspect only if the user is not holding a ticket. The user, in contrast, prefers to buy a ticket only if he will be inspected.*

	C	D
L	$\boxed{0}$, 0	10, $\boxed{-10}$
I	1, $\boxed{0}$	$\boxed{6}$, 90

4 Mixed Nash Equilibrium

Definition 1.11. *A* mixed strategy *σ_i for player i is a probability distribution over the set of pure strategies S_i.*

We will only consider the case of finitely many pure strategies and finitely many players. In this case, we can write a mixed strategy σ_i as $(\sigma_{i,s_i})_{s_i \in S_i}$ with $\sum_{s_i \in S_i} \sigma_{i,s_i} = 1$. The cost of a mixed state σ for player i is

$$c_i(\sigma) = \sum_{s \in S} p(s) \cdot c_i(s) \quad ,$$

where $p(s) = \prod_{i \in \mathcal{N}} \sigma_{i,s_i}$ is the probability that the outcome is pure state s.

Definition 1.12. *A mixed strategy σ_i is a* (mixed) best-response strategy *against a collection of mixed strategies σ_{-i} if $c_i(\sigma_i, \sigma_{-i}) \leq c_i(\sigma'_i, \sigma_{-i})$ for all other mixed strategies σ'_i.*

Definition 1.13. *A mixed state σ is called a* mixed Nash equilibrium *if σ_i is a best-response strategy against σ_{-i} for every player $i \in \mathcal{N}$.*

Note that every pure strategy is also a mixed strategy and every pure Nash equilibrium is also a mixed Nash equilibrium.

It is enough to only consider deviations to pure strategies.

Lemma 1.14. *A mixed strategy σ_i is a best-response strategy against σ_{-i} if and only if $c_i(\sigma_i, \sigma_{-i}) \leq c_i(s'_i, \sigma_{-i})$ for all pure strategies $s'_i \in S_i$.*

Proof. The "only if" part is trivial: Every pure strategy is also a mixed strategy.

For the "if" part, let σ_{-i} be an arbitrary mixed strategy profile for all players except for i. Furthermore, let σ_i be a mixed strategy for player i such that $c_i(\sigma_i, \sigma_{-i}) \leq c_i(s'_i, \sigma_{-i})$ for all pure strategies $s'_i \in S_i$.

Observe that for any mixed strategy σ'_i, we have $c_i(\sigma'_i, \sigma_{-i}) = \sum_{s'_i \in S_i} \sigma'_{i,s'_i} c_i(s'_i, \sigma_{-i}) \geq \min_{s'_i \in S_i} c_i(s'_i, \sigma_{-i})$. Using $\min_{s'_i \in S_i} c_i(s'_i, \sigma_{-i}) \geq c_i(\sigma_i, \sigma_{-i})$, we are done. □

A very important property of mixed best responses is that they are always probability distributions over pure best responses.

Lemma 1.15. *A mixed strategy σ_i is a best-response strategy against σ_{-i} if and only if every strategy in the support of σ_i, i.e., every $s_j \in S_i$ with $\sigma_{i,s_j} > 0$, is a best response against σ_{-i}.*

Proof. First suppose σ_i is a distribution over pure best responses. Then for every strategy $s_i \in S_i$ with $\sigma_{i,s_i} > 0$ and every pure strategy $s'_i \in S_i$ we have

$$c_i(\sigma_i, \sigma_{-i}) = c_i(s_i, \sigma_{-i}) \leq c_i(s'_i, \sigma_{-i}).$$

Now we can invoke Lemma 1.14 to conclude that σ_i is a best response to σ_{-i}.

Next suppose σ_i is *not* a distribution over pure best responses. Then there exists strategies $s_i, s'_i \in S_i$ such that $\sigma_{i,s_i} > 0$ but $c_i(s_i, \sigma_{-i}) > c_i(s'_i, \sigma_{-i})$. We construct distribution σ'_i from σ_i by increasing σ_{i,s'_i} by σ_{i,s_i} and decreasing σ_{i,s_i} to zero. Then,

$$c_i(\sigma'_i, \sigma_{-i}) < c_i(\sigma_i, \sigma_{-i}),$$

and so σ_i cannot be a best response. □

Hence we can compute mixed Nash equilibria by choosing probabilities for one player that will make the other player indifferent between his pure strategies.

Example 1.16 (Inspection Game)**.** *Recall the game from Example 1.10.*

	C	D	
L	0 $\boxed{0}$	$\boxed{-10}$ 10	$1-p$
I	$\boxed{0}$ 1	90 $\boxed{6}$	p
	$1-q$	q	

Let us compute probabilities $(1-p, p)$ for the row player and $(1-q, q)$ for the column player that make the other player indifferent.

To determine the probabilities of the column player, we compute the expected costs for the pure strategies of the row player, equate them, and solve for q:

$$\begin{aligned} & c_{row}(D, (1-q, q)) = c_{row}(I, (1-q, q)) \\ \Leftrightarrow \quad & 0 \cdot (1-q) + 10 \cdot q = 1 \cdot (1-q) + 6 \cdot q \\ \Leftrightarrow \quad & q = \frac{1}{5}. \end{aligned}$$

Similarly, to determine the probabilities for the row player:

$$
\begin{aligned}
& c_{col}(C,(1-p,p)) = c_{col}(D,(1-p,p)) \\
\Leftrightarrow\quad & 0\cdot(1-p) + 0\cdot p = 10\cdot(1-p) + (-90)\cdot p \\
\Leftrightarrow\quad & p = \frac{1}{10}.
\end{aligned}
$$

We obtain the mixed Nash equilibrium in which the row player mixes between D and I with probabilities $(9/10, 1/10)$ and the column player mixes between C and D with probabilities $(4/5, 1/5)$.

There are a number of useful techniques that help compute mixed Nash equilibria.

Technique 1 (Elimination of Strictly Dominated Strategies)**.** *It is generally advisable to check for strictly dominated strategies. A strategy s_i is strictly dominated by some strategy s_i' if $c_i(s_i, s_{-i}) > c_i(s_i', s_{-i})$ for every s_{-i}. Strictly dominated strategies can be eliminated from the game without losing Nash equilibria.*

Technique 2 (The Difference Trick)**.** *In a 2×2 game there is a quick way to determine the probabilities that are needed for a mixed Nash equilibrium in which the players use both strategies. The trick works in the same way for both players, so we only show how to use it to compute the probabilities of the row player.*

	L	R		
T	$A_{T,L}$, $B_{T,L}$	$A_{T,R}$, $B_{T,R}$	$\Delta_T = \lvert B_{T,L} - B_{T,R}\rvert$	$1-p = \frac{\Delta_B}{\Delta_T+\Delta_B}$
B	$A_{B,L}$, $B_{B,L}$	$A_{B,R}$, $B_{B,R}$	$\Delta_B = \lvert B_{B,L} - B_{B,R}\rvert$	$p = \frac{\Delta_T}{\Delta_T+\Delta_B}$

That is, we first take the absolute difference between the payoffs of the column player for each strategy of the row player to obtain $\Delta_T = |B_{T,L} - B_{T,R}|$ and $\Delta_B = |B_{B,L} - B_{B,R}|$. Then the probability for the top strategy is $1-p = \Delta_B/(\Delta_T + \Delta_B)$ and the probability for the bottom strategy is $p = \Delta_T/(\Delta_T + \Delta_B)$.

Technique 3 (The Goalpost Method)**.** *We can also find the probabilities at which a player becomes indifferent between two of his strategies in a graphical manner. For the parametrized game considered for the difference trick:*

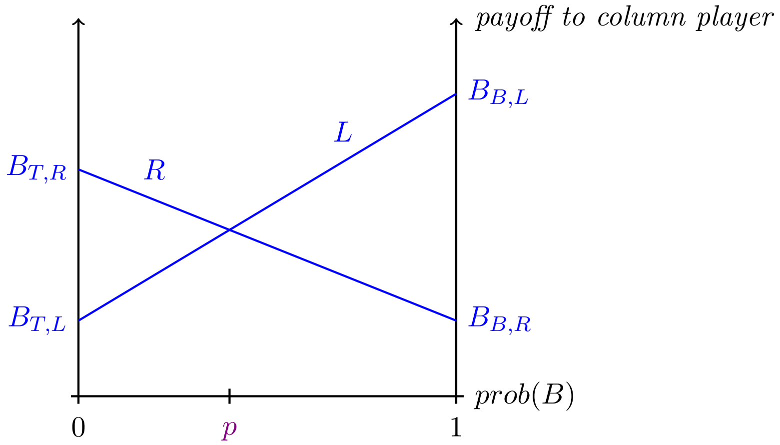

We can read off the best responses of the column player from the lower envelope. At prob$(B) = 0$ the column player prefers L, at prob$(B) = 1$ he prefers R, and at prob$(B) = p$ he is indifferent between L and R.

While dominant strategy equilibria and pure Nash equilibria do not necessarily exist, mixed Nash equilibria always exist if the number of players and the number of strategies is finite.

Theorem 1.17 (Nash' Theorem)**.** *Every finite normal form game has a mixed Nash equilibrium.*

5 Bonus Material: Proof of Nash' Theorem

We will use Brouwer's fixed point theorem to prove Nash' theorem.

Theorem 1.18 (Brouwer's Fixed Point Theorem)**.** *Every continuous function $f\colon D \to D$ 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 $\lambda x + (1-\lambda) y \in D$.

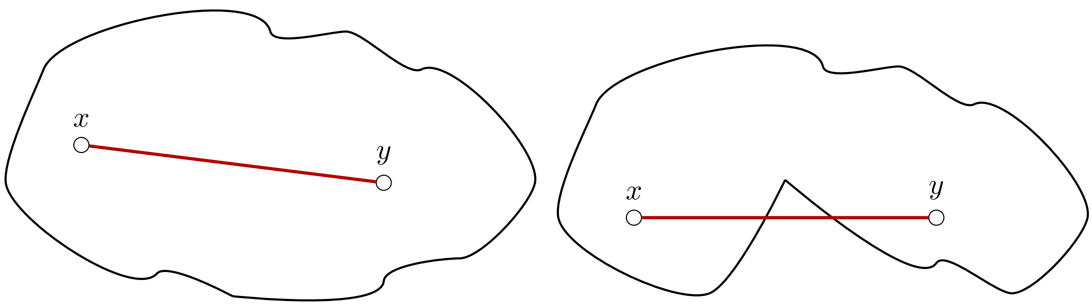

convex not convex

- 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., $\lim_{n\to\infty} x_n$ exists and $x_n \in D$ for all $n \in \mathbb{N}$. Then $\lim_{n\to\infty} x_n \in D$.

$$\begin{aligned} [0,1] &\text{ is closed and bounded} \\ (0,1] &\text{ is not closed but bounded} \\ [0,\infty) &\text{ is closed and unbounded} \end{aligned}$$

- A function $f\colon D \to \mathbb{R}^m$ 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 $f\colon [a,b] \to [a,b]$ has a fixed point.

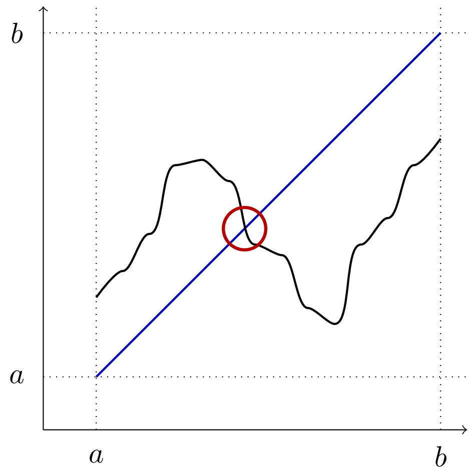

Proof of Theorem 1.17. Consider a finite normal form game. Without loss of generality let $\mathcal{N} = \{1, \ldots, n\}$, $S_i = \{1, \ldots, m_i\}$. So the set of mixed states X can be considered a subset of $\mathbb{R}^m$ with $m = \sum_{i=1}^n m_i$.

Exercise: Show that X is convex and compact.

We will define a function $f\colon X \to X$ 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

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

So, $\phi_{i,j}(x)$ is the amount by which player i's cost would reduce when unilaterally moving from x to j if this quantity is positive, otherwise it is 0.

Observe that by Lemma 1.14 a mixed state x is a Nash equilibrium if and only if $\phi_{i,j}(x) = 0$ for all $i = 1, \ldots, n$, $j = 1, \ldots, m_i$.

Define $f\colon X \to X$ with $f(x) = x' = (x'_{1,1}, ..., x'_{n,m_n})$ by

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

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

Observe that $x' \in X$. That means, $f\colon X \to X$ is well defined. Furthermore, f is continuous. Therefore, by Theorem 1.18, 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 $i = 1, \ldots, n$, $j = 1, \ldots, m_i$.

Fix some $i \in \mathcal{N}$. Once we have shown that $\phi_{i,j}(x^*) = 0$ for $j = 1, \ldots, m_i$, we are done. We observe that there is j' with $x^*_{i,j'} > 0$ and $c_i(x^*) \leq c_i(j', x^*_{-i})$ because $c_i(x^*)$ is defined to be $\sum_{j=1}^{m_i} x^*_{i,j} \cdot c_i(j, x^*_{-i})$. So, it is the weighted average of all costs and it is impossible that every pure strategy has strictly smaller cost then the weighted average. For this j', $\phi_{i,j'}(x^*) = \max\{0, c_i(x^*) - c_i(j, x^*_{-i})\} = 0$.

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. □

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' original paper)