NORMAL FORM GAMES: Strategies, dominance, and Nash

Heinrich H. Nax & Bary S. R. Pradelski

hnax@ethz.ch bpradelski@ethz.ch

March 12, 2018: Lecture 4

ETH

Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Plan

- Introduction normal form games
- Dominance in pure strategies
- Nash equilibrium in pure strategies
- Best replies
- Dominance, Nash, best replies in mixed strategies
- Nash's theorem and proof via Brouwer

The Prisoner's Dilemma

"Two suspects are arrested and interviewed separately. If they both keep quiet (i.e., *cooperate*) they go to prison for one year. If one suspect supplies evidence (*defects*) then that one is freed, and the other one is imprisoned for eight years. If both defect then they are imprisoned for five years."

PLAYERS The players are the two suspects $N = \{1, 2\}$.

STRATEGIES The strategy set for player 1 us $S_1 = \{C, D\}$, and for player 2 is $S_2 = \{C, D\}$.

PAYOFFS For example, $u_1(C, D) = -8$ and $u_2(C, D) = 0$. All payoffs are represented in this matrix:

	Cooperate	*Defect*
Cooperate	$-1, -1$	$-8, \quad 0$
Defect	$0, -8$	$-5, -5$

Definition: Normal form game

A normal form (or strategic form) game consists of three object:

1. *Players:* $N = \{1, \ldots, n\}$, with typical player $i \in N$.
2. *Strategies:* For every player i, a finite set of strategies, S_i, with typical strategy $s_i \in S_i$.
3. *Payoffs:* A function $u_i : (s_1, \ldots, s_n) \to \mathbb{R}$ mapping strategy profiles to a payoff for each player i. $u : S \to \mathbb{R}^n$.

Thus a normal form game is represented by the triplet:

$$G = \langle N, \{S_i\}_{i \in N}, \{u_i\}_{i \in N} \rangle$$

Strategies

Definition: strategy profile

$s = (s_1, \ldots, s_n)$ is called a *strategy profile*.
It is a collection of strategies, one for each player. If s is played, player i receives $u_i(s)$.

Definition: opponents strategies

Write s_{-i} for all strategies except for the one of player i. So a strategy profile may be written as $s = (s_i, s_{-i})$.

Dominance

A strategy strictly dominates another if it is always better whatever others do.

STRICT DOMINANCE A strategy s_i *strictly dominates* s'_i if $u_i(s_i, s_{-i}) > u_i(s'_i, s_{-i})$ for all s_{-i}.

	Cooperate	*Defect*
~~*Cooperate*~~	~~$-1, -1$~~	~~$-8, 0$~~
Defect	$0, -8$	$-5, -5$

Dominance

A strategy strictly dominates another if it is always better whatever others do.

STRICT DOMINANCE A strategy s_i *strictly dominates* s_i' if $u_i(s_i, s_{-i}) > u_i(s_i', s_{-i})$ for all s_{-i}.

DOMINATED STRATEGY A strategy s_i' is *strictly dominated* if there is an s_i that strictly dominates it.

DOMINANT STRATEGY A strategy s_i is *strictly dominant* if it strictly dominates all $s_i' \neq s_i$.

If players are rational they should never play a strictly dominated strategy, no matter what others are doing, they may play weakly dominated strategies:

WEAK DOMINANCE A strategy s_i *weakly dominates* s_i' if $u_i(s_i, s_{-i}) \geq u_i(s_i', s_{-i})$ for all s_{-i}.

Dominant-Strategy Equilibrium

Definition: Dominant-Strategy Equilibrium

The strategy profile s^* is a *dominant-strategy equilibrium* if, for every player i, $u_i(s_i^*, s_{-i}) \geq u_i(s_i, s_{-i})$ for all strategy profiles $s = (s_i, s_{-i})$.

Example: Prisoner's dilemma

	Cooperate	*Defect*
Cooperate	$-1, -1$	$-8, \quad 0$
Defect	$0, -8$	$-5, -5$

(D, D) is the (unique) *dominant-strategy equilibrium*.

Common knowledge of rationality and the game

Suppose that players are rational decision makers and that mutual rationality is common knowledge, that is:

- I know that she knows that I will play rational
- She knows that "I know that she knows that I will play rational"
- I know that "She knows that "I know that she knows that I will play rational""
- ...

Further suppose that all players know the game and that again is common knowledge.

Iterative deletion of strictly dominated strategies

If the game and rationality of players are common knowledge, iterative deletion of strictly dominated strategies yields the set of "rational" outcomes.

	~~L~~	C	~~R~~
~~T~~	1, 1	2, 0	2, 2
~~M~~	0, 3	1, 5	4, 4
B	2, 4	3, 6	3, 0

Note: Iteratively deletion of strictly dominated strategies is independent of the order of deletion.

Battle of the Sexes

PLAYERS The players are the two students $N = \{row, column\}$.

STRATEGIES Row chooses from $S_{row} = \{Cafe, Pub\}$
Column chooses from $S_{column} = \{Cafe, Pub\}$.

PAYOFFS For example, $u_{row}(Cafe, Cafe) = 4$. The following matrix summarises:

	Cafe	*Pub*
Cafe	4, 3	1, 1
Pub	0, 0	3, 4

Battle of the Sexes

In this game, nothing is dominated, so profiles like (Cafe, Pub) are not eliminated. Should they be?

- Column player would play Cafe if row player played Cafe!
- Row player would play Pub if column player played Pub!

In other words, after the game, both players may "regret" having played their strategies.

This a truly interactive game – best responses depend on what other players do ... next slides!

Nash Equilibrium

Definition: Nash Equilibrium

A *Nash equilibrium* is a strategy profiles s^* such that for every player i,

$$u_i(s_i^*, s_{-i}^*) \geq u_i(s_i, s_{-i}^*) \text{ for all } s_i$$

At s^*, no i regrets playing s_i^*. Given all the other players' actions, i could not have done better

Best-reply functions

What should each player do given the choices of their opponents? They should "best reply".

Definition: best-reply function

The *best-reply function* for player i is a function B_i such that:

$$B_i(s_{-i}) = \{s_i | u_i(s_i, s_{-i}) \geq u_i(s_i', s_{-i}) \text{ for all } s_i'\}$$

Best-reply functions in Nash

Nash equilibrium can be redefined using best-reply functions:

Definition: Nash equilibrium

s^* is a *Nash equilibrium* if and only if $s_i^* \in B_i(s_{-i}^*)$ for all i.

In words: a Nash equilibrium is a strategy profile of mutual best responses each player picks a best response to the combination of strategies the other players pick.

Example

For the Battle of the Sexes:

- $B_{row}(Cafe) = Cafe$
- $B_{row}(Pub) = Pub$
- $B_{column}(Cafe) = Cafe$
- $B_{column}(Pub) = Pub$

So (Cafe, Cafe) is a Nash equilibrium and so is (Pub, Pub) …

Cook book: how to find pure-strategy Nash equilibria

The best way to find (pure-strategy) Nash equilibria is to underline the best replies for each player:

	L	C	R
T	$\underline{5}, 1$	$2, 0$	$2, \underline{2}$
M	$0, 4$	$1, \underline{5}$	$\underline{4}, \underline{5}$
B	$2, 4$	$\underline{3}, \underline{6}$	$1, 0$

Hawk-dove game

		Player 2	
		Hawk	Dove
Player 1	Hawk	-2,-2	4,0
	Dove	0,4	2,2

Harmony game

		Company B	
		Cooperate	Not Cooperate
Company A	Cooperate	9,9	4,7
	Not Cooperate	7,4	3,3

A three player game

L

	l	r
T	$0, \underline{21}, 0$	$-10, 11, \underline{1}$
B	$\underline{10}, 0, -10$	$\underline{0}, \underline{10}, \underline{11}$

R

	l	r
T	$\underline{1}, \underline{11}, \underline{10}$	$\underline{11}, 1, -9$
B	$-9, 10, \underline{0}$	$1, \underline{20}, 1$

Matching Pennies

"Each player has a penny. They simultaneously choose whether to put their pennies down heads up (H) or tails up (T). If the pennies match, column receives row's penny, if they don't match, row receives columns' penny."

PLAYERS The players are $N = \{row, column\}$.

STRATEGIES Row chooses from $\{H, T\}$; Column from $\{H, T\}$.

PAYOFFS Represented in the strategic-form matrix:

	H	T
H	$\underline{1}, -1$	$-1, \underline{1}$
T	$-1, \underline{1}$	$\underline{1}, -1$

- Best replies are: $B_{row}(H) = H, B_{row}(T) = T, B_{column}(T) = H$, and $B_{column}(H) = T$
- There is no pure-strategy Nash equilibrium in this game

Randomizing the strategy

Let one player toss her coin and hence play H with probability 0.5 and L with probability 0.5.

	H	T
H	$\underline{1}, -1$	$-1, \underline{1}$
T	$-1, \underline{1}$	$\underline{1}, -1$

Expected utility of column player when playing H:

$$\frac{1}{2} \cdot (\ \ 1) + \frac{1}{2} \cdot (-1) = 0$$

Expected utility of column player when playing T:

$$\frac{1}{2} \cdot (-1) + \frac{1}{2} \cdot (\ \ 1) = 0$$

Column is indifferent! He might decide to also toss a coin!

Mixed strategies

Definition: Mixed strategy

A *mixed strategy* σ_i for a player i is any probability distribution over his or her set S_i of pure strategies. The set of mixed strategies is:

$$\Delta(S_i) = \left\{ x_i \in \mathbb{R}_+^{|S_i|} : \sum_{h \in S_i} x_{ih} = 1 \right\}$$

Mixed extension

Definition: Mixed extension

The mixed extension of a game G has players, strategies and payoffs: $\Gamma = \langle N, \{S_i\}_{i \in N}, \{U_i\}_{i \in N} \rangle$, where

1. Strategies are probability distributions in the set $\Delta(S_i)$.
2. U_i is player i's expected utility function assigning a real number to every strategy profile $\sigma = (\sigma_1, ..., \sigma_n)$.

Mixed Profiles

Suppose player i plays mixed strategy σ_i (that is, a list of probabilities). Denote their probability that this places on pure strategy s_i as $\sigma_i(s_i)$. Then:

$$U_i(\sigma) = \sum_s u_i(s) \prod_{j \in N} \sigma_j(s_j)$$

Definition: opponents' strategies

σ_{-i} is a vector of mixed strategies, one for each player, except i. So $\sigma = (\sigma_i, \sigma_{-i})$.

Example: Matching pennies

	H	T
H	$\underline{1}, -1$	$-1, \underline{1}$
T	$-1, \underline{1}$	$\underline{1}, -1$

- If row player plays $(1, 0)$ what should column play?
- If row player plays $(0.3, 0.7)$ what should column play?
- If row player plays $(0.5, 0.5)$ what should column play?

Which mixed strategy should each player use?

Best-reply function

The definition extends in a straightforward way:

Definition: best-reply function

The *best-reply function* for player i is a function β_i such that:

$$\beta_i(\sigma_{-i}) = \{\sigma_i | U_i(\sigma_i, \sigma_{-i}) \geq U_i(\sigma_i', \sigma_{-i}), \text{ for all } \sigma_i'\}$$

Example: Matching pennies

	H	T
H	$\underline{1}, -1$	$-1, \underline{1}$
T	$-1, \underline{1}$	$\underline{1}, -1$

If column player plays $(q, 1-q)$ what should row play?

- $U_{row}(H, q) = q - (1-q) = 2q - 1$, and ...
- $U_{row}(T, q) = -q + (1-q) = 1 - 2q$, so ...
- play H if $q > \frac{1}{2}$, play T if $q < \frac{1}{2}$, and ...
- indifferent if $q = \frac{1}{2}$: any $(p, 1-p)$ will do for the row player!

Best-reply graph

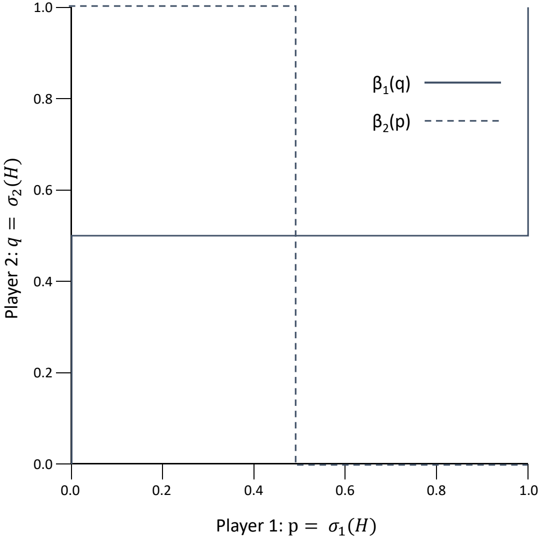

Mixed-Strategy Nash Equilibrium

Definition: Mixed-Strategy Nash Equilibrium

A *mixed-strategy Nash equilibrium* is a profile σ^* such that,

$$U_i(\sigma_i^*, \sigma_{-i}^*) \geq U_i(\sigma_i, \sigma_{-i}^*) \text{ for all } \sigma_i \text{ and } i.$$

Best replies and Nash equilibrium

Proposition

$x \in \Delta(S)$ is a Nash equilibrium if $x \in \beta(x)$.

Note that if $x \in \Delta(S)$ is a mixed Nash equilibrium, then every pure strategy in the support of each strategy x_i is a best reply to x:

$$s_i \in supp(x_i) \Rightarrow s_i \in \beta_i(x)$$

Indifference and Matching Pennies

	H	T
H	$\underline{1}, -1$	$-1, \underline{1}$
T	$-1, \underline{1}$	$\underline{1}, -1$

Suppose row player mixes with probability p and $1-p$ on H and T:

$$U_{column}(H, p) = p \cdot (-1) + (1-p) \cdot (\ \ 1) = 1 - 2p,$$
$$U_{column}(T, p) = p \cdot (\ \ 1) + (1-p) \cdot (-1) = 2p - 1$$

Column player is indifferent when $2p - 1 = 1 - 2p \Leftrightarrow p = \frac{1}{2}$.

Similarly for row player.

The only Nash equilibrium involves both players mixing with probability $\frac{1}{2}$.

Indifference and Matching Pennies

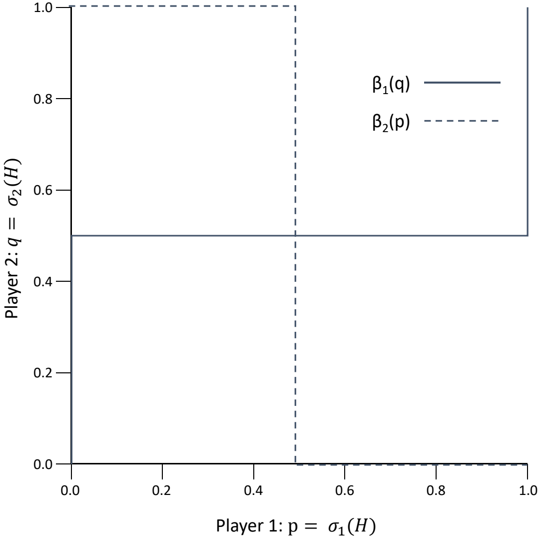

Battle of the Sexes revisited

PLAYERS The players are the two students $N = \{row, column\}$.

STRATEGIES Row chooses from $S_{row} = \{Cafe, Pub\}$
Column chooses from $S_{column} = \{Cafe, Pub\}$.

PAYOFFS For example, $u_{row}(Cafe, Cafe) = 4$. The following matrix summarises:

	$Cafe(q)$	$Pub(1-q)$	Expected
$Cafe(p)$	$4, 3$	$1, 1$	$4q+(1-q)$
$Pub(1-p)$	$0, 0$	$3, 4$	$3(1-q)$
Expected	$3p$	$p + 4(1-p)$	

Column chooses $q = 1$ whenever $3p > p + 4(1-p) \Leftrightarrow 6p > 4 \Leftrightarrow p > \frac{2}{3}$.

Row chooses $p = 1$ whenever $4q + (1-q) > 3(1-q) \Leftrightarrow 6q > 2 \Leftrightarrow q > \frac{1}{3}$.

Battle of the Sexes: Best-reply graph

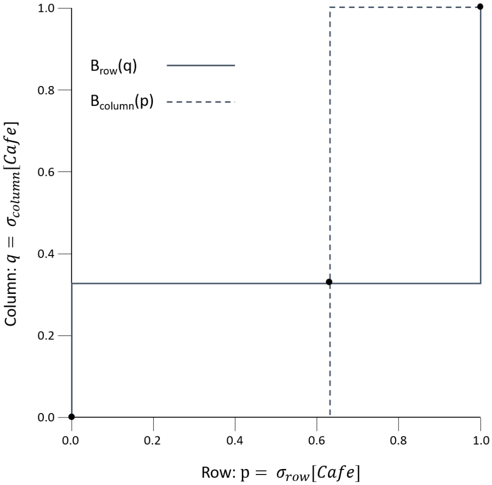

There is a mixed Nash equilibrium with $p = \frac{2}{3}$ and $q = \frac{1}{3}$.

Battle of the Sexes: Expected payoff

	$Cafe(1/3)$	$Pub(2/3)$	Expected
$Cafe(2/3)$	$4, 3$	$1, 1$	$4{\cdot}1/3{+}2/3$
$Pub(1/3)$	$0, 0$	$3, 4$	$3{\cdot}2/3$
Expected	$3 \cdot 2/3$	$2/3 + 4 \cdot 1/3$	

Frequency of play:

	$Cafe(1/3)$	$Pub(2/3)$
$Cafe(2/3)$	$2/9$	$4/9$
$Pub(1/3)$	$1/9$	$2/9$

Expected utility to row player: 2

Expected utility to column player: 2

Cook book: How to find mixed Nash equilibria

- Find all pure strategy NE.

Check whether there is an equilibrium in which row mixes between several of her strategies:

- Identify candidates:
  - If there is such an equilibrium then each of these strategies must yield the same expected payoff given column's equilibrium strategy.
  - Write down these payoffs and solve for column's equilibrium mix.
  - Reverse: Look at the strategies that column is mixing on and solve for row's equilibrium mix.
- Check candidates:
  - The equilibrium mix we found must indeed involve the strategies for row we started with.
  - All probabilities we found must indeed be probabilities (between 0 and 1).
  - Neither player has a positive deviation.

Nash's equilibrium existence theorem

Theorem (Nash 1951)

Every finite game has at least one [Nash] equilibrium in mixed strategies.

Original paper is this week's reading.

Nash's contribution – remarks

- Put economics beyond one branch of social sciences but allowed it to today encompass all analytical fields in social sciences and beyond
  - political sciences: strategic interactions, contracts, ...
  - biology: evolution
  - economics: auctions, trading, contracts, ...
  - computer sciences: cloud computing, car routing, ...
  - sociology: opinion formation, political polarization, ...
  - ...
- Nash recognized that his equilibrium concept can be used to study
  - non-cooperative games
  - cooperative games – bargaining
  - does not need to assume perfect rationality – mass-action interpretation and evolutionary game theory

THANKS EVERYBODY

See you next week!

And keep checking the website for new materials as we progress:

http://www.coss.ethz.ch/education/GT.html

Voluntary extra reading: Proof of Nash's theorem; not part of exam

Brouwer's fixed point theorem

Theorem (Brouwer)

Given $S \subset \mathbb{R}^n$ convex and compact (bounded and closed), $f : S \to S$ continuous. Then f has at least one fixed point $s \in S$ with $f(s) = s$.

Example $S = [0, 1]$

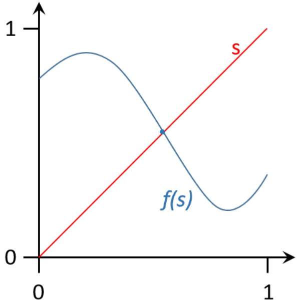

Puzzle: the football which cannot be moved

Theorem (Brouwer)

Given $S \subset \mathbb{R}^n$ convex and compact (bounded and closed), $f : S \to S$ continuous. Then f has at least one fixed point $s \in S$ with $f(s) = s$.

Can you move a football on its spot such that no point on its sphere (surface) remains in the same spot?

Proof of Nash via Bouwer

The polyhedron $\Delta(S)$ is non-empty, convex and compact.

Hence, by Bouwer, every continuous function that maps $\Delta(S)$ into itself has at least one fix point.

We thus have to find a continuous function $f : \Delta(S) \to \Delta(S)$ such that every fix point under f is a Nash equilibrium.

Nash's construction

For each player i and strategy profile σ define the *excess payoff* player i receives when playing pure strategy $h \in S_i$ in comparison with σ_i

$$v_{ih}(\sigma) = \max\{0, U_i(e_i^h, \sigma_{-i}) - U_i(\sigma)\}$$

where e_i^h is the unit vector with position h equal to 1.

Let for all $i \in N, h \in S_i$:

$$f_{ih}(\sigma) = \frac{1 + v_{ih}(\sigma)}{1 + \sum_{k \in S_i} x_{ik} v_{ik}(\sigma)} x_{ih}$$

where $\sigma_i = (x_{i1}, x_{i2}, ..., x_{i|S_i|})$.

Nash's construction

$$f_{ih}(\sigma) = \frac{1 + v_{ih}(\sigma)}{1 + \sum_{k \in S_i} x_{ik} v_{ik}(\sigma)} x_{ih}$$

We have

- $f_{ih}(\sigma) \geq 0$
- $\sum_h f_{ih}(\sigma) = 1$ for all $i \in N$ and $\sigma \in \Delta(S)$
- $f_{ih}(\sigma)$ is continuous in σ

Thus f is a continuous mapping of $\Delta(S)$ to itself

$\Rightarrow f$ has at least one fix point

Nash's construction

Suppose that σ is a fixpoint of f, that is $\sigma = f(\sigma)$. We must have

$$\begin{aligned} 0 &= f_{ih}(\sigma) - x_{ih} \\ &= \frac{1 + v_{ih}(\sigma)}{1 + \sum_{k \in S_i} x_{ik} v_{ik}(\sigma)} x_{ih} - x_{ih} \\ &= \frac{x_{ih} + v_{ih}(\sigma) x_{ih} - x_{ih} - x_{ih} \sum_{k \in S_i} x_{ik} v_{ik}(\sigma)}{1 + \sum_{k \in S_i} x_{ik} v_{ik}(\sigma)} \\ &= [v_{ih}(\sigma) - \sum_{k \in S_i} x_{ik} v_{ik}(\sigma)] x_{ih} = 0 \end{aligned}$$

for all $i \in N, h \in S_i$.

Nash's construction: fixpoint $\iff$ equilibrium

$$[v_{ih}(\sigma) - \sum_{k \in S_i} x_{ik} v_{ik}(\sigma)] x_{ih} = 0$$

"$\Rightarrow$": This equation is satisfied for $v_{ih}(\sigma) = 0$ for all $i \in N, h \in S_i$, that is, σ is a [Nash] equilibrium.

"$\Leftarrow$": Suppose the equation is satisfied by some $\sigma \in \Delta(S)$ which is not a Nash equilibrium:

$$v_{ih}(\sigma) = \sum_{k \in S_i} x_{ik} v_{ik}(\sigma)$$

for all i, h with $x_{ih} > 0$.
But this implies that $v_{ih} = 0$ for all such i, h, since otherwise all used pure strategies would earn above average, an impossibility.

□