Best Response

Definition 33
A strategy $\sigma_i \in \Sigma_i$ of player i is a *best response* to a strategy profile $\sigma_{-i} \in \Sigma_{-i}$ of his opponents if

$$u_i(\sigma_i, \sigma_{-i}) \geq u_i(\sigma_i', \sigma_{-i}) \qquad \text{for all } \sigma_i' \in \Sigma_i$$

We denote by $BR_i(\sigma_{-i}) \subseteq \Sigma_i$ the set of all best responses of player i to the strategy profile of opponents $\sigma_{-i} \in \Sigma_{-i}$.

Best Response – Example

Consider a game with the following payoffs of player 1:

	X	Y
A	2	0
B	0	2
C	1	1

- Player 1 (row) plays $\sigma_1 = (a(A), b(B), c(C))$.
- Player 2 (column) plays $(q(X), (1-q)(Y))$ (we write just q).

Compute $BR_1(q)$.

Rationalizability in Mixed Strategies (Two Players)

For simplicity, we temporarily switch to **two-player** setting $N = \{1, 2\}$.

Definition 34

A *(mixed) belief* of player $i \in \{1, 2\}$ is a mixed strategy σ_{-i} of his opponent.

(A general definition works with so called *correlated beliefs* that are arbitrary distributions on S_{-i}, the notion of the expected payoff needs to be adjusted, we are not going in this direction ....)

Assumption: *Any rational player with a belief σ_{-i} always plays a best response to σ_{-i}.*

Definition 35

A strategy $\sigma_i \in \Sigma_i$ of player $i \in \{1, 2\}$ is *never best response* if it is not a best response to any belief σ_{-i}.

No rational player plays a strategy that is never best response.

Rationalizability in Mixed Strategies (Two Players)

Define a sequence $R_i^0, R_i^1, R_i^2, \ldots$ of strategy sets of player i.
(Denote by G^k_{Rat} the game obtained from G by restricting the pure strategy sets to R_i^k, $i \in N$.)

1. Initialize $k = 0$ and $R_i^0 = S_i$ for each $i \in N$.
2. For all players $i \in N$: Let R_i^{k+1} be the set of all strategies of R_i^k that are *best responses to some (mixed) beliefs in* G^k_{Rat}.
3. Let $k := k + 1$ and go to 2.

We say that $s_i \in S_i$ is *rationalizable* if $s_i \in R_i^k$ for all $k = 0, 1, 2, \ldots$

Definition 36
A strategy profile $s = (s_1, \ldots, s_n) \in S$ is a *rationalizable equilibrium* if each s_i is rationalizable.

Rationalizability vs IESDS (Two Players)

	X	Y
A	3	0
B	0	3
C	1	1

- Player 1 (row) plays $\sigma_1 = (a(A), b(B), c(C))$
- player 2 (column) plays $(q(X), (1-q)(Y))$ (we write just q)

What strategies of player 1 are never best responses?

What strategies of player 1 are strictly dominated?

Observation: The set of strictly dominated strategies coincides with the set of never best responses!

... and this holds in general for two player games:

Theorem 37

Assume $N = \{1, 2\}$. A pure strategy s_i is never best response to any belief $\sigma_{-i} \in \Sigma_{-i}$ **iff** *s_i is strictly dominated by a strategy $\sigma_i \in \Sigma_i$.*

It follows that a strategy of S_i survives IESDS **iff** it is rationalizable.

(The theorem is true also for an arbitrary number of players but correlated beliefs need to be used.)

Mixed Nash Equilibrium

Definition 38

A mixed-strategy profile $\sigma^* = (\sigma^*_1, \ldots, \sigma^*_n) \in \Sigma$ is a (mixed) Nash equilibrium if σ^*_i is a best response to σ^*_{-i} for each $i \in N$, that is

$$u_i(\sigma^*_i, \sigma^*_{-i}) \geq u_i(\sigma_i, \sigma^*_{-i}) \quad \text{for all } \sigma_i \in \Sigma_i \text{ and all } i \in N$$

An interpretation: each σ^*_{-i} can be seen as a belief of player i against which he plays a best response σ^*_i.

Given a mixed strategy profile of opponents $\sigma_{-i} \in \Sigma_{-i}$, we denote by $BR_i(\sigma_{-i})$ the set of all $\sigma_i \in \Sigma_i$ that are best responses to σ_{-i}.

Then σ^* is a Nash equilibrium iff $\sigma^*_i \in BR_i(\sigma^*_{-i})$ for all $i \in N$.

Theorem 39 (Nash 1950)

Every finite game in strategic form has a Nash equilibrium.

This is THE fundamental theorem of game theory.

Example: Matching Pennies

	H	T
H	$1, -1$	$-1, 1$
T	$-1, 1$	$1, -1$

Player 1 (row) plays $(p(H), (1-p)(T))$ (we write just p) and player 2 (column) plays $(q(H), (1-q)(T))$ (we write q).

Compute all Nash equilibria.

What are the expected payoffs of playing pure strategies for player 1?

$$v_1(H, q) = 2q - 1 \text{ and } v_1(T, q) = 1 - 2q$$

Then
$v_1(p, q) = pv_1(H, q) + (1-p)v_1(T, q) = p(2q-1) + (1-p)(1-2q)$.

We obtain the best-response correspondence BR_1:

$$BR_1(q) = \begin{cases} p = 0 & \text{if } q < \frac{1}{2} \\ p \in [0, 1] & \text{if } q = \frac{1}{2} \\ p = 1 & \text{if } q > \frac{1}{2} \end{cases}$$

Example: Matching Pennies

	H	T
H	$1, -1$	$-1, 1$
T	$-1, 1$	$1, -1$

Player 1 (row) plays $(p(H), (1-p)(T))$ (we write just p) and player 2 (column) plays $(q(H), (1-q)(T))$ (we write q).

Compute all Nash equilibria.

Similarly for player 2 :

$$v_2(p, H) = 1 - 2p \text{ and } v_1(p, T) = 2p - 1$$

We obtain best-response relation BR_2:

$$BR_2(p) = \begin{cases} q = 1 & \text{if } p < \frac{1}{2} \\ q \in [0, 1] & \text{if } p = \frac{1}{2} \\ q = 0 & \text{if } p > \frac{1}{2} \end{cases}$$

The only "intersection" of BR_1 and BR_2 is the only Nash equilibrium $\sigma_1 = \sigma_2 = (\frac{1}{2}, \frac{1}{2})$.

Computing Mixed Nash Equilibria

Lemma 40

$\sigma^* = (\sigma_1^*, \ldots, \sigma_n^*) \in \Sigma$ is a Nash equilibrium **iff** *there exist $w_1, \ldots, w_n \in \mathbb{R}$ such that the following holds:*

- *For all $i \in N$ and all $s_i \in supp(\sigma_i^*)$ we have $u_i(s_i, \sigma_{-i}^*) = w_i$.*
- *For all $i \in N$ and all $s_i \notin supp(\sigma_i^*)$ we have $u_i(s_i, \sigma_{-i}^*) \leq w_i$.*

Here, the right hand side implies $u_i(\sigma^*) = w_i$.

Proof.

The fact that the right hand side implies $u_i(\sigma^*) = w_i$ follows immediately from Lemma 23:

$$u_i(\sigma^*) = \sum_{s_i \in S_i} \sigma^*(s_i) u_i(s_i, \sigma_{-i}^*) = \sum_{s_i \in supp(\sigma_i^*)} \sigma^*(s_i) u_i(s_i, \sigma_{-i}^*)$$

$$= \sum_{s_i \in supp(\sigma_i^*)} \sigma^*(s_i) w_i = w_i \sum_{s_i \in supp(\sigma_i^*)} \sigma^*(s_i) = w_i$$

Computing Mixed Nash Equilibria

Lemma 41
$\sigma^* = (\sigma_1^*, \ldots, \sigma_n^*) \in \Sigma$ is a Nash equilibrium **iff** *there exist $w_1, \ldots, w_n \in \mathbb{R}$ such that the following holds:*

- *For all $i \in N$ and all $s_i \in supp(\sigma_i^*)$ we have $u_i(s_i, \sigma_{-i}^*) = w_i$.*
- *For all $i \in N$ and all $s_i \notin supp(\sigma_i^*)$ we have $u_i(s_i, \sigma_{-i}^*) \leq w_i$.*

Here, the right hand side implies $u_i(\sigma^*) = w_i$.

Proof. (Cont.)
"$\Leftarrow$": Use the first equality of Lemma 23 to obtain for every $i \in N$ and every $\sigma_i' \in \Sigma_i$

$$u_i(\sigma_i', \sigma_{-i}^*) = \sum_{s_i \in S_i} \sigma_i'(s_i) u_i(s_i, \sigma_{-i}^*) \leq \sum_{s_i \in S_i} \sigma_i'(s_i) u_i(\sigma^*) = u_i(\sigma^*)$$

Thus σ^* is a Nash equilibrium.

Computing Mixed Nash Equilibria

Lemma 42

$\sigma^* = (\sigma_1^*, \ldots, \sigma_n^*) \in \Sigma$ is a Nash equilibrium **iff** *there exist $w_1, \ldots, w_n \in \mathbb{R}$ such that the following holds:*

- *For all $i \in N$ and all $s_i \in supp(\sigma_i^*)$ we have $u_i(s_i, \sigma_{-i}^*) = w_i$.*
- *For all $i \in N$ and all $s_i \notin supp(\sigma_i^*)$ we have $u_i(s_i, \sigma_{-i}^*) \leq w_i$.*

Here, the right hand side implies $u_i(\sigma^*) = w_i$.

Proof (Cont.)

Idea for "$\Rightarrow$": Let $w_i := u_i(\sigma^*)$.

Clearly, every $i \in N$ and $s_i \in S_i$ satisfy $u_i(s_i, \sigma_{-i}^*) \leq u_i(\sigma^*) = w_i$.

By Corollary 24, there is *at least one* $s_i \in supp(\sigma_i^*)$ satisfying $u_i(s_i, \sigma_{-i}^*) = u_i(\sigma^*) = w_i$.

Now if there is $s_i' \in supp(\sigma_i^*)$ such that

$$u_i(s_i', \sigma_{-i}^*) < u_i(\sigma^*) \quad (= u_i(s_i, \sigma_{-i}^*))$$

then increasing the probability $\sigma_i^*(s_i)$ and decreasing (in proportion) $\sigma_i^*(s_i')$ strictly increases of $u_i(\sigma^*)$, a contradiction with σ^* being NE.

Example: Matching Pennies

	H	T
H	$1, -1$	$-1, 1$
T	$-1, 1$	$1, -1$

Player 1 (row) plays $(p(H), (1-p)(T))$ (we write just p) and player 2 (column) plays $(q(H), (1-q)(T))$ (we write q).

Compute all Nash equilibria.

There are no pure strategy equilibria.

There are no equilibria where only player 1 randomizes:

Indeed, assume that (p, H) is such an equilibrium. Then by Lemma 42,

$$1 = u_1(H, H) = u_1(T, H) = -1$$

a contradiction. Also, (p, T) cannot be an equilibrium.

Similarly, there is no NE where only player 2 randomizes.

Example: Matching Pennies

	H	T
H	$1,-1$	$-1,1$
T	$-1,1$	$1,-1$

Player 1 (row) plays $(p(H),(1-p)(T))$ (we write just p) and player 2 (column) plays $(q(H),(1-q)(T))$ (we write q).

Compute all Nash equilibria.

Assume that both players randomize, i.e., $p,q \in (0,1)$.

The expected payoffs of playing pure strategies for player 1:

$$v_1(H,q) = 2q-1 \text{ and } v_1(T,q) = 1-2q$$

Similarly for player 2 :

$$v_2(p,H) = 1-2p \text{ and } v_1(p,T) = 2p-1$$

By Lemma 42, Nash equilibria must satisfy:

$$2q-1 = 1-2q \qquad \text{and} \qquad 1-2p = 2p-1$$

That is $p = q = \frac{1}{2}$ is the only Nash equilibrium.

Example: Battle of Sexes

	O	F
O	$2,1$	$0,0$
F	$0,0$	$1,2$

Player 1 (row) plays $(p(O), (1-p)(F))$ (we write just p) and player 2 (column) plays $(q(O), (1-q)(F))$ (we write q).

Compute all Nash equilibria.

There are two pure strategy equilibria $(2,1)$ and $(1,2)$, no Nash equilibrium where only one player randomizes.

Now assume that

- player 1 (row) plays $(p(H), (1-p)(T))$ (we write just p) and
- player 2 (column) plays $(q(H), (1-q)(T))$ (we write q)

where $p, q \in (0,1)$.

By Lemma 42, any Nash equilibrium must satisfy:

$$2q = 1 - q \qquad \text{and} \qquad p = 2(1-p)$$

This holds only for $q = \frac{1}{3}$ and $p = \frac{2}{3}$.

An Algorithm?

What did we do in the previous examples?

We went through all support combinations for both players.
(pure, one player mixing, both mixing)

For each pair of supports we tried to find equilibria in strategies with these supports.
(in Battle of Sexes: two pure, no equilibrium with just one player mixing, one equilibrium when both mixing)

Whenever one of the *supports* was non-singleton, we reduced computation of Nash equilibria to *linear equations*.

Support Enumeration (Idea)

Recall Lemma 42: $\sigma^* = (\sigma_1^*, \ldots, \sigma_n^*) \in \Sigma$ is a Nash equilibrium **iff** there exist $w_1, \ldots, w_n \in \mathbb{R}$ such that the following holds:

- For all $i \in N$ and all $s_i \in supp(\sigma_i^*)$ we have $u_i(s_i, \sigma_{-i}^*) = w_i$.
- For all $i \in N$ and all $s_i \notin supp(\sigma_i^*)$ we have $u_i(s_i, \sigma_{-i}^*) \leq w_i$.

Suppose that we somehow know the supports $supp(\sigma_1^*), \ldots, supp(\sigma_n^*)$ for some Nash equilibrium $\sigma_1^*, \ldots, \sigma_n^*$ (which itself is unknown to us).

Now we may consider all $\sigma_i^*(s_i)$'s and all w_i's as variables and use the above conditions to design a system of inequalities capturing Nash equilibria with the given support sets $supp(\sigma_1^*), \ldots, supp(\sigma_n^*)$.

Support Enumeration

To simplify notation, assume that for every i we have $S_i = \{1, \ldots, m_i\}$.
Then $\sigma_i(j)$ is the probability of the pure strategy j in the mixed strategy σ_i.

Fix supports $supp_i \subseteq S_i$ for every $i \in N$ and consider the following system of constraints with variables $\sigma_1(1), \ldots, \sigma_1(m_1), \ldots, \sigma_n(1), \ldots, \sigma_n(m_n), w_1, \ldots, w_n$:

1. For all $i \in N$ and all $k \in supp_i$ we have
$$(u_i(k, \sigma_{-i}) =) \qquad \sum_{s \in S \wedge s_i = k} \left(\prod_{j \neq i} \sigma_j(s_j) \right) u_i(s) = w_i$$
2. For all $i \in N$ and all $k \notin supp_i$ we have
$$(u_i(k, \sigma_{-i}) =) \qquad \sum_{s \in S \wedge s_i = k} \left(\prod_{j \neq i} \sigma_j(s_j) \right) u_i(s) \leq w_i$$
3. For all $i \in N$: $\sigma_i(1) + \cdots + \sigma_i(m_i) = 1$.
4. For all $i \in N$ and all $k \in supp_i$: $\sigma_i(k) \geq 0$.
5. For all $i \in N$ and all $k \notin supp_i$: $\sigma_i(k) = 0$.

Support Enumeration

Consider the system of constraints from the previous slide.

The following lemma follows immediately from Lemma 42.

Lemma 43
Let $\sigma^* \in \Sigma$ be a strategy profile.

- *If σ^* is a Nash equilibrium and $supp(\sigma_i^*) = supp_i$ for all $i \in N$, then assigning $\sigma_i(k) := \sigma_i^*(k)$ and $w_i := u_i(\sigma^*)$ solves the system.*
- *If $\sigma_i(k) := \sigma_i^*(k)$ and $w_i := u_i(\sigma^*)$ solves the system, then σ^* is a Nash equilibrium with $supp(\sigma_i^*) \subseteq supp_i$ for all $i \in N$.*

Support Enumeration (Two Players)

The constraints are *non-linear* in general, but *linear* for two player games! Let us stick to two players.

How to find $supp_1$ and $supp_2$? ... Just guess!

Input: A two-player strategic-form game G with strategy sets $S_1 = \{1, \ldots, m_1\}$ and $S_2 = \{1, \ldots, m_2\}$ and rational payoffs u_1, u_2.

Output: A Nash equilibrium σ^*.

Algorithm: For all possible $supp_1 \subseteq S_1$ and $supp_2 \subseteq S_2$:

- Check if the corresponding system of linear constraints (from the previous slide) has a feasible solution $\sigma^*, w_1^*, \ldots, w_n^*$.
- If so, STOP: the feasible solution σ^* is a Nash equilibrium satisfying $u_i(\sigma^*) = w_i^*$.

Question: How many possible subsets $supp_1, supp_2$ are there to try?
Answer: $2^{(m_1+m_2)}$

So, unfortunately, the algorithm requires worst-case exponential time.

Remarks on Support Enumeration

- The algorithm combined with Theorem 39 and properties of linear programming imply that every finite two-player game has a rational Nash equilibrium (furthermore, the rational numbers have polynomial representation in binary).
- The algorithm can be used to compute *all* Nash equilibria.

  (There are algorithms for computing (a finite representation of) a set of all feasible solutions of a given linear constraint system.)
- The algorithm can be used to compute "good" equilibria.

  For example, to find a Nash equilibrium maximizing the sum of all expected payoffs (the "social welfare") it suffices to solve the system of constraints while maximizing $w_1 + \cdots + w_n$. More precisely, the algorithm can be modified as follows:
  - Initialize $W := -\infty$ (W stores the current maximum welfare)
  - For all possible $supp_1 \subseteq S_1$ and $supp_2 \subseteq S_2$:
    - Find the maximum value $\max(\sum w_i)$ of $w_1 + \cdots + w_n$ so that the constraints are satisfiable (using linear programming).
    - Put $W := \max\{W, \max(\sum w_i)\}$.
  - Return W.

Remarks on Support Enumeration (Cont.)

Similar trick works for any notion of "good" NE that can be expressed using a linear objective function and (additional) linear constraints in variables $\sigma_i(j)$ and w_i.

(e.g., maximize payoff of player 1, minimize payoff of player 2 and keep probability of playing the strategy 1 below $1/2$, etc.)

Complexity Results – (Two Players)

Theorem 44

All the following problems are NP-complete: Given a two-player game in strategic form, does it have

1. *a NE in which player 1 has utility at least a given amount v ?*
2. *a NE in which the sum of expected payoffs of the two players is at least a given amount v ?*
3. *a NE with a support of size greater than a given number?*
4. *a NE whose support contains a given strategy s ?*
5. *a NE whose support does not contain a given strategy s ?*
6. *....*

Membership to NP follows from the support enumeration:
For example, for 1., it suffices to guess supports $supp_1, supp_2$ and add $w_1 \geq v$ to the constraints; the resulting NE σ^* satisfies $u_1(\sigma^*) \geq v$.

Complexity Results (Two Players)

Theorem 45
All the following problems are NP-complete: Given a two-player game in strategic form, does it have

1. *a NE in which player 1 has utility at least a given amount v ?*
2. *a NE in which the sum of expected payoffs of the two players is at least a given amount v ?*
3. *a NE with a support of size greater than a given number?*
4. *a NE whose support contains a given strategy s ?*
5. *a NE whose support does not contain a given strategy s ?*
6. *....*

NP-hardness can be proved using reduction from SAT
(The reduction is not difficult but we are not going into it.

It is presented in "New Complexity Results about Nash Equilibria" by

V. Conitzer and T. Sandholm (pages 6–8))

The Reduction (It's Short and Sweet)

Definition 4 *Let ϕ be a Boolean formula in conjunctive normal form (representing a SAT instance). Let V be its set of variables (with $|V| = n$), L the set of corresponding literals (a positive and a negative one for each variable[6]), and C its set of clauses. The function $v : L \to V$ gives the variable corresponding to a literal,* e.g., $v(x_1) = v(-x_1) = x_1$. *We define $G_\epsilon(\phi)$ to be the following finite symmetric 2-player game in normal form. Let $\Sigma = \Sigma_1 = \Sigma_2 = L \cup V \cup C \cup \{f\}$. Let the utility functions be*

- $u_1(l^1, l^2) = u_2(l^2, l^1) = n - 1$ *for all* $l^1, l^2 \in L$ *with* $l^1 \neq -l^2$;
- $u_1(l, -l) = u_2(-l, l) = n - 4$ *for all* $l \in L$;
- $u_1(l, x) = u_2(x, l) = n - 4$ *for all* $l \in L$, $x \in \Sigma - L - \{f\}$;
- $u_1(v, l) = u_2(l, v) = n$ *for all* $v \in V$, $l \in L$ *with* $v(l) \neq v$;
- $u_1(v, l) = u_2(l, v) = 0$ *for all* $v \in V$, $l \in L$ *with* $v(l) = v$;
- $u_1(v, x) = u_2(x, v) = n - 4$ *for all* $v \in V$, $x \in \Sigma - L - \{f\}$;
- $u_1(c, l) = u_2(l, c) = n$ *for all* $c \in C$, $l \in L$ *with* $l \notin c$;
- $u_1(c, l) = u_2(l, c) = 0$ *for all* $c \in C$, $l \in L$ *with* $l \in c$;
- $u_1(c, x) = u_2(x, c) = n - 4$ *for all* $c \in C$, $x \in \Sigma - L - \{f\}$;
- $u_1(x, f) = u_2(f, x) = 0$ *for all* $x \in \Sigma - \{f\}$;
- $u_1(f, f) = u_2(f, f) = \epsilon$;
- $u_1(f, x) = u_2(x, f) = n - 1$ *for all* $x \in \Sigma - \{f\}$.

Theorem 1 *If $(l_1, l_2, \ldots, l_n)$ (where $v(l_i) = x_i$) satisfies ϕ, then there is a Nash equilibrium of $G_\epsilon(\phi)$ where both players play l_i with probability $\frac{1}{n}$, with expected utility $n - 1$ for each player. The only other Nash equilibrium is the one where both players play f, and receive expected utility ϵ each.*

... But What is The Exact Complexity of *Computing* Nash Equilibria in Two Player Games?

Let us concentrate on the problem of computing one Nash equilibrium (sometimes called the *sample equilibrium problem*).

As the class NP consists of decision problems, it cannot be directly used to characterize complexity of the sample equilibrium problem.

We use complexity classes of *function problems* such as FP, FNP, etc.

The support enumeration gives a deterministic algorithm which runs in exponential time. Can we do better?

In what follows we show that

- the sample equilibrium problem can be solved in polynomial time for zero-sum two-player games,

  (Using a beautiful characterization of all Nash equilibria)
- the sample equilibrium problem belongs to the complexity class PPAD (which is a subclass of FNP) for two-player games.

  (... to be defined later)

MaxMin

Is there a better characterization of Nash equilibria than Lemma 42 ?

Definition 46
$\sigma_i^* \in \Sigma_i$ is a *maxmin* strategy of player i if

$$\sigma_i^* \in \operatorname*{argmax}_{\sigma_i \in \Sigma_i} \min_{\sigma_{-i} \in \Sigma_{-i}} u_i(\sigma_i, \sigma_{-i})$$

(Intuitively, a *maxmin* strategy σ_1^* maximizes player 1's worst-case payoff in the situation where player 2 strives to cause the greatest harm to player 1.)

(Since u_i is continuous and Σ_{-i} compact, $\min_{\sigma_{-i} \in \Sigma_{-i}} u_i(\sigma_i, \sigma_{-i})$ is well defined and continuous on Σ_i, which implies that there is at least one maxmin strategy.)

MaxMin

Lemma 47

σ_i^* is maxmin iff

$$\sigma_i^* \in \operatorname*{argmax}_{\sigma_i \in \Sigma_i} \min_{s_{-i} \in S_{-i}} u_i(\sigma_i, s_{-i})$$

Proof.

By Corollary 24, for every $\sigma \in \Sigma$ we have $u_i(\sigma_i, \sigma_{-i}) \geq u_i(\sigma_i, s_{-i})$ for some $s_{-i} \in S_{-i}$.

Thus $\min_{\sigma_{-i} \in \Sigma_{-i}} u_i(\sigma_i, \sigma_{-i}) = \min_{s_{-i} \in S_{-i}} u_i(\sigma_i, s_{-i})$. Hence,

$$\operatorname*{argmax}_{\sigma_i \in \Sigma_i} \min_{\sigma_{-i} \in \Sigma_{-i}} u_i(\sigma_i, \sigma_{-i}) = \operatorname*{argmax}_{\sigma_i \in \Sigma_i} \min_{s_{-i} \in S_{-i}} u_i(\sigma_i, s_{-i})$$

□

Question: Assume a strategy profile where both players play their maxmin strategies? Does it have to be a Nash equilibrium?

Zero-Sum Games: von Neumann's Theorem

Assume that G is zero sum, i.e., $u_1 = -u_2$.

Then $\sigma_2^* \in \Sigma_2$ is maxmin of player 2 **iff**

$$\sigma_2^* \in \operatorname*{argmin}_{\sigma_2 \in \Sigma_2} \max_{\sigma_1 \in \Sigma_1} u_1(\sigma_1, \sigma_2) \quad \left(= \operatorname*{argmin}_{\sigma_2 \in \Sigma_2} \max_{s_1 \in S_1} u_1(s_1, \sigma_2)\right)$$

(Intuitively, maxmin of player 2 minimizes the payoff of player 1 when player 1 plays his best responses. Such strategy of player 2 is often called minmax.)

Theorem 48 (von Neumann)

Assume a two-player **zero-sum** *game. Then*

$$\max_{\sigma_1 \in \Sigma_1} \min_{\sigma_2 \in \Sigma_2} u_1(\sigma_1, \sigma_2) = \min_{\sigma_2 \in \Sigma_2} \max_{\sigma_1 \in \Sigma_1} u_1(\sigma_1, \sigma_2)$$

Morever, $\sigma^* = (\sigma_1^*, \sigma_2^*) \in \Sigma$ is a Nash equilibrium **iff** *both σ_1^* and σ_2^* are maxmin.*

So to compute a Nash equilibrium it suffices to compute (arbitrary) maxmin strategies for both players.

Proof of Theorem 48 (Homework)

Homework: Prove von Neumann's Theorem in 4 easy steps:

1. Prove this inequality:

$$\max_{\sigma_1 \in \Sigma_1} \min_{\sigma_2 \in \Sigma_2} u_1(\sigma_1, \sigma_2) \leq \min_{\sigma_2 \in \Sigma_2} \max_{\sigma_1 \in \Sigma_1} u_1(\sigma_1, \sigma_2)$$

2. Prove that (σ_1^*, σ_2^*) is a Nash equilibrium iff

$$\min_{\sigma_2 \in \Sigma_2} u_1(\sigma_1^*, \sigma_2) \geq u_1(\sigma_1^*, \sigma_2^*) \geq \max_{\sigma_1 \in \Sigma_1} u_1(\sigma_1, \sigma_2^*)$$

   Hint: One of the inequalities is trivial and the other one almost.

3. Use 1. and 2. together with Theorem 39 to prove

$$\max_{\sigma_1 \in \Sigma_1} \min_{\sigma_2 \in \Sigma_2} u_1(\sigma_1, \sigma_2) \geq \min_{\sigma_2 \in \Sigma_2} \max_{\sigma_1 \in \Sigma_1} u_1(\sigma_1, \sigma_2)$$

4. Use the above to prove the rest of the theorem.
   Hint: Use the characterization of NE from 2., do not forget that you already have $\max_{\sigma_1 \in \Sigma_1} \min_{\sigma_2 \in \Sigma_2} u_1(\sigma_1, \sigma_2) = \min_{\sigma_2 \in \Sigma_2} \max_{\sigma_1 \in \Sigma_1} u_1(\sigma_1, \sigma_2)$
   You may already have proved one of the implications when proving 3.

Zero-Sum Two-Player Games – Computing NE

Assume $S_1 = \{1, \ldots, m_1\}$ and $S_2 = \{1, \ldots, m_2\}$.

We want to compute

$$\sigma_1^* \in \operatorname*{argmax}_{\sigma_1 \in \Sigma_1} \min_{\ell \in S_2} u_1(\sigma_1, \ell)$$

Consider a linear program with variables $\sigma_1(1), \ldots, \sigma_1(m_1), v$:

$$\begin{array}{rll} \textbf{maximize:} & v & \\ \textbf{subject to:} & \displaystyle\sum_{k=1}^{m_1} \sigma_1(k) \cdot u_1(k, \ell) \geq v & \ell = 1, \ldots, m_2 \\ & \displaystyle\sum_{k=1}^{m_1} \sigma_1(k) = 1 & \\ & \sigma_1(k) \geq 0 & k = 1, \ldots, m_1 \end{array}$$

Lemma 49
$\sigma_1^* \in \operatorname{argmax}_{\sigma_1 \in \Sigma_1} \min_{\ell \in S_2} u_1(\sigma_1, \ell)$ **iff** *assigning* $\sigma_1(k) := \sigma_1^*(k)$ *and* $v := \min_{\ell \in S_2} u_1(\sigma_1^*, \ell)$ *gives an optimal solution.*

Zero-Sum Two-Player Games – Computing NE

Summary:

- We have reduced computation of NE to computation of maxmin strategies for both players.
- Maxmin strategies can be computed using linear programming in polynomial time.
- That is, Nash equilibria in zero-sum two-player games can be computed in polynomial time.

IESDS vs Rationalizability Revisited

We get Theorem 37 as a simple corollary of Theorem 48.

Let s_1^* be a strategy of player 1. Consider a zero-sum game $G' = (\{1,2\}, (S'_1, S'_2), (u'_1, u'_2))$ where

- $S'_1 = S_1 \setminus \{s_1^*\}$ and $S'_2 = S_2$,
- $u'_1(s_1, s_2) = u_1(s_1, s_2) - u_1(s_1^*, s_2)$ and $u'_2(s_1, s_2) = u_1(s_1^*, s_2) - u_1(s_1, s_2)$.

Now s_1^* is never best resp. in G **iff**
for every $\sigma_2 \in \Sigma_2$ exists $\sigma_1 \in \Sigma_1 : u_1(\sigma_1, \sigma_2) - u_1(s_1^*, \sigma_2) > 0$ **iff**
for every $\sigma_2 \in \Sigma_2$ exists $s_1 \in S_1 : u_1(s_1, \sigma_2) - u_1(s_1^*, \sigma_2) > 0$ **iff**
$\min_{\sigma_2 \in \Sigma_2} \max_{s_1 \in S_1} u'_1(s_1, \sigma_2) > 0$ **iff**
$\min_{\sigma_2 \in \Sigma_2} \max_{\sigma_1 \in \Sigma_1} u'_1(\sigma_1, \sigma_2) > 0$ **iff**
$\max_{\sigma_1 \in \Sigma_1} \min_{\sigma_2 \in \Sigma_2} u'_1(\sigma_1, \sigma_2) > 0$ **iff**
there is $\sigma_1 \in \Sigma_1$ such that for all $\sigma_2 \in \Sigma_2$ we have
$0 < u'_1(\sigma_1, \sigma_2) = u_1(\sigma_1, \sigma_2) - u_1(s_1^*, \sigma_2)$ **iff**
s_1^* is strictly dominated (by σ_1) in G.