Equilibres de Nash

Wiesław Zielonka

www.liafa.jussieu.fr/~zielonka

mail: zielonka@liafa.jussieu.fr

LIAFA, Universit Denis Diderot

Strategic-form games

Strategic game (or normal form game)

$$\Gamma = (N, \{S_i\}_{i \in N}, \{u_i\}_{i \in N}).$$

N – the set of players,
S_i — the set of strategies of player $i \in I$,
$u_i : S \to \mathbb{R}$ — the utility payoff mapping for player $i \in N$.

Here

$$S = \prod_{i \in I} S_i$$

the set of strategy profiles, i.e. all possible strategy combinations obtained when each player chooses his strategy.

Intuitively, if each player fixes his strategy then their choices determine the play outcome and then

$$u_i(s) < u_i(s') \quad \text{for } s, s' \in S$$

means that player i prefers the outcome resulting from the strategy profile s' to the outcome resulting from the profile s.

Another interpretation, $u_i(s)$ is the payment received by player i if the strategy profile is $s \in S$. (If $u_i(s) < 0$ then player i loses $|u_i(s)|$.)

The aim of each player : maximise his gain.

Notation

Let $s = (s_1, \ldots, s_i, \ldots, s_n) \in S$.

Then

$$s_{-i} = (s_1, \ldots, s_{i-1}, s_{i+1}, \ldots, s_n) \quad \text{and} \quad S_{-i} = \prod_{j \in N \setminus \{i\}} S_j$$

and for $s'_i \in S_i$

$$(s_{-i}, s'_i) = (s_1, \ldots, s_{i-1}, s'_i, s_{i+1}, \ldots, s_n).$$

The prisoner's dilemma

I \ II	don't confess	confess
don't confess	$-2, -2$	$-10, -1$
confess	$-1, -10$	$-5, -5$

Two suspects are arrested for allegedly engaging in a serious crime and are put into separate cells. They can either cooperate with the police or not. The payoff represent the number of years they will spend in prison in each case.

If both refuse to cooperate then they will be convicted of minor offence for two years of prison. If one cooperates and the other not then the first can be used as a witness against the other who will receive the sentence of ten years. If both cooperate then they get five years each.

There is only one plausible outcome (cooperate,cooperate). This is because cooperating is each player's best strategy *regardless of what the other player does*.

Dominant and Dominated Strategies

A strategy $s_i \in S_i$ is a strictly dominant strategy for player i in game $\Gamma = (N, \{S_i\}_{i \in N}, \{u_i\}_{i \in N})$ if for all $s_i^* \neq s_i$, we have

$$u_i(s_i, s_{-i}) > u_i(s_i^*, s_{-i})$$

for all $s_{-i} \in S_{-i}$.

A strategy $s_i \in S_i$ is a strictly dominated for player i in game $\Gamma = (N, \{S_i\}_{i \in N}, \{u_i\}_{i \in N})$ if there exists another strategy $s_i^\sharp \in S_i$ such that for all $s_{-i} \in S_{-i}$

$$u_i(s_i^\sharp, s_{-i}) > u_i(s_i, s_{-i}) \ .$$

A strategy $s_i \in S_i$ is weakly dominated for player i in game $\Gamma = (N, \{S_i\}_{i \in N}, \{u_i\}_{i \in N})$ if there exists another strategy $s_i^\sharp \in S_i$ such that for all $s_{-i} \in S_{-i}$

$$u_i(s_i^\sharp, s_{-i}) \geq u_i(s_i, s_{-i})$$

with strict inequality for some s_{-i}.

In this case we say that $s_i^\sharp$ weakly dominates s_i.

A strategy is weakly dominant for player i if it weakly dominates every other strategy in S_i.

Iterated deletion of strictly dominated strategies

A variant of prisoner's dilemma.

I \ II	don't confess	confess
don't confess	$0, -2$	$-10, -1$
confess	$-1, -10$	$-5, -5$

Exercise

Show that the order of deletion does not matter for the set of strategies surviving a process of iterated deletion of strictly dominated strategies.

Nash equilibrium

Definition. *The strategy profile $s^* \in S$ is a Nash equilibrium if for each player $i \in I$ and each strategy $s_i \in S_i$,*

$$u_i(s^*_{-i}, s_i) \leq u_i(s^*) \ .$$

Examples of games

Battle of the Sexes

Marc \ Julie	football	shopping
football	3,1	0,0
shopping	0,0	1,3

Marc et Julie decide what to do on Saturday afternoon. Neither of them will derive any pleasure from being without the other but Marc prefers football match while Julie prefers to go to a shopping center.

A coordination game

Marc \ Julie	Mozart	Mahler
Mozart	2,2	0,0
Mahler	0,0	1,1

Marc and Julie wish to go out together but now they have the same preferences.. Both prefer Mozart but (Mahler,Mahler) is also a Nash equilibrium! Nash equilibrium can be highly inefficient!

Hawk and Dowe

I \ II	Dowe	Hawk
Dowe	3, 3	1, 4
Hawk	4, 1	0, 0

Two animals fight over a prey. Nash equilibria (Hawk,Dowe) and (Dowe,Hawk).

Matching pennies

I \ II	head	tail
head	$1, -1$	$-1, 1$
tail	$-1, 1$	$1, -1$

Two players choose either head or tail.. If the choices differ then I pays to II 1 euro, otherwise II pays to I 1 euro.

No Nash equilibrium if only deterministic (pure) strategies admitted.

An auction

An object is to be attributed to one of n players: $N = \{1, \dots, n\}$.

The object value for player i is v_i, to simplify let us assume that $v_1 > v_2 > \dots > v_n > 0$.

Players submit simultaneously their bids (nonnegative numbers). The object is given to the player submitting the highest bid, if several players announce the same highest bid the player with the greatest index receives the object. The winner pays the amount of his bid (first price auction).

How to formulate this game as a strategic game?

Strategy of player i = his bid.

If i obtains the object and pays p_i then his gain is $v_i - p_i$.

If i does not get the object his gain is 0.

Second price auction

The winner is determined as previously but he pays the highest bid among those submitted by the players that do not win.

A war of attrition

Two players dispute an object, the valuer of the object for player i is $v_i > 0$.

Time is a continuous variable starting at 0 and going to infinity. At any moment each player can decide to concede the object to his adversary i which gains in this way v_i, this terminates the game (if both decide to concede the object the same moment then the object is split and therefore player i receives $v_i/2$).

Time is money, for each time unit until the end of the game each player looses 1 euro.

Exercise.
Formulate this game as a strategic game.
Show that in a Nash equilibrium one player concedes immediately the object.

A location game

n players choose if they will present themselves as candidates at the elections and if so which political position to take (they are opportunists without real political convictions, the only thing that counts for them is the result of the election).

There is a continuum of citizens, each with his or her well-defined favorite political opinion. The distribution of favorite opinions is given by a density function $f > 0$, the opinions from left to right are represented by the interval $[0; 1]$, if $0 \leq a < b \leq 1$ then

$$\int_a^b f(x)dx$$

gives the number of citizens with opinions in the interval $[a; b]$.

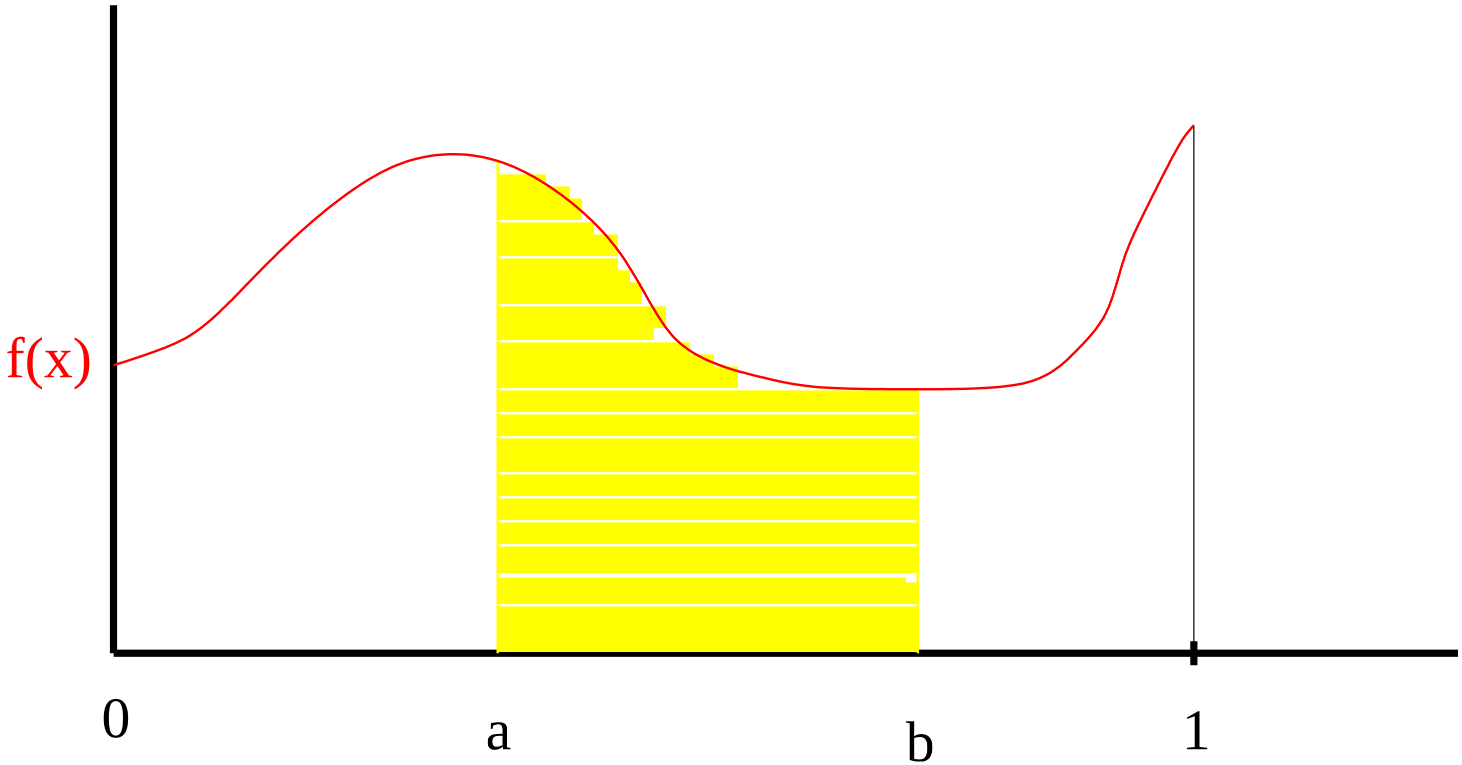
f(x)
0
a
b
1

A candidate attracts all the votes that are closer to his position then to the position of any other candidate. If k candidates choose the same position then each receives the fraction $1/k$ of the votes that this position attracts.

Example: two candidates J and C with their respective positions.

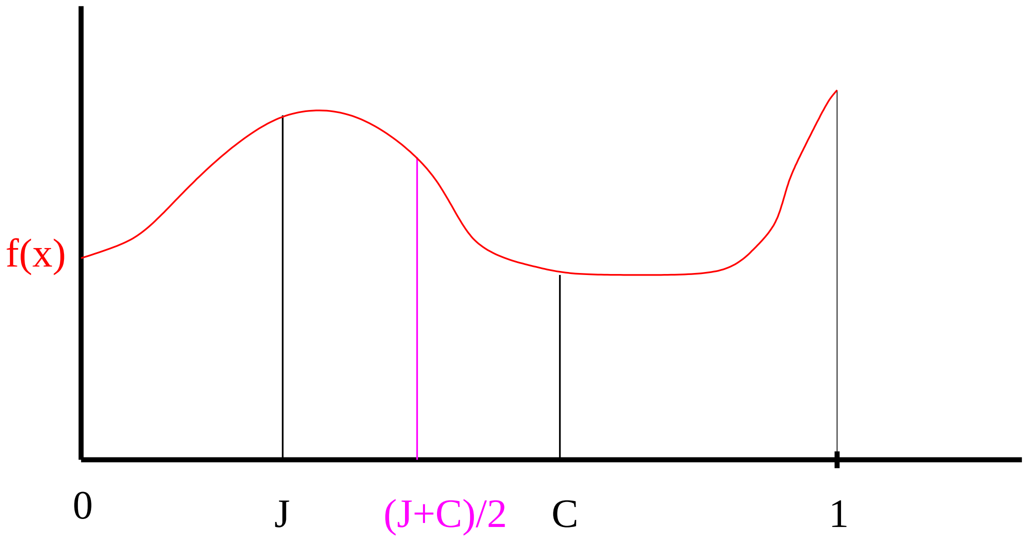

The candidate J receives all votes from 0 to $(J + C)/2$.
The candidate C receives all votes from $(J + C)/2$ to 1.

The winner is the candidate who receives the majority of votes.

Each person prefers to be the unique winner than to tie for the first place, prefers to tie for the first place than to stay out of the competition, prefers to stay out of the competition than to enter and loose.

The utility mapping can give for example the gain $\frac{1}{k}$ for each of k winners, 0 for a player that prefers to stay out and -1 for a players that enters the competition but looses.

Exercise
Find Nash equilibria for $n = 2$.

Show that there is no Nash equilibrium for $n = 3$.

Mixed strategies

Suppose that the set S_i of strategies of player i is finite. We call strategies from S_i - pure strategies.

A mixed strategy for player i is a probability distribution over S_i, i.e. a function

$$\sigma_i : S_i \longrightarrow [0, 1]$$

such that

$$\sum_{s_i \in S_i} \sigma_i(s_i) = 1.$$

If each player $i \in N$ uses some mixed strategy σ_i then the utility expectation for player k in the game $\Gamma = (N, \{S_i\}_{i \in N}, \{u_i\}_{i \in N})$ is

$$\sum_{(s_1,\ldots,s_n) \in S} \sigma_1(s_1) \cdots \sigma_n(s_n) u_k(s_1, \ldots, s_n) \ .$$

If mixed strategies are allowed then the utility expectation is interpreted as the player's gain.

Let

$$\mathcal{D}(S_i)$$

be the set of all mixed strategies for player i. Formally, we have replaced here the game $\Gamma = (N, \{S_i\}_{i\in N}, \{u_i\}_{i\in N})$ by a new game

$$\Gamma^\sharp = (N, \{\mathcal{D}(S_i)\}_{i\in N}, \{u_i\}_{i\in N})$$

where for any mixed strategy profile $\sigma \in \prod_{i\in N} \mathcal{D}(S_i)$, the payoff for player i is given by

$$u_i(\sigma) = \sum_{(s_1,\ldots,s_n)\in S} \sigma_1(s_1)\cdots\sigma_n(s_n)u_k(s_1,\ldots,s_n) \ .$$

Domination in mixed strategies

1 \ 2	L	R
U	10,1	0,4
M	4,2	4,3
D	0,5	10,2

The mixed strategy $\frac{1}{2}U + \frac{1}{2}D$ strictly dominates M.

No pure strategy of player 1 is dominated by any other pure strategy.

Conclusion:
Mixed strategies allow to eliminate more pure strictly dominated strategies.

Existence of Nash equilibria in mixed strategies

Theorem. [John Nash (1951)] *If, for all $i \in N$, the set S_i of pure strategies for player i is finite then there exists a mixed strategy profile $\sigma \in \prod_{i \in N} \mathcal{D}(S_i)$ which is a Nash equilibrium.*

Proof. For each $s_{-i} \in S_{-i}$ define $B_i(s_i)$ to be the set of player's i best strategies given the strategy profile s_{-i} for the other players:

$$B_i(s_{-i}) = \{s_i \in S_i \mid \forall s_i^* \in S_i, u_i(s_{-i}, s_i) \geq u_i(s_{-i}, s_i^*)\}$$

We call the set-valued function the best-response function of player i.

A Nash equilibrium is a strategy profile $s \in S$ such that

$$s_i \in B_i(s_{-i}) \quad \text{for all } i \in N \ .$$

Theorem. [Kakutani's fixed point theorem] *Let* X *be a compact convex subset of* $\mathbb{R}^n$ *and let* $f : X \to \mathcal{P}(X)$ *be a set valued function such that*

- *for all* $x \in X$*, the set* $f(x)$ *is nonempty and convex,*
- *the graph of* f *is closed (i.e. for all sequences* $\{x_k\}$*,* $\{y_k\}$ *such that* $y_k \in f(x_k)$ *for all* k*, if* $x_k \to x$ *and* $y_k \to y$ *then* $y \in f(x)$*).*

Then there exists $x^* \in X$ *such that* $x^* \in f(x^*)$*.*

Now take as X the set $\prod_{i \in N} \mathcal{D}(S_i)$ of all mixed strategy profiles and the mapping f the mapping $B = (B_1, \ldots, B_n)$ where B_i is the best response mapping for player i (best response in mixed strategies). Then check that the conditions of Kakutani's fixed point theorem are satisfied.

□

Calculating Nash equilibria in mixed strategies

If we have three or more players then it is possible to have a game such that

- each player has a finite number of pure strategies,
- the payoff values for pure strategies profiles are all rational numbers,
- there exists Nash equilibria only in mixed strategies and all such equilibria are non-rational (i.e. have non-rational probability distributions and non-rational payoffs).

For two players (bimatrix games) if all entries in both matrices are rational then there are Nash equilibria with rational payoffs for both players. However, even in the case of bimatrix games best known algorithm of calculating Nash equilibria has an exponential complexity.

Open problem. Find an efficient algorithm calculating equilibria for bimatrix games.

3-players games are PPAD complete (Daskalakis, Papadimitriou).

Practical calculations of Nash equilibria

For any mixed strategy $\sigma_i \in \mathcal{D}(S_i)$ of player i the support of σ_i is the set

$$\mathrm{support}(\sigma_i) = \{s_i \in S_i \mid \sigma_i(s_i) > 0\}$$

Theorem. *Let σ be a Nash equilibrium and $i \in N$ a player. Then*

$$\forall s_i \in \text{support}(\sigma_i), \quad u_i(\sigma_{-i}, s_i) = u_i(\sigma_{-i}, \sigma_i) \tag{1}$$

Moreover for all pure strategies $s_i^* \in S_i$,

$$u_i(\sigma_{-i}, s_i^*) \leq u_i(\sigma_{-i}, \sigma_i) \tag{2}$$

Proof. Eq. (2) follows immediately from the definition of a Nash equilibrium.

To prove (1) note that

$$u_i(\sigma) = u_i(\sigma_{-i}, \sigma_i) = \sum_{s_i \in \text{support}(\sigma_i)} \sigma_i(s_i) u_i(\sigma_{-i}, s_i)$$

Since $\sum_{s_i \in \text{support}(\sigma_i)} \sigma_i(s_i) = 1$ and

$$\forall s_i \in \text{support}(\sigma_i), 0 < \sigma_i(s_i) \leq 1$$

if (1) does not hold then we could eliminate from the support the pure strategy s_i for which

$$u_i(\sigma_{-i}, s_i) < u_i(\sigma_{-i}, \sigma_i)$$

and increment the probability of some other pure strategy from the support by $\sigma_i(s_i)$. This would strictly increase the payoff of player i showing that σ is not a Nash equilibrium profile.

□

Battle of the Sexes revisited

Marc \ Julie	football	shopping
football	3,1	0,0
shopping	0,0	1,3

Let us look at mixed strategies for Marc and Julie with the support $\{\text{football}, \text{shopping}\}$. Note in the sequel

$$F = \text{football} \quad \text{and} \quad S = \text{shopping}.$$

Let us take a mixed strategy σ_M for Marc with and σ_J for Julie such that

$$\text{support}(\sigma_M) = \{\text{football}, \text{shopping}\}$$

and

$$\text{support}(\sigma_J) = \{\text{football}, \text{shopping}\}.$$

Then Marc wins

$$u_M(\sigma_M, \sigma_J) = 3 \cdot \sigma_M(F) \cdot \sigma_J(F) + \sigma_M(S) \cdot \sigma_J(S)$$

while Julie wins

$$u_J(\sigma_M, \sigma_J) = \sigma_M(F) \cdot \sigma_J(F) + 3 \cdot \sigma_M(S) \cdot \sigma_J(S).$$

By the preceding theorem

$$\sigma_M(F) = u_J(\sigma_M, F) = u_J(\sigma_M, S) = 3 \cdot \sigma_M(S) = 3 \cdot (1 - \sigma_M(F)).$$

Solving this we get

$$\sigma_M(F) = \frac{3}{4} \quad \text{and} \quad \sigma_M(S) = \frac{1}{4}.$$

Thus

$$u_J(\sigma_M, \sigma_J) = \frac{3}{4} \cdot \sigma_J(F) + 3 \cdot \frac{1}{4} \cdot \sigma_J(S) = \frac{3}{4} \cdot (\sigma_J(F) + \sigma_J(S)) = \frac{3}{4}$$

In a similar way we get

$$\sigma_J(F) = \frac{1}{4} \quad \text{and} \quad \sigma_J(S) = \frac{3}{4}$$

and

$$u_M(\sigma_M, \sigma_J) = \frac{3}{4}.$$

Exercise. Show that (σ_M, σ_J) is effectively a Nash equilibrium.

exercise. Compute Nash equilibria for the following bimatrix game

I \ II	L	M	R
T	7,2	2,7	3,6
B	2,7	7,2	4,5

Zero-sum two-person games

A two-person game Γ, $N = \{1, 2\}$, is a zero-sum game if for each strategy profile s

$$u_2(s) = -u_1(s).$$

A zero-sum game can be noted as $\Gamma = (X, Y, u)$, where X and Y the set of strategies for players 1 and 2 and u is the payoff mapping of player 1.

Matrix games

are zero-sum games where both players have a finite number of pure strategies.

Morra game

Each player hides either 1 or 2 euros and tries to guess the sum hidden by the adversary.

The player guessing correctly wins the sum of money hidden by both players. This amount is payed by his adversary.

Each player has 4 strategies: $[i, j]$, $i, j = 1, 2$, where i is the amount of money that he has hidden and j is his guess concerning the amount of money hidden by his adversary.

Marc \ Julie	[1,1]	[1,2]	[2,1]	[2,2]
[1,1]	0	2	−3	0
[1,2]	−2	0	0	3
[2,1]	3	0	0	−4
[2,2]	0	−3	4	0

Marc can assure for himself the payoff of -2. Julie can assure that her loss does not exceed 2.

In general in zero-sum games player 1 can always secure for himself the payoff of

$$\max_{x \in X} \min_{y \in Y} u(x, y)$$

while player 2 can assure that his loss will not be greater than

$$\min_{y \in Y} \max_{x \in X} u(x, y)$$

(Replace $\min$, $\max$ by $\inf$, $\sup$ if X, Y are infinite).

Always

$$\underline{\mathrm{val}} = \max_{x \in X} \min_{y \in Y} u(x, y) \leq \min_{y \in Y} \max_{x \in X} u(x, y) = \overline{\mathrm{val}}$$

If

$$\underline{\mathrm{val}} = \overline{\mathrm{val}}$$

then we say that the zero-sum game Γ has a value

$$\mathrm{val} = \underline{\mathrm{val}} = \overline{\mathrm{val}}.$$

$\underline{\mathrm{val}}$ and $\overline{\mathrm{val}}$ are called respectively the lower and the upper value of a zero-sum game.

$(x, y) \in X \times Y$ is a saddle point if

$$u(x', y) \leq u(x, y) \leq u(x, y') \quad \text{for each } x' \in X \text{ and } y' \in Y$$

In this case we say that x and y are optimal strategies.

Exercise (Exchangeability of optimal strategies for zero-sum games)
If (x, y) and (x^*, y^*) are optimal strategies in a zero-sum game then (x, y^*) and (x^*, y) are also optimal and $u(x, y) = u(x^*, y^*)$.

Minimax theorem

Let $U = [u_{ij}]_{1 \le i \le m, 1 \le j \le n}$ a matrix game.

A mixed strategy of player 1 is then a row vector $\mathbf{X}$ of length m while a mixed strategy of player 2 is given by a column vector of length n $\mathbf{Y}$, all elements of vectors $\mathbf{X}$ and $\mathbf{Y}$ are nonnegative and $\sum_{i=1}^{m} \mathbf{X}_i = \sum_{i=1}^{n} \mathbf{Y}_i = 1$.

Theorem. [minimax, John von Neumann 1928]

$$\max_X \min_Y XUY = \min_Y \max_X XUY$$

where X and Y are taken from the sets of mixed strategies of players 1 and 2.

Conclusion, matrix games have always optimal strategies if we admit mixed strategies.

Matrix games and linear programming

If player 1 uses a mixed strategy X then the best response of player 2 is to use Y that minimizes XUY.

Now we should note that

$$\min_Y XUY = \min_j \sum_{i=1}^{m} X_i U_{ij} \tag{3}$$

To prove formally (3), let j be the column for which the right hand side of (3) is minimal and let us set

$$\alpha := \sum_{i=1}^{m} X_i U_{ij}$$

Then

$$XUY = \sum_{i=1}^{m} \sum_{k=1}^{n} X_i U_{ik} Y_k = \sum_{k=1}^{n} Y_k (\sum_{i=1}^{m} X_i U_{ik}) \geq$$

$$= \sum_{k=1}^{n} Y_k \cdot \alpha = \alpha \cdot \sum_{k=1}^{n} Y_k = \alpha$$

On the other hand if $Y := e_j$ is such that $Y_j = 1$ and $Y_k = 0$ for all $k \neq j$ then we get $XUe_j = \alpha$. This ends the proof of (3).

Eq. (3) implies that finding the row player's optimal strategy X reduces to the following problem:

maximize $\min_j \sum_{i=1}^{m} X_i U_{ij}$

subject to: $\sum_{i=1}^{m} X_i = 1$ and
$X_i \geq 0$ for all i.

The last problem is equivalent to the following linear programming problem:

maximize $\qquad z$

subject to: $\quad z - \sum_{i=1}^{m} X_i U_{ij} \leq 0$, $j = 1, \ldots, n$

$\sum_{i=1}^{m} X_i = 1$ et

$X_i \geq 0$ for all i.

To see that equivalence note that the first set of constraints can be written as $z \leq \sum_{i=1}^{m} X_i U_{ij}$, for all j, and these are the only constraints for z. Thus the optimal (i.e. maximal) value z^* of z satisfies all these inequalities and for at least one of them we get actually an equality. Thus $z^* = \min_j \sum_{i=1}^{m} X_i U_{ij}$.

In a similar way we can look for an optimal strategy for the column player. It turns out that this problem reduces to the linear programming problem dual to the one considered above.

Thus let X^* and Y^* be the optimal solutions for both primal and dual problems. The equivalence of primal and dual yields

$$\min_Y X^*UY = \max_X XUY^*$$

which ends the proof of the minmax theorem.