Games in strategic form

Wiesław Zielonka

www.irif.fr/~zielonka

mail: zielonka@irif.fr

November 20, 2018

1 Games in Strategic Form

Let O be a non-empty set of outcomes. A *preference relation* over O is a binary relation $\preccurlyeq$ over O which reflexive ($a \preccurlyeq a$ for all $a \in O$), transitive ($a \preccurlyeq b$ and $b \preccurlyeq c$ imply $a \preccurlyeq c$) and total (for all $a, b \in O$, either $a \preccurlyeq b$ or $b \preccurlyeq a$). We write $a \prec b$ if $a \preccurlyeq b$ but not $b \preccurlyeq a$. Intuitively, $a \prec b$ means that outcome b is preferable to a.
Suppose that N is a finite non-empty set of players. Each player i has a non-empty set A_i of strategies or actions. The elements of $A = \prod_{i \in N} A_i$ are called action profiles (or strategy profiles).
Let us suppose that there is a mapping $g : A \longrightarrow O$ from the set of action profiles to a set O of outcomes each player i has his preference relation $\preccurlyeq_i$ over the set of outcomes. The game is played in the following way: the players choose simultaneously and independently actions a_i, $i \in N$, which yields an outcome $g(a_1, \ldots, a_n)$.
We assume that the whole structure of the game, i.e. the set of outcomes, the players' preference relations are known to all players, the only uncertainty concerns the actions chosen by the players.
In fact we can omit the outcomes from this description and assume that $\preccurlyeq_i$ are preference relations over the set A of action profiles.
The triple $(N, (A_i)_{i \in N}, (\preccurlyeq_i)_{i \in N})$ is a strategic-form game.
Suppose that players can quantitatively measure their preferences, i.e. each player has a payoff or utility mapping $u_i : O \longrightarrow \mathbb{R}$ such that, for all outcomes $a, b \in O$, $a \preccurlyeq_i b$ iff $u_i(a) \leq u_i(b)$. In this case we get rid of outcomes as well and assume that u_i map directly the set A of action profiles into $\mathbb{R}$.
If we use payoff mappings rather than preference relations then strategic-form game is a triple $(N, (A_i)_{i \in N}, (u_i)_{i \in N})$, where A_i are sets of actions/strategies and $u_i : A \longrightarrow \mathbb{R}$ payoff mappings.

Notation. Let $(X_i)_{i \in N}$ be any family of sets indexed by elements of N (by players). Then, for each player $i \in N$, $X_{-i} = \prod_{j \in N \setminus \{i\}} X_j$ denotes the product of all X_j except X_i. By the same token, $x_{-i} = (x_1, \ldots, x_{i-1}, x_{i+1}, \ldots, x_n)$ will denote an element of X_{-i} and, for any $x'_i \in X_i$, $(x_{-i}, x'_i) = (x_1, \ldots, x_{i-1}, x'_i, x_{i+1}, \ldots, x_n)$ is an element of $X = \prod_{j \in N} X_j$. This notation is standard in game theory.

1.1 Nash equilibria

Let $\Gamma = (N, (A_i), (\preccurlyeq_i))$ be a strategic game.

For $i \in N$ and $a_{-i} \in A_{-i}$ we define the set $B_i(a_{-i})$ of best responses of player i given a_{-i} to be

$$B_i(a_{-i}) = \{a_i \in A_i \mid (a_{-i}, a'_i) \preccurlyeq_i (a_{-i}, a_i) \text{ for all } a'_i \in A_i\}.$$

A *Nash equilibrium* is an action profile $a^* \in A$ such that, for each player $i \in N$, $a^*_i \in B_i(a^*_{-i})$, i.e. no player can do better by changing unilaterally his action at a^*.

1.2 Examples

Battle of the sexes. John and Mary wish to go out together. Their priority is to be together but John prefers to go to a concert (he likes classic music) and Mary prefers to go to a football match.
This situation is modeled by the following game where John chooses the line and Mary chooses the column. The entries inside show John's and Mary's payoffs.
Intuitively they both want to coordinate their actions but have conflicting interests.

	concert	match
concert	$2, 1$	$0, 0$
match	$0, 0$	$1, 2$

There are two equilibria (concert, concert) and (match,match).

Coordination game. Now John and Mary both prefer the concert:

	concert	match
concert	$2, 2$	$0, 0$
match	$0, 0$	$1, 1$

Note that there still two equilibria, in particular (match, match) is a Nash equilibrium even if both players would prefer the outcome (concert, concert).

The prisoners' dilemma. Two gang members are interrogated by prosecutors (the interrogation takes place in separate rooms). The thugs are suspected of committing together two crimes, one minor and one serious. The proofs for the minor crime are sufficient to convict both gangsters. However for the serious crime the proofs are weak and without the confession of at least one of the gangsters the prosecutors will not be able to convince the jury. If both gangsters confess their serious crime then each gets 9 years of prison. If one gangster confesses and cooperates with the prosecutors while the other refuses to cooperate then the cooperating gangster will get only 1 year of prison while his accomplice will spend 11 years in jail. If both gangsters do not confess then they will spend 2 years in jail for the minor crime.

	don't confess	confess
don't confess	$-2, -2$	$-11, -1$
confess	$-1, -11$	$-9, -9$

Note that the only equilibrium is (confess, confess) even if both players would certainly prefer the outcome resulting from (don't confess, don't confess).

Matching pennies. Two players show coins. If both show the same face then one of them wins 1 euro, if both show different faces then the other wins 1 euros. This is a strictly antagonistic game, what one player wins the other loses.

	head	tail
head	$1,-1$	$-1,1$
tail	$-1,1$	$1,-1$

No Nash equilibrium.

1.3 Existence of Nash equilibria

A game in strategic form $\Gamma = (N, (A_i), (u_i))$ is said to be finite if actions sets A_i are finite.
Let $\Delta(A_i)$ be the set of probability distributions (measures) over A_i, i.e. the elements of $\Delta(A_i)$ are mappings $\sigma_i : A_i \longrightarrow [0,1]$ such that $\sum_{a_i \in A_i} \sigma_i(a_i) = 1$.
A *mixed* strategy of player i is an element of $\Delta(A_i)$. In this setting a *pure* strategy is a strategy that assigns probability 1 to one of the actions (and 0 to all the other actions).
We extend payoff mapping from pure strategy profiles to mixed strategy profiles by taking the expectation. Thus if $\sigma = (\sigma_1, \dots, \sigma_n) \in \prod_{i \in N} \Delta(A_i)$ is a mixed strategy profile then the payoff of player i is

$$u_i(\sigma) = \sum_{a_1 \in A_1} \dots \sum_{a_n \in A_n} \sigma_1(a_1) \cdots \sigma_n(a_n) u_i(a_1, \dots, a_n).$$

The notion of Nash equilibria transfers to mixed strategies in the obvious way. A mixed strategy profile $\sigma^* = (\sigma_1^*, \dots, \sigma_n^*)$ is a Nash equilibrium if for each player $i \in N$ and each mixed strategy $\sigma_i \in \Delta(A_i)$, $u_i(\sigma^*) \geq u_i(\sigma_{-i}^*, \sigma_i)$.
Abusing the notation, for $a_i \in A_i$, we shall note in the sequel by a_i the strategy of $\Delta(A_i)$ that assigns the probability 1 to a_i and probability 0 to all actions of $A_i \setminus \{a_i\}$

Theorem 1 (Nash)**.** *Finite games have Nash equilibria in mixed strategies.*

For bimatrix games[1] we can calculate Nash equilibria using the Lemke-Howson algorithm. The problem of finding Nash equilibria in bimatrix games is PPAD complete (a complexity class), hence probably no polynomial algorithm exists.
For three players the situation is even worse. There exist finite strategic games for three players such that the payoffs are rational for pure strategy profiles, but for each Nash equilibrium σ, for each player i and each action $a_i \in A_i$, $\sigma_i(a_i)$ is either 0 or irrational.

Let $\sigma_i \in \Delta(A_i)$ be a mixed strategy. The *support* of σ_i is the set $\{a_i \in A_i \mid \sigma_i(a_i) > 0\}$ of actions that have positive probability under σ_i.
A trivial but very important observation concerning Nash equilibria is given in the following lemma:

Lemma 2. *Let $\sigma^* = (\sigma_1^*, \dots, \sigma_n^*)$ be a Nash equilibrium of a finite strategic game $\Gamma = (N, (A_i), (u_i))$. Then*

[1] Two-player finite strategic games.

(i) for each player $i \in N$ and each action $a_i \in A_i$, $u_i(\sigma^*_{-i}, a_i) \le u_i(\sigma^*_{-i}, \sigma^*_i) = u_i(\sigma^*)$.

(ii) moreover, for each $a_i \in A_i$ belonging to the support of σ_i, $u_i(\sigma^*_{-i}, a_i) = u_i(\sigma^*_{-i}, \sigma_i) = u_i(\sigma^*)$.

Proof. (i) follows directly from the definition of a Nash equilibrium.
Indeed if $u_i(\sigma^*_{-i}, a_i) > u_i(\sigma^*_{-i}, \sigma^*_i)$ then player i would win strictly more by deviating and playing a_i with probability 1 instead of σ^*_i.
(ii) Let $\mathrm{supp}(\sigma^*_i)$ be the support of σ^*_i. Let $b_i \in \mathrm{supp}(\sigma^*_i)$ be such that for all $a_i \in \mathrm{supp}(\sigma^*_i), u_i(\sigma^*_{-i}, a_i) \le u_i(\sigma^*_{-i}, b_i)$, i.e. b_i is an action of $\mathrm{supp}(\sigma^*_i)$ maximizing the utility of player i against the profile σ^*_{-i} of his adversaries. Then

$$u_i(\sigma^*) = u_i(\sigma^*_{-i}, \sum_{a_i \in A_i} \sigma^*_i(a_i) a_i) = \sum_{a_i \in A_i} \sigma^*_i(a_i) \cdot u_i(\sigma^*_{-i}, a_i) =$$
$$\sum_{a_i \in \mathrm{supp}(\sigma^*_i)} \sigma^*_i(a_i) \cdot u_i(\sigma^*_{-i}, a_i) \le u_i(\sigma^*_{-i}, b_i) \cdot (\sum_{a_i \in \mathrm{supp}(\sigma^*_i)} \sigma^*_i(a_i)) = u_i(\sigma^*_{-i}, b_i).$$

If the inequality above were strict then player i would profit from changing his strategy from σ^*_i to the pure strategy b_i and σ^* would not be in equilibrium. Thus in equilibrium the last inequality is an equality which proves (ii). □

1.4 Computing Nash equilibria of bimatrix games – an example

Lemma 2 can be used to calculate equilibria of two-person games.
Consider the game

	L	M	R
T	$7, 2$	$2, 7$	$3, 6$
B	$2, 7$	$7, 2$	$4, 5$

Easy case analysis allows to exclude the existence of equilibria where both players use pure (non randomized) strategies.

If player 1 uses pure strategy T then the best response of player 2 is strategy M. Similarly if player 1 plays B then the best response of player 2 is L. Since we have already excluded pure equilibria we can see that in any equilibrium player 1 should choose both T and B with positive probabilities.
Similarly if player 2 chooses one of the strategies L,M,B with probability 1 then the best response of player 1 is a non randomized strategy.
Thus in any equilibrium (σ, τ), $\mathrm{supp}(\sigma) = \{T, B\}$ and $\mathrm{supp}(\tau)$ contains at least two elements.

Let us check whether there exists an equilibrium (σ, τ) such that $\mathrm{supp}(\sigma) = \{T, B\}$ and $\mathrm{supp}(\tau) = \{L, M, R\}$.
By Lemma 2 we should have $u_1(T, \tau) = u_1(B, \tau)$ and $u_2(\sigma, L) = u_2(\sigma, M) = u_2(\sigma, R)$, which yields

$$7\tau(L) + 2\tau(M) + 3\tau(R) = 2\tau(L) + 7\tau(M) + 4\tau(R)$$

and

$$2\sigma(T) + 7\sigma(B) = 7\sigma(T) + 2\sigma(B) = 6\sigma(T) + 5\sigma(B)$$

with two probability equations $\sigma(T) + \sigma(B) = 1$ and $\tau(L) + \tau(R) + \tau(M) = 1$.
$2\sigma(T) + 7\sigma(B) = 7\sigma(T) + 2\sigma(B)$ yields $\sigma(T) = \sigma(B)$, whence $\sigma(T) = \sigma(B) = \frac{1}{2}$. This contradict however $7\sigma(T) + 2\sigma(B) = 6\sigma(T) + 5\sigma(B)$.

Our next guess is $\text{supp}(\sigma) = \{T, B\}$ and $\text{supp}(\tau) = \{M, R\}$. This yields the equation $u_1(T, \tau) = u_1(B, \tau)$, i.e.

$$2\tau(M) + 3\tau(R) = 7\tau(M) + 4\tau(R)$$

with $\tau(M) + \tau(R) = 1$ and $u_2(\sigma, M) = u_2(\sigma, R)$, i.e.

$$7\sigma(T) + 2\sigma(B) = 6\sigma(T) + 5\sigma(B)$$

with $\sigma(T) + \sigma(B) = 1$.
The equations with τ have a unique solution with $\tau(M) = -\frac{1}{4}$. Thus our guess was again incorrect, no equilibrium with the supports prescribed above.

Let us try $\text{supp}(\sigma) = \{T, B\}$ and $\text{supp}(\tau) = \{L, M\}$. We get $u_1(T, \tau) = u_1(B, \tau)$, i.e.

$$7\tau(L) + 2\tau(M) = 2\tau(L) + 7\tau(M)$$

with $\tau(L) + \tau(M) = 1$ and $u_2(\sigma, L) = u_2(\sigma, M)$, i.e.

$$7\sigma(T) + 2\sigma(B) = 7\sigma(T) + 2\sigma(B)$$

with $\sigma(T) + \sigma(B) = 1$. The unique solution is $\sigma(T) = \sigma(B) = \frac{1}{2}$ and $\tau(L) = \sigma(M) = \frac{1}{2}$. Is this an equilibrium? Note that $u_2(\sigma, \tau) = \frac{9}{2}$, while $u_2(\sigma, R) = \frac{11}{2}$, i.e. if player 1 plays σ then player 2 would be better playing R rather than τ, (σ, τ) is not an equilibrium.

The last guess, $\text{supp}(\sigma) = \{T, B\}$ and $\text{supp}(\tau) = \{L, R\}$ is left as an exercise (since we know that there must exist a Nash equilibrium we should obtain one in this case).

Exercise 1. In the middle of the night n persons are waked up by a noise coming from the street. Looking through the window each of them can see a burglar entering a shop through a broken window. (And they see also the lights in other windows thus each of them knows who is waked up – this is the common knowledge shared by all of them.)
All n persons are good citizens so they prefer to see the burglar arrested than to be left free. It is sufficient that at least one person calls the police to be sure that the burglar is arrested. If nobody calls the police then the burglar will get free for sure.
Calling the police is however a nuisance (the caller should testify in the court which will take his/her precious time and already dealing with the police is maybe not a pleasure even if we are just witnesses).
Each person should choose one of the two actions: either C - call the police or I - ignore the incident and go to sleep pretending not to see anything.
We assume that for each person P the payoff (measuring the satisfaction) is as follows:

- 0 – if nobody calls the police and the burglar goes free,

- 50 – if the burglar is arrested because somebody calls the police but the action chosen by P is I (P is satisfied because the burglar is arrested and he/she does not suffer from the inconveniences of dealing with the police),
- 45 if P called the police (the burglar is arrested but P's satisfaction is diminished by the trouble that P has because of his/her call). This payoff is the same even if there were other persons that called the police.

All n person should choose between actions C and I simultaneously (this is a game in strategic form).
Find Nash symmetric[2] equilibria.
Hint. Use Lemma 2.

Exercise 2 (Ice cream vendors)**.** On a sunny day n ice cream vendors choose their positions on a beach. Each vendor attracts all clients that are closer to him than to any of the fellow vendors. If k vendors choose the same point of the beach then they would attract the same clients as one vendor situated at this point, however now they should share these clients equitably. For which values of n there exist Nash equilibria? Describe Nash equilibria in this game. (We are looking here for pure strategy equilibria, thus a strategy of each vendor is to choose deterministically the point x where he puts his stand. Note that there is an infinity of strategies now.)
To present the problem mathematically we suppose that the beach is represented by an interval $[0, a]$ of length $a > 0$. The (potential) clients are distributed uniformly on the beach, i.e. the number of clients on the interval $[c, d] \subset [0, a]$ is proportional to $(d - c)/a$. Suppose that vendors are positioned at points $x_j, j \in N$, where N the set of vendors. Then for each j, $0 \leq x_j \leq a$. For the i-th vendor let $l_i = \max\{0\} \cup \{x_j \mid j \in N \text{ and } x_j < x_i\}$ be the position of the closest vendor situated strictly on the left to i (or 0 if the is no vendor left to i). Let $r_i = \min\{a\} \cup \{x_j \mid j \in N \text{ and } x_j > x_i\}$ be the position of the closest vendor on the right side of i (or a if there is no vendor on the right to i). Finally let $k_i = |\{j \in N \mid x_j = x_i\}|$ be the number of vendors sharing exactly the same position as vendor i. Then vendor's i market share is $\frac{r_i - l_i}{2k_i}$.

Exercise 3. Army A has a single plane with which it can attack one of three possible targets. Army B has one anti-aircraft gun that can be assigned to defend one of the three targets. The value of target i is v_i with $v_1 > v_2 > v_3 > 0$ for each army. The attack of army A is successful only if the attacked target is not defended. Army A wishes to maximize expected damages while army B wishes to minimize expected damages. What are mixed strategy Nash equilibria?

1.5 Final remarks

A game can have many Nash equilibria, some of them seem intuitively to be more plausible than others. For this reason different refinements of Nash equilibria were proposed. Such refinements tend to be more selective and to choose only some of Nash equilibria. All these refinement concepts have their proper drawbacks and we do not consider them here.

[2] Symmetric means that each person uses the same randomized strategy $p \times C + (1 - p) \times I$.

One can wonder also if rational players always will play a Nash equilibrium. Consider the game Γ

	L	R
U	$2,2$	$4,1$
D	$4,1$	$3,3$

This game has only one Nash equilibrium when player 1 plays $\sigma^* = 2/3U + 1/3D$ while player 2 plays $\tau^* = 1/3L + 2/3R$. This equilibrium yields the payoffs 10/3 to player 1 and 5/3 to player 2, $u_1(\sigma^*, \tau^*) = 10/3$ and $u_2(\sigma^*, \tau^*) = 5/3$.
Note however that 10/3 is only the expected payoff of player 1 if player 2 plays τ^*. However if player 2 plays τ^* then he can play L with a positive probability and then player 1 receives only $u_1(\sigma^*, L) = 8/3 < 10/3$. Of course, if player 2 plays τ^* then he can play also R and then player 1 would receive $u_1(\sigma^*, R) = 11/8 > 10/8$. Thus if player 1 plays σ^* while player 2 plays τ^* then there is some risk for player 1 to get less than 10/3 but also there is some chance to gain more than 10/3. However, this is only the case when player 2 plays the randomized strategy τ^*.
Consider however the worst case scenario, player 2 instead of playing τ^* plays another strategy, in fact consider the case when player 2 plays the strategy which is the worst possible from the point of view of player 1.
How much player 1 can assure himself against any possible strategy of player 2? This can be seen by solving the zero-sum game Γ'

	L	R
U	$2,-2$	$4,-4$
D	$4,-4$	$3,-3$

where player 2 is replaced by a player which is the the most hostile against 1 (i.e. with payoffs which are the exact inverse of payoffs of player 1).
The equilibrium of this games is achieved when player 1 plays the strategy $\sigma = 1/3U + 2/3D$ (the strategy of player 2 in Γ' has no interest for our problem but you can compute it if you are interested). With this strategy player 1 wins in the original game Γ (as well as in the modified game Γ') $u_1(\sigma, L) = 10/3$ and $u_1(\sigma, R) = 10/3$, i.e. he always guarantees the payoff 10/3 against all possible (deterministic or mixed) strategies of player 2.
If you are player 1 what would you prefer to play in game Γ, σ^* or σ? Strategy σ^* assures for player 1 that he wins 10/3 but only if player 2 plays τ^*, strategy σ assures for him to win 10/3 against all strategies of player 2.
A similar reasoning applies to player 2.

Exercise 4. Find the best strategy τ for player 2 against any, even the most hostile strategy of player 1. Which payment can player 2 guarantee for himself with this strategy τ?

2 Zero-sum games

A zero-sum game is a two player game such that for each pair $a_1 \in A_1$ and $a_2 \in A_2$, $u_1(a_1, a_2) = -u_2(a_1, a_2)$.

Thus the players have here exactly opposite objectives, what one player wins the other one loses.
In the sequel we denote by u the payoff mapping of player 1. Thus the objective of player 1 is to maximize u while player 2 wants to minimize u.
Suppose for a moment that the sets A_1, A_2 of strategies are finite.
Suppose that player 1 chooses a strategy $a_1 \in A_1$. Then he is sure to win at least $\min_{a_2 \in A_2} u(a_1, a_2)$ (consider the move of player 2 that is the worst possible from the point of view of player 1).
If player 1 is cautious then he will choose a strategy maximizing this expression and choosing such a strategy he is sure to win at least

$$\underline{v} = \max_{a_1 \in A_1} \min_{a_2 \in A_2} u(a_1, a_2).$$

We call this quantity *the lower value of the game.*
By the symmetric reasoning, by choosing an appropriate strategy, player 2 can limit his losses to

$$\overline{v} = \min_{a_2 \in A_2} \max_{a_1 \in A_1} u(a_1, a_2).$$

This is the upper value of the game.
Note that always

$$\underline{v} = \max_{a_1 \in A_1} \min_{a_2 \in A_2} u(a_1, a_2) \leq \min_{a_2 \in A_2} \max_{a_1 \in A_1} u(a_1, a_2) = \overline{v}.$$

If the sets A_1, A_2 of strategies are infinite then we should replace min and max by inf and sup but we still have the inequality

$$\underline{v} = \sup_{a_1 \in A_1} \inf_{a_2 \in A_2} u(a_1, a_2) \leq \inf_{a_2 \in A_2} \sup_{a_1 \in A_1} u(a_1, a_2) = \overline{v}.$$

If $\underline{v} = \overline{v}$ then we say that the game has the value

$$v = \sup_{a_1 \in A_1} \inf_{a_2 \in A_2} u(a_1, a_2) = \inf_{a_2 \in A_2} \sup_{a_1 \in A_1} u(a_1, a_2).$$

If the game has a value v then a strategy $a_1 \in A_1$ of player 1 is said to be *optimal* if for each strategy $a_2 \in A_2$ of player 2 we have $u(a_1, a_2) \geq v$. Thus player 1 obtains at least v against any strategy of player 2.
Similarly a strategy $a_2 \in A_2$ of player 2 is said to be optimal if for each strategy $a_1 \in A_1$ of player 1 we have $u(a_1, a_2) \leq v$. Thus player 2 loses at most v against any strategy of player 1.
Note that strategies $a_1 \in A_1, a_2 \in A_2$ are optimal for a zero sum game if and only if (a_1, a_2) is a Nash equilibrium.
Let us return to the case where A_1 and A_2 are finite. Matching pennies game shows that for such games the lower and the upper value can be different when players use only pure strategies.
What happens if players use mixed strategies, i.e. the set of strategies of player i is $\Delta(A_i)$? Our aim is to find the lower value $\underline{v}$ of such a game.

First of all note that, by the same argument as in the proof of Lemma 2, for each mixed strategy $\sigma_1 \in \Delta(A_1)$,

$$\underline{v} = \min_{\sigma_2 \in \Delta(A_2)} \sum_{a_1 \in A_1} \sum_{a_2 \in A_2} \sigma_2(a_2)\sigma_1(a_1)u(a_1, a_2) = \min_{a_2 \in A_2} \sum_{a_1 \in A_1} \sigma_1(a_1)u(a_1, a_2)$$

implying

$$\underline{v} = \max_{\sigma_1 \in \Delta(A_1)} \min_{\sigma_2 \in \Delta(A_2)} \sum_{a_1 \in A_1} \sum_{a_2 \in A_2} \sigma_2(a_2)\sigma_1(a_1)u(a_1, a_2) = \max_{\sigma_1 \in \Delta(A_1)} \min_{a_2 \in A_2} \sum_{a_1 \in A_1} \sigma_1(a_1)u(a_1, a_2). \tag{1}$$

Without loss of generality we can assume that $A_1 = \{1, \ldots, k\}$, $A_2 = \{1, \ldots, m\}$. e set $u(i, j) = u_{ij}$, $1 \le i \le k, 1 \le j \le m$. Then $U = (u_{ij})$ is the payoff matrix.
A mixed strategy of player 1 is represented by a vector $x = (x_1, \ldots, x_k)$ such that $x_i \ge 0$ and $\sum_{i=1}^{k} x_i = 1$ and a mixed strategy of player 2 is a vector $y = (y_1, \ldots, y_m)$ such that $\sum_{i=1}^{m} y_i = 1$ (x_i, y_j give the probabilities of choosing pure strategies i and j). In this notation the payoff received by player 1 is players use strategies x and y is xUy^T, where y^T is the column vector — the transpose of y. Then (1) takes the form

$$\underline{v} = \max_{x \in \Delta_k} \min_{j \in \{1,\ldots,m\}} \sum_{i=1}^{k} x_i u_{ij}$$

where $\Delta_k = \{(x_1, \ldots, x_k) \mid \sum_i x_i = 1, \text{all } x_i \ge 0\}$.
In other words the aim of player 1 is to choose a strategy $x = (x_1, \ldots, x_k) \in \Delta_k$ such that

$$\min_{j \in \{1,\ldots,m\}} \sum_{i=1}^{k} x_i u_{ij}$$

has the greatest possible value.
This amounts to solving the following linear programming problem:

$$\begin{aligned} \text{maximize } & \underline{v} \\ \text{subject to } & \\ \sum_{i=1}^{k} x_i u_{ij} \ge \underline{v}, \quad & \forall j, 1 \le j \le m \\ & \sum_{i=1}^{k} x_i = 1 \\ x_i \ge 0, \quad & \forall 1 \le i \le k \end{aligned}$$

where $\underline{v}$ and x_i are unknown and u_{ij} are fixed constants given by the payoff mapping.

Calculating the upper value $\overline{v}$ leads to the dual linear programming problem:

$$\begin{aligned} \text{minimize} \quad & \overline{v} \\ \text{subject to} \\ \sum_{j=1}^{m} u_{ij} y_j \leq \overline{v}, \quad & \forall i, 1 \leq i \leq k \\ \sum_{j=1}^{m} y_j = 1 \\ y_j \geq 0, \quad & \forall 1 \leq i \leq m \end{aligned}$$

where $\overline{v}$ and y_j are unknown.
Since the direct and the dual problems have the same optimal solution we deduce that $\underline{v} = \overline{v}$ i.e. the value exists for finite games if players can use mixed strategies. This result was first proved by John von Neumann. Note that solving the corresponding linear programming problems we obtain also the optimal strategies for both players.