NASH EQUILIBRIA

CLARA STROHM

ABSTRACT. In this talk, the paper *Non-cooperative games* of John Nash is discussed. Nash's article provides a generalisation of the game theory developed by von Neumann and Morgenstern. The central notion are equilibrium points of non-cooperative games. This can be used to describe rational behaviour in various situations and to develop profitable strategies.

Game theory – an introduction

In this section, we introduce the notion of finite games, as defined by von Neumann and Morgenstern [2].

Games and Strategies. Intuitively, a game consists of a set of players, rules according to which the players may act and probably some chance element (for example dices). Moreover, every player assigns to each possible end position of the game a certain value ("payoff") and tries to optimise this value. There are many types of games depending on how much information the players get during play etc.

However, von Neumann and Morgenstern achieved a profound simplification and were able to deduce the following abstract description of games [2; Section 11]:

Definition (1.1). Let $n \in \mathbf{N}$. A **finite n-person game** is a pair (S, p) consisting of a product $S = S_1 \times \cdots \times S_n$ of non-empty finite sets and a function $p\colon S \longrightarrow \mathbf{R}^n$. For $j \in \{1, \dots, n\}$ the elements of the set S_j are called **pure strategies** of the player j. The function p is called **payoff function**. ♣

Here the player might be individuals but also groups, nations or even species. Each player tries to find a strategy which maximises his component of the payoff function. How do we have to interpret the "strategies"? A strategy is not a single move in a game but a collection of moves for every possible position in the game. The game is non-cooperative in the sense that all players must choose their strategies without communication with the other players.

To obtain a satisfying theory, one has to introduce the possibility that the players make random moves, i.e., that they choose their strategy according to some probability distribution:

Date: May 4, 2005.
Key words and phrases. game theory, Nash equilibrium.

Definition (1.2). Let (S,p) be a finite n-person game and let $j \in \{1,\dots,n\}$. The set of **mixed strategies** of player j is given by

$$\overline{S_j} := \Big\{ \sum_{\pi \in S_j} a_\pi \cdot \pi \;\Big|\; (a_\pi)_{\pi \in S_j} \subset [0,1],\ \sum_{\pi \in S_j} a_\pi = 1 \Big\}.$$

In addition, we use the abbreviation

$$\overline{S} := \overline{S_1} \times \cdots \times \overline{S_n}.$$

A mixed strategy $\sum_{\pi \in S_j} a_\pi \cdot \pi \in \overline{S_j}$ **uses** the pure strategy $\pi \in S_j$ if $a_\pi \neq 0$. Of course, p possesses a unique n-linear extension to $\overline{S}$. We will denote this extension also by p. ♠

Remark (1.3). In the situation of the previous definition, $\overline{S_j}$ is the convex hull of $\S_j$ (if we identify S_j with the canonical base of $\mathbf{R}^{|S_j|+1}$), i.e., $\overline{S_j}$ is a simplex, and $\overline{S}$ is a product of simplices (and hence a convex polytope). The extension of p can be interpreted as the expected payoff. ♡

Example (1.4). In the game "matching pennies," given by the table

p_1			p_2
1	a	α	-1
-1	a	β	1
-1	b	α	1
1	b	β	-1,

player 1 can guarantee that he wins (in expectation) at least 0, when playing the mixed strategy $\frac{1}{2}a + \frac{1}{2}b$. No pure strategy can achieve this. ♢

The central problems in game theory. Unfortunately, games are mostly not played for fun but for profit. Thus each player's objective is to maximise his component of the payoff function. Game theory is concerned both with the description of "real" situations (or rational behaviour) and with the search for "optimal" strategies.

Zero-sum two-person games. Before diving into Nash's general theory for non-cooperative games, we briefly recall the results of von Neumann and Morgenstern on zero-sum two-person games:

Definition (1.5). A **zero-sum two-person game** is a finite two-person game (S,p) with $p_1 + p_2 = 0$. ♣

Example (1.6). The game "matching pennies" (Example (1.4)) is a zero-sum two-person game. More prominent examples include "chess" and "go." ♠

The following theorem [2; Section 17.6] is the central part of this theory.

Theorem (1.7) (Minimax Theorem)**.** *Let (S,p) be a zero-sum two-person game. Then there exists a so-called* ***pair of good (mixed) strategies*** *$(r_1,r_2) \in \overline{S}$, i.e.,*

$$\max_{r \in \overline{S_1}} p_1(r,r_2) = p_1(r_1,r_2) = \min_{r \in \overline{S_2}} p_1(r_1,r).$$

Moreover, for any pair $(r_1,r_2) \in \overline{S}$ of good strategies

$$\max_{r \in \overline{S_1}} \min_{s \in \overline{S_2}} p_1(r,s) = p_1(r_1,r_2) = \min_{s \in \overline{S_2}} \max_{r \in \overline{S_1}} p_1(r,s).$$

This number is the ***value*** *of* (S, p).

A proof will be given in Section 2.1.

Remark (1.8). We can translate this theorem as follows: Let (S, p) be a zero-sum two-person game with value v and let (r_1, r_2) be a pair of good strategies for (S, p). Then player 1 ensures by using the strategy r_1 that he gains (in expectation) at least v and that he prevents player 2 from gaining more than $-v$. A zero-sum two-person game with value 0 is considered to be **fair**. ♡

Example (1.9). For example, the single pair of good strategies in "matching pennies" is $(\frac{1}{2}a + \frac{1}{2}b, \frac{1}{2}\alpha + \frac{1}{2}\beta)$ (which is mixed). The value of this game is 0 (and hence this game is fair).

The game

p_1			p_2
0	a	α	0
-1	a	β	1
0	b	α	0
1	b	β	-1

possesses exactly one good pair of strategies, namely (b, α) (which is pure). The value of this game is 0 (and hence this game is also fair). ♢

The theory of von Neumann and Morgenstern deals effectively with the case of zero-sum two-person games. However, their discussion of games with more than two players or non-zero-sum games is not very convincing.

Nash's paper [1] presents a generalisation of von Neumann's and Morgenstern's theory to the case of non-cooperative games (not necessarily zero-sum and possibly with a large number of players). Of course, these results cannot be as complete as in the zero-sum two-person case. However, they are powerful enough to describe rational behaviour and to indicate feasible strategies.

Nash's solution – equilibrium points

The central notion of Nash's theory [1] are equilibrium points. They are positions in the game which are "locally optimal" for each player.

Definition (2.1). Let (S, p) be a finite n-person game.

- For $s \in \overline{S}$, $j \in \{1, \ldots, n\}$ and $r \in \overline{S_j}$ we introduce the substitution notation
$$s[j]r := (s_1, \ldots, s_{j-1}, r, s_{j+1}, \ldots, s_n).$$
I.e., the strategy of the j-th player in s is replaced by r.
- For $j \in \{1, \ldots, n\}$, a pure strategy $\pi \in S_j$ and $s \in \overline{S}$, we write
$$p_{j,\pi}(s) := p\big(s[j]\pi\big).$$
♣

Definition (2.2). Let (S, p) be a finite game. An n-tuple $s \in \overline{S}$ of mixed strategies is an **equilibrium point** if

$$\forall_{j\in\{1,\ldots,n\}} \quad p_j(s) = \max_{r\in \overline{S_j}} p_j\big(s[j]r\big). \qquad \spadesuit$$

An equilibrium point is hence an n-tuple of mixed strategies such that all players play optimal (i.e., maximise their payoff) when the strategies of the other players are kept fixed. Thus equilibrium points are "stable" in the sense that no single player can gain by choosing another strategy.

Of course, any pair of good strategies in a zero-sum two-person game constitutes an equilibrium point – hence equilibrium points form a generalisation of good strategies in the sense of von Neumann and Morgenstern.

In general, equilibrium points need not be unique (this problem is addressed in Section 3). Furthermore, equilibrium points need not represent a desirable outcome of the game.

Example (2.3).
- The game (b means "defect," a means "remain silent")

p_1			p_2
-2	a	α	-2
-5	b	β	-5
-30	a	β	-1
-1	b	α	-30

  is called "prisoner's dilemma." Its single equilibrium point is (b, β). However, it would be "better" to play (a, α).
- In the game

p_1			p_2
1	a	α	1
0	a	β	1
1	b	α	0
0	b	β	0

  all pairs of mixed strategies are equlibrium points.
- The equilibrium points of the game

p_1			p_2
1	a	α	1
0	a	β	0
0	b	α	0
0	b	β	0

  are (a, α) and (b, β). Hence the players cannot reach an equilibrium point without communication. ♡

We have the following – obvious – characterisation of equilibrium points:

Lemma (2.4). *Let (S, p) be a finite n-person game. An n-tuple $s \in \overline{S}$ of mixed strategies is an equilibrium point of (S, p) if and only if*

$$\forall_{j\in\{1,\ldots,n\}} \quad p_j(s) = \max_{\pi\in S_j} p_{j,\pi}(s).$$

Proof. It is clear that the above condition is necessary. But the condition is also sufficient: Let $j \in \{1, \dots, n\}$ and let $r \in \overline{S_j}$ be some mixed strategy of player j. By definition, there are $(a_\pi)_{\pi \in S_j} \subset [0,1]$ satisfying

$$\sum_{\pi \in S_j} a_\pi = 1 \quad \text{and} \quad r = \sum_{\pi \in S_j} a_\pi \cdot \pi.$$

Therefore, we obtain

$$\begin{aligned} p_j\big(s[j]r\big) &= \sum_{\pi \in S_j} a_\pi \cdot p_j\big(s[j]\pi\big) \\ &\leq \sum_{\pi \in S_j} a_\pi \cdot p_j(s) \\ &= p_j(s), \end{aligned}$$

as desired. $\diamondsuit$

Remark (2.5). The set of equilibrium points of a finite game (S, p) is a closed subset of $\overline{S}$ (by the above Lemma). ♣

Existence of equilibrium points. Using Brouwer's fixed point theorem, Nash showed that every finite game possesses an equilibrium point [1; Theorem 1]:

Theorem (2.6). *Every finite game has an equilibrium point.*

Proof. Let (S, p) be a finite n-person game. In order to apply the Brouwer fixed point theorem, we consider for each $j \in \{1, \dots, n\}$ the transformation

$$\begin{aligned} \varphi \colon \overline{S} &\longrightarrow \overline{S} \\ s &\longmapsto \left(\frac{s_j + \sum_{\pi \in S_j} \varphi_{j,\pi}(s) \cdot \pi}{1 + \sum_{\pi \in S_j} \varphi_{j,\pi}(s)} \right)_{j \in \{1,\dots,n\}} \end{aligned}$$

where

$$\begin{aligned} \varphi_{j,\pi} \colon \overline{S} &\longrightarrow \mathbf{R} \\ s &\longmapsto \max\big(0, p_{j,\pi}(s) - p_j(s)\big) \end{aligned}$$

for all pure strategies $\pi \in S_j$. Thus, φ pushes a mixed strategy towards a more equilibrial strategy.

The fixed points of φ are equilibrium points of the game (S, p): Let $s \in \overline{S}$ be a fixed point of φ, i.e., $\varphi(s) = s$. Let $j \in \{1, \dots, n\}$ and suppose $\pi_0 \in S_j$ is the "least profitable" pure strategy for player j, that is

$$p_{j,\pi_0}(s) \leq p_j(s).$$

In particular, $\varphi_{j,\pi_0}(s) = 0$. Since s is a fixed point, we obtain

$$\begin{aligned} (s_j)_{\pi_0} &= \big(\varphi_j(s)\big)_{\pi_0} \\ &= \frac{(s_j)_{\pi_0} + 0 \cdot \pi_0}{1 + \sum_{\pi \in S_j} \varphi_{j,\pi}(s)} \end{aligned}$$

and hence $1 + \sum_{\pi \in S_j} \varphi_{j,\pi}(s) = 1$. Therefore

$$0 = \varphi_{j,\pi}(s) = \max\big(0, p_{j,\pi}(s) - p_j(s)\big)$$

for all $\pi \in S_j$. Now the characterisation of Lemma (2.4) shows that s is an equilibrium point.

Since $\overline{S} = \overline{S_1} \times \cdots \times \overline{S_n}$ is homeomorphic to a closed ball, the Brouwer fixed point theorem implies that φ has a fixed point and hence that (S, p) possesses an equilibrium point. ♠

We can use this theorem in particular, to prove the minimax theorem (1.7) (this is of course not the original proof by von Neumann and Morgenstern):

Proof (of Theorem (1.7)). By the existence theorem, the game possesses a pair of good strategies.

On the other hand it is easy to see that for all pairs (r_1, r_2) of good strategies (i.e., saddle points) the equality

$$\max_{r \in \overline{S_1}} \min_{s \in \overline{S_2}} p_1(r, s) = p_1(r_1, r_2) = \min_{s \in \overline{S_2}} \max_{r \in \overline{S_1}} p_1(r, s)$$

holds. □

There exist many generalisation of the above existence theorem to more general games.

However, existence of equilibrium points is only one part of the story. To get useful results, we also need some kind of uniqueness of equilibrium points. This is investigated in Section 3.

Solvable games

In a finite game, each player can compute the set of equilibrium points if he knows the payoff functions of all other players (i.e., all players are honest about their objectives). However, since equilibrium points need not be unique and our games are conisdered to be non-cooperative, the players may not be able to reach an equilibrium point without communication. Nash [1] introduced a special class of games in which the players can always obtain an equilibrium point (provided that all players act rationally): solvable games.

Definition (3.1). Let $\mathcal{S}$ be the set of equilibrium points of the n-person game (S, p).

- The game (S, p) is called **solvable** if for all $j \in \{1, \ldots, n\}$, all $t \in \overline{S}$ and all $r \in \overline{S_j}$ the condition
$$(t[j]r) \in \mathcal{S} \text{ and } s \in \mathcal{S} \Longrightarrow (s[j]r) \in \mathcal{S}$$
is satisfied.
- If (S, p) is solvable, then the set $\mathcal{S}$ of equilibrium points is called **solution** of (S, p).
- If (S, p) is solvable and $j \in \{1, \ldots, n\}$, then
$$\mathcal{S}_j := \{s_j \mid s \in \mathcal{S}\}$$
is the set of **equilibrium strategies** of player j. ♡

As mentioned above, in a solvable game, the players can reach an equilibrium point without communication:

Proposition (3.2). *Let (S, p) be a solvable finite n-person game with solution $\mathcal{S}$. Then*

$$\mathcal{S} = \mathcal{S}_1 \times \cdots \times \mathcal{S}_n.$$

Proof. By definition, $\mathcal{S} \subset \mathcal{S}_1 \times \cdots \times \mathcal{S}_n$. Conversely, let $(r_1, \ldots, r_n) \in \mathcal{S}_1 \times \cdots \times \mathcal{S}_n$. In particular, for each $j \in \{1, \ldots, n\}$ there is an $s_{(j)} \in \mathcal{S}$ such that

$$s_{(j)}[j]r_j \in \mathcal{S}.$$

Therefore,

$$\begin{aligned} s_{(1)}[1]r_1 &\in \mathcal{S} \\ s_{(1)}[1]r_1[2]r_2 &\in \mathcal{S} \\ &\vdots \\ (r_1, \ldots, r_n) = s_{(1)}[1]r_1 \cdots r_{n-1}[n]r_n &\in \mathcal{S}. \end{aligned}$$

♢

Example (3.3). The game

p_1			p_2
1	a	α	1
0	a	β	0
0	b	α	0
0	b	β	0

is unsolvable (its equilibrium points are (a, α) and (b, β)).

The game

p_1			p_2
1	a	α	1
0	a	β	1
1	b	α	0
0	b	β	0

also has more than one equilibrium point (all pairs of mixed strategies are equilibrium points) but is solvable. ♣

Proposition (3.4). *Every zero-sum two-person game is solvable. The sets of equilibrium strategies are in this case just the sets of good strategies.*

Proof. Let (S, p) be a zero-sum two-person game, and let $\mathcal{S}$ be its set of equilibrium points. For the first part, let $(r_1, r_2) \in \mathcal{S}$ and $(s_1, s_2) \in \mathcal{S}$. Hence (r_1, r_2) and (s_1, s_2) are pairs of good strategies. Why is $(r_1, s_2) \in \mathcal{S}$? The minimax theorem (Theorem (1.7)) yields

$$\begin{aligned} p_1(r_1, s_2) &\geq \min_{t \in \overline{S_2}} p_1(r_1, t) \\ &= p_1(r_1, r_2) \\ &= \max_{t_1 \in \overline{S_1}} \min_{t_1 \in \overline{S_2}} p_1(t_1, t_2) \\ &= p_1(s_1, s_2) \\ &= \max_{t \in \overline{S_1}} p_1(t, s_2) \end{aligned}$$

and thus

$$p_1(r_1, s_2) = \max_{t \in \overline{S_1}} p_1(t, r_2).$$

Analogously,

$$p_1(r_1, s_2) = \min_{t \in \overline{S_1}} p_1(r_1, t)$$

which implies $(r_1, s_2) \in \mathcal{S}$. Thus, (S, p) is solvable.

The second part is clear by definition. ♠

Variations of solvability. Since not all finite games are solvable, Nash introduced a weaker notion replacing solutions:

Definition (3.5). Let (S, p) be a finite game. A set $\mathcal{S}$ of equilibrium points is a **sub-solution** of (S, p) if it satisfies the solvability condition as in Definition (3.1) and is maximal with respect to inclusion. ♡

Of course, the solution of a solvable game always is a sub-solution. Since any set containing exactly one equilibrium strategy satisfies the solvability condition in Definition (3.1), every finite game possesses at least one sub-solution. Furthermore, a finite game is solvable if and only if it has exactly one sub-solution.

The same proof as in Proposition (3.2) shows that any sub-solution is the product of its projections. Moreover, the shape of the projections is not very wild [1; Theorem 4]:

Theorem (3.6). *The projections of a sub-solution of a finite game are closed and convex as subsets of the mixed strategy spaces.*

Proof. Let (S, p) be a finite n-person game with sub-solution $\mathcal{S}$ and $j \in \{1, \dots, n\}$. In order to prove this theorem, we will make use of the following observation:

Lemma (3.7). *If $r \in \overline{S_j}$ such that $s[j]r$ is an equilibrium point of (S, p) for all $s \in \mathcal{S}$, then already $s[j]r \in \mathcal{S}$ for all $s \in \mathcal{S}$.*

Proof. It is not hard to see that in this case, the set

$$\mathcal{S} \cup \left\{ s[j]r \;\middle|\; s \in \mathcal{S} \right\}$$

satisfies the solvability condition of Definition (3.1). Thus maximality of $\mathcal{S}$ implies taht $s[j]r \in \mathcal{S}$ for all $s \in \mathcal{S}$. ♢

It suffices to show that convex combinations and limit points of points in $\mathcal{S}$ satisfy the assumption in the above lemma. But this is a direct consequence of linearity. ♣

Sub-solutions are a weaker form of solutions. But there are also stronger concepts than solvability, namely strong solvability [1]:

Definition (3.8). A solvable finite n-person game (S, p) with solution $\mathcal{S}$ is called **strongly solvable** if

$$s \in \mathcal{S} \text{ and } p_j\big(s[j]r\big) = p_j(s) \Longrightarrow \big(s[j]r\big) \in \mathcal{S}$$

for all $j \in \{1, \dots, n\}$ and all $r \in \overline{S_j}$. Then $\mathcal{S}$ is the **strong solution** of (S, p). ♠

Proposition (3.9). *If a finite game is strongly solvable, then it possesses an equilibrium point consisting of pure strategies.*

Proof. Let (S, p) be a strongly solvable n-person game with solution $\mathcal{S}$. Let $s \in \mathcal{S}$ be any equilibrium point and let $j \in \{1, \ldots, n\}$. By Lemma (2.4), there is a pure strategy $\pi \in S_j$ such that $p_j(s[j]\pi) = p_j(s)$. Now strong solvability implies $s[j]\pi \in \mathcal{S}$. Applying this argument successively to all players, shows that (S, p) possesses an equilibrium point consisting of pure strategies. ♡

Hence in strongly solvable games it is relatively easy to find an equilibrium point: one has just to check all tuples of *pure* strategies (and there are only finitely many of those!).

Remark (3.10). However, the converse is not true: The game

p_1			p_2
0	a	α	0
-1	a	β	1
0	b	α	0
1	b	β	-1

is solvable, possesses an equilibrium point consisting of pure strategies but is not strongly solvable. ♢

Values. It is always possbible to assign to a finite game upper and lower values corresponding to each player. In solvable games, values indicate how much profit each player can expect when engaging in the game.

Definition (3.11). Let $\mathcal{S}$ be the set of equilibrium points of the finite n-person game (S, p). For $j \in \{1, \ldots, n\}$ we call

$$v_j^+ := \max_{s \in \mathcal{S}} p_j(s)$$
$$v_j^- := \min_{s \in \mathcal{S}} p_j(s)$$

the **upper value** and **lower value of player** j respectively. If $v_j^+ = v_j^-$, then $v_j := v_j^+ = v_j^-$ is the **value of player** j. ♣

If (S, p) is a solvable finite n-person game, then player $j \in \{1, \ldots, n\}$ wins (in expectation) at least v_j^- and at most v_j^+ if all players act rationally.

If a finite game has only one equilibrium point, then obviously values exist for all players. But in general values do not have to exist (even strongly solvable games need not have values).

Geometry of equilibrium strategies. Analoguous to the zero-sum two-person case, where the good strategies form polyhedral subsets of the respective mixed strategy spaces, Nash proved a corresponding result for solvable games describing the geometry of the sets of equilibrium strategies [1; Theorem 5]:

Theorem (3.12). *The sets of equilibrium strategies of a solvable finite game are polyhedral convex subsets of the respective mixed strategy spaces.*

Proof. Let (S, p) be a solvable finite n-person game with solution $\mathcal{S}$. According to Lemma (2.4), an n-tuple $s \in \overline{S}$ is an equilibrium point (i.e., lies in $\mathcal{S}$) if and only if

$$\forall_{j \in \{1, \ldots, n\}} \quad p_j(s) = \max_{\pi \in S_j} p_{j,\pi}(s).$$

This condition is equivalent to

$$\forall_{j\in\{1,\ldots,n\}} \quad \forall_{\pi\in S_j} \quad p_j(s) - p_{j,\pi}(s) \geq 0.$$

By Theorem (2.6), the game (S, p) possesses an equilibrium point s. Using Proposition (3.2), we obtain that a mixed strategy $r \in \overline{S_j}$ is an equilibrium strategy if and only $(s[j]r)$ is. By the above discussion, this is equivalent to saying that

$$\forall_{k\in\{1,\ldots,n\}} \quad p_k\big(s[j]r\big) - p_{k,\pi}\big(s[j]r\big) \geq 0.$$

Since for each $k \in \{1,\ldots,n\}$ the function

$$\begin{aligned} \overline{S_j} &\longrightarrow \mathbf{R} \\ r &\longmapsto p_k\big(s[j]r\big) - p_{k,\pi}\big(s[j]r\big) \end{aligned}$$

is linear and $\overline{S_j}$ is a simplex, the claim follows. ♠

Finding equilibria

In a finite game, each player can “compute” the set of all equilibrium points provided he knows the payoff function of all other players. For example, in a zero-sum two-person game the equilibrium points can be computed using linear programming. In the general case, we can use symmetries and the analysis of dominances.

Linear programming. Let (S, p) be a zero-sum two-person game. Then the first component $r_1 = \sum_{\pi\in S_1} a_\pi \cdot \pi$ of an equilibrium point (and the corresponding value v) can be computed via linear programming as follows (analogously for the second component):

$$\begin{array}{ll} \text{maximise} & v + \sum_{\pi\in S_1} 0\cdot a_\pi \\ \text{with} & v \in \mathbf{R},\ (a_\pi)_{\pi\in S_1} \subset \mathbf{R} \\ \text{under} & \forall_{\pi\in S_1} \quad a_\pi \geq 0 \\ & \sum_{\pi_1} a_\pi = 1 \\ & \forall_{\pi_2\in S_2} \quad v \leq \sum_{\pi_1\in S_1} a_\pi \cdot p_1(\pi_1, \pi_2). \end{array}$$

I.e.,

$$v = \max_{r\in\overline{S_1}} \min_{\pi_2\in S_2} p_1(r, \pi_2).$$

Symmetries of games. Symmetries form an additional structure on games. Sometimes symmetries can help to understand a game better. In particular, Theorem (4.3) shows that symmetries give hints for equilibrium points.

Definition (4.1). Let (S, p) be a finite n-person game and let $\chi\colon \coprod_{j=1}^n S_j \longrightarrow \coprod_{j=1}^n S_j$ be a bijection such that for each $j \in \{1,\ldots,n\}$ there is a $\psi(j) \in \{1,\ldots,n\}$ satisfying

$$\chi(S_j) \subset S_{\psi(j)}.$$

The bijection χ induces a bijection

$$\chi\colon S \longrightarrow S.$$

Then χ is a **symmetry** of (S, p) if

$$\forall_{j\in\{1,\dots,n\}} \quad \forall_{s\in S} \quad p_{\psi(j)}\big(\chi(s)\big) = p_j(s).$$

A **symmetric n-tuple** of (S, p) is a point $s \in \overline{S}$ such that

$$\chi(s) = s$$

for all symmetries χ. ♡

Example (4.2). The zero-sum two-person game given by the matrix

p_1			p_2
0	a	α	0
-1	a	β	1
0	b	α	0
1	b	β	-1

has no non-trivial symmetries at all.

The game "matching pennies" has exactly one non-trivial symmetry, namely

$$\begin{aligned} a &\longmapsto \beta \\ b &\longmapsto \alpha \\ \alpha &\longmapsto a \\ \beta &\longmapsto b \end{aligned}$$

The only symmetric pair of this game is $(\frac{1}{2}a + \frac{1}{2}b, \frac{1}{2}\alpha + \frac{1}{2}\beta)$. ♢

Analogously to Theorem (2.6), Nash showed taht each finite game has a symmetric equilibrium point [1; Theorem 2]:

Theorem (4.3). *Any finite game has a symmetric equilibrium point.*

Proof. Let (S, p) be a finite n-person game.

Then (S, p) has a symmetric n-tuple: Of course,

$$\Big(\frac{1}{|\overline{S_j}|} \cdot \sum_{\pi\in S_j} \pi\Big)_{j\in\{1,\dots,n\}}$$

is an element of $\overline{S}$ and symmetric.

The set S_{sym} of symmetric equilibrium points is homeomorphic to a closed ball: It suffices to show that s_{sym} is closed and convex. But this easily follows from the various linearity conditions.

It is easy to see that the transformation φ in the proof of the existence theorem maps symmetric points to symmetric points. Hence the Brouwer fixed point theorem implies that (S, p) possesses a symmetric equilibrium point. ♣

Example (4.4). In particular, the pair $(\frac{1}{2}a + \frac{1}{2}b, \frac{1}{2}\alpha + \frac{1}{2}\beta)$ must be an equilibrium point of "matching pennies." ♠

Dominance. Strategies can be partially ordered with respect to a relation called "dominance." However, dominance is a very restrictive concept, i.e., many strategies are not comparable with respect to dominance.

Definition (4.5). Let (S, p) a finite n-person game, let $j \in \{1, \ldots, n\}$ and $r, r' \in \overline{S_j}$. The strategy r **dominates** r' if

$$\forall_{s \in \overline{S}} \quad p_j\big(s[j]r\big) > p_j\big(s[j]r'\big).$$

The strategy r' is **undominated** if there is no strategy r dominating r'. ♡

Obviously, no equilibrium point can use a dominated strategy. This fact can be used to find equilibria, namely all dominated strategies can be excluded from the search. This in particular applies to solvable games (where the set of equilibrium points decomposes as a product). Moreover, Nash [1] claims the following result about the shape of the set of undominated strategies (but he does not give a proof):

Proposition (4.6). *The set of undominated strategies of a finite game is simply connected and formed by the union of some collection of faces of the strategy simplex.*

References

[1] John F. Nash. Non-cooperative games. Ann. of Math. (2), 54, (1951), 286–195.

[2] John von Neumann, Oskar Morgenstern. *Theory of Games and Economic Behaviour*. Second edition. Princeton University Press, (1947).

GRADUIERTENKOLLEG "ANALYTISCHE TOPOLOGIE UND METAGEOMETRIE", WWU MÜNSTER, EINSTEINSTR. 62, 48149 MÜNSTER, GERMANY

E-mail address: `strohm@math.uni-muenster.de`