Game Theory.
Mathematical and Conceptual Aspects

Jürgen Jost[1]

July 19, 2013

[1]Max Planck Institute for Mathematics in the Sciences, Inselstr.22, 04103 Leipzig, Germany, jost@mis.mpg.de

Contents

Chapter 1

Examples and basic concepts

1.1 An example

We begin with an example of a game. We have two players, Alice (abbreviated as A and referred to by the pronoun "she") and Bob (B, "he") each of which has the choice between two actions. For the choice a_i of A and b_j of B, the (i, j) entry in the table lists the pay-off of A before, and the one of B after the comma.

	b_1	b_2
a_1	3,2	1,3
a_2	2,1	0,0

(1.1.1)

So, how should the two players reason and act in order to maximize their pay-off, assuming that both know the structure of the game and the pay-off matrix? We assume at this point that the two players play simultaneously.
The first observation is that for player A, given an action of B, the first row is always better than the second. One says that action a_1 dominates a_2. So, let us reason that therefore she should disregard action a_2 and play a_1 in any case. When B realizes this, he should play b_2. Thus, A will get the pay-off 1, whereas B gets 3. This represents a so-called Nash equilibrium, meaning that neither player can unilaterally change her/his action without reducing her/his pay-off. If A changed from a_1 to a_2, but B keeps b_2, her pay-off would be reduced from 1 to 0. If B switched from b_2 to b_1, while A continues to play a_1, his pay-off would be reduced from 3 to 2.

Obviously, this equilibrium leaves B better off than A. If the game were played sequentially instead of simultaneously, with A playing first, she should choose a_2 in place of a_1, even though that action is dominated in the simultaneous game, as this would force B to play b_1, giving A the pay-off 2 which is higher than 1 as achieved in the Nash equilibrium.

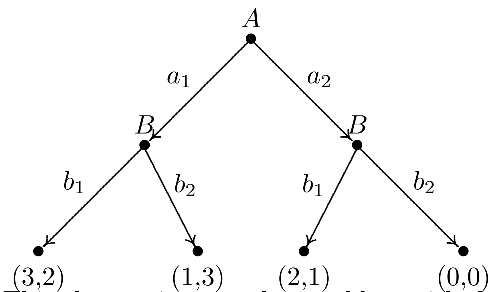

(1.1.2)

This observation reveals a problem with the dominance argument that A should disregard a_2 in the simultaneous game. After all, if she can make it plausible to B that she will play a_2, this would force him to play b_1. The dominance argument compares the outcome of each action against the same action of the opponent. The opponent, however, will react differently to the different actions of A. He will play b_2 against a_1, but b_1 against a_2, and the latter is better for A.

When, however, B can move first, he should play b_2, forcing A to play a_1, which is the above Nash equilibrium which is optimal for B.

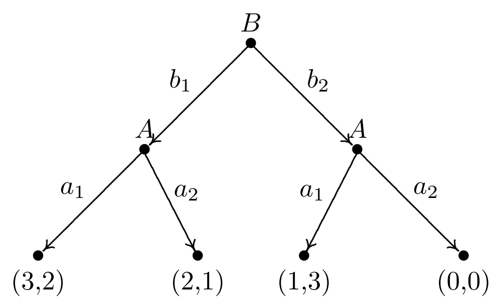

(1.1.3)

We now consider the situation where A can move first, but B cannot observe A's move, as indicated by the dashed line in the following diagram,

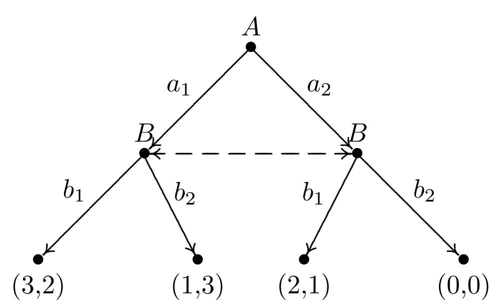

(1.1.4)

B might reason that A had played a_2 in order to force him to play b_1, and he should correspondingly do so. Now, however, A could think that because of this reasoning, B will play b_1 anyway, and he could therefore decide to play a_2 to maximize her pay-off. When then, in turn, B anticipates that reasoning, he could then play b_2. However, whenever A believes B to play b_2, she should play a_1. And hence the consistent beliefs repeat themselves. Whenever B believes

A to play a_1, he should play b_2, and whenever A expects that B will play b_2, she should play a_1. And since any chain of consistent higher order beliefs will eventually arrive at that point, it seems that (a_1, b_2) will be the only equilibrium. We need a formal definition of an equilibrium here, however, in order to substantiate that claim.
Incidentally, in this example, it helps B to be ignorant about A's move, as otherwise A could have forced him to play b_1 which leads to a worse pay-off for him. Thus, in such games, it can be disadvantageous to acquire more information about the opponent. – One should be careful with the interpretation of this finding, however. What harms B is not that he has the information, but the fact that A knows that he has that information. That is, A also possesses some additional information. If A did not know that B knows her move, then B would not be at a disadvantage. – Conversely, A should try to transmit the information about her move to B, but when B cannot verify the correctness of that information, we are again in a situation to which the preceding analysis applies.
We can gain further insight by modifying the pay-off matrix. For instance, when we consider

	b_1	b_2
a_1	1,2	1,3
a_2	2,1	0,0

(1.1.5)

that is, we lower the pay-off for A when playing a_1 against b_1, then it is prefectly reasonable for A to play a_2 to which B should react with b_1. Thus, when we remove the temptation for A to play a_1 instead of a_2 against b_1, we improve her position.

1.2 Formalization and further examples

The basic model is a simple game where two or more agents perform certain actions and receive a pay-off as a result of their own and their opponent's action. Each player wishes to maximize her pay-off, knowing that her opponent is trying to do the same. In particular, both players know each other's pay-offs.
We now introduce the basic formalism for finite games in normal form. In this and the following section, we discuss the case of two players only. This case already exhibits most of the pertinent phenomena while avoiding certain technical complications. The general case will then be treated in Section 1.4.

In order to better make formal use of the symmetry between the two players, we now label the players by $i = 1, 2$ (in order to make contact with the notation above, let 1 stand for A (Alice), 2 for B (Bob)). For a player i, we denote by $-i$ the other player, her opponent. Each player can perform a finite number of actions, or strategies as they are usually called. Thus, she has a finite set $S_i = \{1, \dots, m_i\}$ of pure strategies. $S := S_1 \times S_2$ is called the set of pure-strategy profiles, or the pure-strategy space. We then have the pay-off function

$$\pi_i : S \to \mathbb{R} \tag{1.2.1}$$

that assigns to player i her pay-off $\pi_i(s)$ for $s = (s_1, s_2) \in S$, that is, when i plays s_i and the opponent plays s_{-i}. We also consider $\pi = (\pi_1, \pi_2) : S \to \mathbb{R}^2$. Thus, a game between two players is given by the pair (S, π). This is also called the representation of the game in extensive form. The key point is that player can only determine her own strategy s_i while her pay-off also depends on the strategy s_{-i} of her opponent.
It is also convenient to write the pay-offs in matrix form. The pay-off matrix of player 1 is usually denoted by $A = (\pi_1(\alpha, \beta))_{\alpha=1,\dots,m_1,\beta=1,\dots,m_2}$, the one of player 2 by $B = (\pi_2(\alpha, \beta))_{\alpha=1,\dots,m_1,\beta=1,\dots,m_2}$. Thus, player 1 is the "row-player", player 2 the "column-player". The game is symmetric iff $B = A^t$.
It is also convenient to combine the two matrices into the single pay-off matrix

$$\Pi = (\pi_1(\alpha, \beta) | \pi_2(\alpha, \beta))_{\alpha=1,\dots,m_1,\beta=1,\dots,m_2}. \tag{1.2.2}$$

We now discuss some further
Examples:

1. We consider the following complementarity game: Each player can contribute 0,1,2 or 3 (units). When the sum of the two contributions is at least 3, each player receives 3 minus her own contribution, else both receive 0. The pay-off matrix then is

0,0	0,0	0,0	3,0
0,0	0,0	2,1	2,0
0,0	1,2	1,1	1,0
0,3	0,2	0,1	0,0

. (1.2.3)

We can already make a simple, but important, observation here: There is no point for player 1 to play the last strategy, that is, contribute 3, because her pay-off for any other strategy is always at least as high as the one for 3, and in several cases higher. The same applies for player 2, of course, as this game is symmetric. Therefore, player i can assume that player $-i$ will never play 3. Therefore, we only need to consider the reduced pay-off matrix

0,0	0,0	0,0
0,0	0,0	2,1
0,0	1,2	1,1

. (1.2.4)

Applying the same reasoning to this new pay-off matrix, it then will pay for neither player to play the first strategy, that is, contribute 0. Thus, we can reduce the pay-off matrix once more to arrive at

0,0	2,1
1,2	1,1

. (1.2.5)

Thus, each player would contribute 1 or 2. For each player, the best situation is if she herself contributes 1 while the other one contributes 2. When she knows, however, that the other one will contribute only 1, then she has no choice but to contribute 2 if she is to maximize her pay-off. That is, this represents a situation where neither player can change her strategy unilaterally without decreasing her pay-off.

2. We consider a game, the so-called matching pennies game, where each player has only two strategies, and player 1 gets a pay-off of 1 when both play the same, -1 when they play different strategies. For player 2, the situation is the opposite, that is, she gains 1 when the two play different strategies and looses 1 when they play the same. In particular, this is a zero-sum game, meaning that for each strategy profile, the pay-offs of the two players sum to 0. Thus, the pay-off matrix is

1,-1	-1,1
-1,1	1,-1

. (1.2.6)

In this game, whatever the actual strategy profile is, it is always advantageous for precisely one of the two players to change her strategy. In particular, if the game were to be played repeatedly, it would be disadvantageous for either player to always play the same strategy because the other one could then choose a winning strategy. More generally, it would be disadvantageous for either player to play in a manner that is predictable for her opponent. Thus, in this game, the best option for either player would be to play the two strategies randomly with probability 1/2 each. At least, if both of them played that way, then for neither of them it would carry an advantage to unilaterally change her strategy because the opponent would respond correspondingly.

3. This game is sometimes called the battle-of-the-sexes game. A girl (player 1) and her boy friend (player 2) would like to spend their time together, but each of them prefers a different activity. The girl likes to attend a heavy weight boxing fight (action 1) whereas the boy prefers to go to the fashion show (action 2) (I may have gotten the story wrong, but never mind). The pay-off matrix is

4,2	1,1
0,0	2,4

. (1.2.7)

For this game, either way of attending the same activity is an equilibrium where neither of them would benefit from changing her/his action unilaterally.
As represented, this game is non-symmetric, but there is an equivalent

symmetric form. Either player could insist on her/his preferred activity or yield to the preference of the partner. When insisting is action 1, yielding action 2, the pay-off matrix becomes

1,1	4,2
2,4	0,0

(1.2.8)

which is symmetric. In fact, the situation is now the same as in Example 1.

4. The next game has the structure of the so-called prisoner's dilemma.

4,4	2,6
6,2	3,3

. (1.2.9)

Even though it is better for both of them if they both play their first strategy ("cooperate") than when they both play the second one ("defect"), each player has an incentive to switch to defecting when the other player cooperates. The cooperating player would then be put at a disadvantage and should also switch to defecting to avoid that. Therefore, the only equilibrium is where they both defect and get only 3 each. This outcome looks somewhat paradoxical because it seems perfectly possible that they agreed to cooperate and received the pay-off of 4.

After having analyzed these examples in detail, let us once more emphasize the key point. Player i can only choose her own strategy, but her pay-off also depends on the strategy of her opponent s_{-i}. She should therefore select her own strategy s_i in such a way that the opponent strategy s_{-i} that is the opponent's best reply to s_i leaves i with the highest pay-off among all her possible strategy choices. In other words, i wants to be best off under the assumption that her opponent chooses her – the opponent's – best response. And the opponent applies the same reasoning.

Keeping these examples and observations in mind, we now develop some general concepts. In particular, the examples may have created the impression that even though a 2×2 game is completely described by 8 numbers, there is a bewildering multitude of possible phenomena. The question therefore arises whether any kind of classification is possible.
First of all, we observe that, so far, only pay-off differences were relevant while their absolute values did not matter. Also, in our analysis of Example 3, we have seen that relabelling the strategies simply interchanges some rows or columns, but does not change the game.In that way, many (but not all) games can be transformed into a symmetric form.
Classification of symmetric 2×2 games: We consider a symmetric 2×2 game with pay-off matrix

$$A = \begin{pmatrix} \pi_{11} & \pi_{12} \\ \pi_{21} & \pi_{22} \end{pmatrix}. \tag{1.2.10}$$

We observe that we obtain an equivalent game, in the sense of ranking of strategies when we subtract fixed numbers from each column. More precisely, it does not make any difference for the choice of strategies of the player i when her pay-off to a particular strategy of her opponent is changed by a fixed amount that does not depend on her own strategy. Thus, we may subtract π_{21} from the first and π_{12} from the second column, to obtain the diagonal pay-off matrix

$$A' = \begin{pmatrix} a_1 & 0 \\ 0 & a_2 \end{pmatrix} \tag{1.2.11}$$

with $a_1 := \pi_{11} - \pi_{21}, a_2 := \pi_{22} - \pi_{12}$. Thus, any symmetric 2×2 game is represented by the pair $(a_1, a_2) \in \mathbb{R}^2$. The classification is then in terms of the signs of a_1 and a_2. Category I consists of those games with $a_1 < 0, a_2 > 0$. Here, strategy 2 strictly dominates strategy 1 (a concept to be defined shortly, but probably already intuitively plausible). Thus, both players will play 2. This category includes the prisoner's dilemma as the reader will easily check. Of course, after this rearrangement, the paradox discussed above is gone. Category IV where $a_1 > 0, a_2 < 0$ is of course equivalent to this one by relabelling the two strategies. Category II means $a_1 > 0, a_2 > 0$. Here, for neither player it is advantageous to unilaterally deviate from strategy α if her opponent plays that strategy. Finally, Category III comprises the games with $a_1 < 0, a_2 < 0$. Here, playing different strategies is the best option from which no unilateral deviation pays off. Example 1 belongs to this category.

1.3 Mixed strategies and Nash equilibria

Except at the end of Example 2, we have considered pure strategies, that is, elements of the finite set S_i for player i. A *mixed* strategy of player i then is a probability distribution on S_i. It can be represented by a vector $p_i = (p_i^1, \dots, p_i^{m_i}) \in \mathbb{R}^{m_i}$ with $p_i^\alpha \geq 0$ for all α and $\sum_\alpha p_i^\alpha = 1$. The support of such a mixed strategy is defined as the set of those α with $p_i^\alpha > 0$. A pure strategy can then be considered as a mixed strategy whose support contains one single element of S_i. Because of the normalization $\sum_\alpha p_i^\alpha = 1$, the space of mixed strategies for player i is the $(m_i - 1)$-dimensional simplex

$$\Sigma_i := \{p_i \in \mathbb{R}_+^{m_i} : \sum_{\alpha=1}^{m_i} p_i^\alpha = 1\}. \tag{1.3.1}$$

The vertices $e_i^\alpha = (0, \dots, 0, 1, 0 \dots, 0)$ where the 1 is at the αth position correspond to the pure strategies, and Σ_i is the convex hull of these vertices. The strategies in the interior of Σ_i are precisely those whose support is all of S_i.
A mixed-strategy profile then is a pair $p := (p_1, p_2) \in \Sigma_1 \times \Sigma_2 =: \Sigma$, the space of mixed-strategy profiles. For such a mixed-strategy profile p, the pure-strategy profile $s = (s_1, s_2)$ is then played with probability

$$p(s) = p_1^{s_1} p_2^{s_2}, \tag{1.3.2}$$

and the expected value of the pay-off of player i is

$$\pi_i(p) := \sum_{s \in S} p(s)\pi_i(s). \tag{1.3.3}$$

We observe for later purposes that $\pi_i(p)$ is continuous as a function of p.
We can rewrite (1.3.3) as

$$\pi_i(p) = \sum_{\alpha=1}^{m_j} p_j^\alpha \pi_i(e_j^\alpha, p_{-j}), \text{ for } j = 1, 2, \tag{1.3.4}$$

that is, as a weighted sum of the pay-offs when j plays the pure strategy α. We may also write[1]

$$\pi_i(p) = \sum_{\alpha=1}^{m_1} \sum_{\beta=1}^{m_2} p_1^\alpha \pi_i(\alpha, \beta) p_2^\beta, \tag{1.3.5}$$

or in matrix notation,

$$\pi_1(p) = p_1 A p_2, \ \pi_2(p) = p_1 B p_2 = p_2 B^t p_1. \tag{1.3.6}$$

The pair (Σ, π) is called the representation of the game in strategic form, as opposed to the extensive-form representation (S, π) introduced earlier. Of course, this represents one and the same game, in terms of its rules, but in the strategic-form representation, the players are also allowed to play mixed strategies.
Let us consider Example 3, as given by (1.3.1), with mixed strategies. When player 2 plays strategy 1 with probability q and correspondingly 2 with probability $1-q$, then the expected pay-off for strategy 1 of player 1 is $4q + (1-q)$, and the expected pay-off for strategy 2 is $0 + 2(1-q)$. For $q < \frac{1}{5}$, the expected pay-off for strategy 2 is higher while for $q > \frac{1}{5}$, strategy 1 is better. For $q = \frac{1}{5}$, both strategies yield the same expected pay-off, and therefore, instead of playing a pure strategy, she could as well play any mixed strategy. Likewise, when player 1 plays strategy 1 with probability q', the expected pay-offs for the strategies 1 and 2 of player 2 are $2q' + 0$ and $q' + 4(1-q')$, resp. Here, the two expected pay-offs are the same for $q' = \frac{4}{5}$, in which case player 2 is indifferent. In particular, when player 1 plays the mixture with $q' = \frac{4}{5}$, and 2 with $q = \frac{1}{5}$, then neither of them can improve her pay-off by unilaterally changing strategy. Thus, this represents another equilibrium, in addition to the two given by the diagonal elements of (1.3.1).
Before addressing the general question of the existence of equilibria, we first turn to the elimination of non-optimal strategies.

Definition 1.3.1. We say that the (mixed) strategy $q_i \in \Sigma_i$ of player i strictly dominates p_i if

$$\pi_i(q_i, r_{-i}) > \pi_i(p_i, r_{-i}) \text{ for any } r_{-i} \in \Sigma_{-i}. \tag{1.3.7}$$

[1] Note that $\pi_i(\alpha, \beta) = \pi_i(e_j^\alpha, e_{-j}^\beta)$, just as a piece of alternative notations employed.

We say that q_i weakly dominates p_i if

$$\pi_i(q_i, r_{-i}) \geq \pi_i(p_i, r_{-i}) \text{ for any } r_{-i} \in \Sigma_{-i}, \tag{1.3.8}$$

with strict equality for at least one r_{-i}.
p_i is said to be undominated if it is not weakly dominated by any other strategy.

As we have already seen in our analysis of Example 1, weakly dominated pure strategies can be eliminated from consideration, until arriving at a reduced version of the game without weakly dominated pure strategies. We should note that this reduction makes crucial use of the rationality assumption, that is, no player will choose strategies that leave her unconditionally worse off than others, and that each player knows that her opponent is behaving that way.There is one sublety here: When we first eliminate strongly dominated strategies, then some other strategies may cease to be weakly dominated, because the r_{-i} where we have the strict inequality in (1.3.8) may belong to a strongly dominated, hence eliminated, strategy of the opponent.

We now come to the concept of a Nash equilibrium.

Definition 1.3.2. A pure-strategy profile $s^* = (s_1^*, s_2^*)$ is called a Nash equilibrium if for both $i = 1, 2$

$$\pi_i(s^*) \geq \pi_i(s_i, s_{-i}^*) \text{ for all } s_i \in S_i. \tag{1.3.9}$$

This means that no player can do better by changing her strategy when the other one keeps her strategy. In other words, no player can gain from a unilateral move. Thus, unless the players would or could make some coordinated move – which, however, is not allowed by the rules of the game –, they should stick to their strategy. In this sense, this represents an equilibrium.
As we have seen from Example 2, however, such a Nash equilibrium in pure strategies need not exist. Therefore, Definition 1.3.2 is generalized to

Definition 1.3.3. A mixed-strategy profile $\sigma^* = (\sigma_1^*, \sigma_2^*)$ is called a Nash equilibrium if for both $i = 1, 2$

$$\pi_i(\sigma^*) \geq \pi_i(\sigma_i, \sigma_{-i}^*) \text{ for all } \sigma_i \in \Sigma_i. \tag{1.3.10}$$

Epistemically, the concept might look somewhat strange. In fact, for a Nash equilibrium, every player assumes that he is the only one that makes a change of strategy while all the others keep their strategies fixed. And, they all simultaneously assume that, that is, they all assume that they make some move while the others don't. So, this looks contradictory. However, since everybody assumes that the others don't move and consequently he cannot improve his pay-off, nobody moves, indeed, at a Nash equilibrium. Thus, an assumption that is contradictory outside the equilibrium is mutually confirmed at the equilibrium. In particular, the rational reasoning of the players leading to the equilibrium needs to be consistent only at the equilibrium itself, but not outside it.

The fundamental theorem of Nash tells us that such an equilibrium does exist, indeed. In order to prepare that result, we now develop an approach via best responses. For each player, we have the best-response correspondance given by all the best responses to each possible strategy of her opponent:

$$\begin{aligned} \rho_i : \Sigma_{-i} &\rightrightarrows \Sigma_i \\ \sigma_{-i} &\mapsto \{\sigma_i \in \Sigma_i : \pi_i(\sigma_i, \sigma_{-i}) \geq \pi_i(\tau_i, \sigma_{-i}) \text{ for all } \tau_i \in \Sigma_i\} \end{aligned} \tag{1.3.11}$$

Here, the symbol $\rightrightarrows$ indicates that the image of an element $\sigma_{-i} \in \Sigma_{-i}$ is in general a subset, in place of a single element, of Σ_i, because the best response need not be unique. We then have the best-response correspondance

$$\rho := \rho_1 \times \rho_2 : \Sigma_2 \times \Sigma_1 \rightrightarrows \Sigma_1 \times \Sigma_2, \tag{1.3.12}$$

which, by changing factors, we may consider as a correspondance

$$\rho : \Sigma \rightrightarrows \Sigma. \tag{1.3.13}$$

We then observe

Lemma 1.3.1. *A profile $\sigma^* = (\sigma_1^*, \sigma_2^*)$ is a Nash equilibrium iff*

$$\sigma^* \in \rho(\sigma^*), \tag{1.3.14}$$

that is, iff for each player, her strategy is a best response to her opponent's.

The *proof* is obvious. The advantage of (1.3.14) is that it characterizes a Nash equilibrium as a fixed point of the correspondance ρ. This suggests to invoke a fixed-point theorem to obtain the existence of a Nash equilibrium.

Theorem 1.3.1. (Nash): *Every finite game in strategic form possesses a Nash equilibrium.*

Proof. The idea of the proof is to reduce the result to a general fixed-point theorem (which we shall not prove here, but which is an extension of Brouwer's fixed point theorem and can be proved along similar lines as the latter, see e.g. [25]),

Theorem 1.3.2. (Kakutani): *Let $K \subset \mathbb{R}^d$ be compact, convex, $\neq \emptyset$. Let $\rho : K \rightrightarrows K$ be a correspondance for which, for all $x \in K$, $\rho(x)$ is convex and $\neq \emptyset$, and which satisfies:*

$$\begin{aligned} &\textit{If } \lim_{n\to\infty} x_n = x_0 \textit{ for some } x_0 \in K \textit{ and some sequence } (x_n) \subset K, \\ &\quad \textit{then every sequence } (y_n) \subset K \textit{ with } y_n \in \rho(x_n) \textit{ for all } n \\ &\quad\quad \textit{has some limit } y_0 \textit{ with } y_0 \in \rho(x_0). \end{aligned} \tag{1.3.15}$$

(This property is called upper hemi-continuity. It is equivalent to the requirement that the graph of ρ be closed as a set.)
Then ρ has a fixed point, i.e., there exists some $x^* \in K$ with $x^* \in \rho(x^*)$.

Most of the assumptions of the Kakutani fixed-point theorem are obviously satisfied for finite games in strategic form. The simplex product Σ which we take as the K in the theorem is compact, convex, non-empty. By construction, the pay-off function π_i is continuous and therefore, $\pi_i(., \sigma_{-i})$ attains its maximum in the compact set Σ_i, for every σ_{-i}. Therefore, for every $\sigma \in \Sigma$, $\rho(\sigma) \neq \emptyset$. Also, since $\pi_i(., \sigma_{-i})$ is linear for every σ_{-i}, convex combinations of best responses are again best responses, which implies the convexity of the image $\rho(\sigma)$ for all σ. It remains to verify the continuity property (1.3.15). It suffices to check this for every component $\rho_i : \Sigma_{-i} \rightrightarrows \Sigma_i$. Let $(\xi_n) \subset \Sigma_{-i}$ converge to ξ_0. Let $\eta_n \in \rho_i(\xi_n)$ for every n. This means that

$$\pi_i(\eta_n, \xi_n) \geq \pi_i(\eta, \xi_n) \text{ for every } \eta \in \Sigma_i. \tag{1.3.16}$$

By the compactness of Σ_i, after taking a subsequence, $\lim_{n\to\infty} \eta_n = \eta_0$ exists. (1.3.16) and the continuity of π_i imply

$$\pi_i(\eta_0, \xi_0) \geq \pi(\eta, \xi_0) \text{ for every } \eta \in \Sigma_i. \tag{1.3.17}$$

This means that $\eta_0 \in \rho_i(\xi_0)$ which is the required continuity property in Kakutani's theorem. Thus, we may apply that theorem to deduce the Theorem of Nash. □

Remark: Many fixed-point theorems are proved by simply iterating a map or a correspondance and obtaining the fixed point in the limit. As Example 2 shows, however, in the situation of games as considered here, such an iteration need not converge, but may oscillate forever.

Some references for this section and the subsequent on iterated games are [6, 5, 9, 20, 24, 21].

1.4 Several players: The Brouwer fixed point theorem and the Nash equilibrium theorem

As already explained, the Nash theorem is an easy consequence of variants of the Brouwer fixed point theorem. In this subsection, we provide a complete treatment, using the references [25, 2]. The starting point is combinatorial lemma of Sperner. We consider the n-dimensional simplex Σ^n with vertices $p_1, \dots, p_{n+1}$. It can be iteratively constructed from its subsimplices or faces. This goes as follows. If we take two of its vertices, $p_i, p_j, i \neq j$, their convex hull is a 1-dimensional simplex p_{ij}, an edge of Σ^n. Likewise, the convex hull of three vertices is a triangle, and so on. For each $p \in \Sigma^n$, we let $\sigma(p)$ denote the unique subsimplex of smallest dimension containing p.
A simplicial decomposition of a simplex Σ is the representation of Σ as a finite union of simplices of the same dimension any two of which are either disjoint or intersect in one common face.

Lemma 1.4.1. (Sperner): *Let* $\Sigma^n = \bigcup_k \Sigma_k$ *be a simplicial decomposition. To every vertex* p *of every* Σ_k*, we assign some number* $i \in \{1, \dots, n+1\}$ *in such a manner that the vertex* p_i *of* Σ^n *is contained in* $\sigma(p)$*. Then there exists a Sperner simplex among the* Σ_k*, that is, a simplex whose vertices cover all the labels* $1, \dots, n+1$*.*

Proof. The proof procedes by induction on n to show that the number of Sperner simplices is odd, and therefore $\neq 0$. The case $n = 0$ is trivial. $n = 1$ is an easy exercise. We present here the case $n = 2$ which gives a reasonable idea of the general, combinatorially more complicated proof. We call an edge of a 2-dimensional simplex Σ_k distinguished if its vertices carry the labels $1, 2$. The distinguished edges in the interior occur in pairs whereas their number on the boundary is odd, by the step for $n = 1$. Thus, their total number is odd. Since each 2-dimensional simplex carries either one or zero or two distinguished sides, there must exist one such simplex with precisely one distinguished side. The third vertex of that simplex then has to carry the label 3, and therefore, we have found a Sperner simplex. Also, since the number of distinguished edges is odd, the number of simplices carrying precisely one of them in their boundary is odd as well. □

We now derive the lemma of Knaster, Kuratowski and Mazurkiewicz which is the essential ingredient of the proofs of the fixed point theorems. Let V be a topological vector space. (That is, V carries both a linear structure, so that we can take convex combinations, and a topology, so that we can speak about closed or compact subsets.) For $x_1, \dots, x_m \in V$, we define their convex hull as

$$C(x_1, \dots, x_m) := \{x = \sum_{\mu=1}^{m} \lambda_\mu x_\mu : \lambda_\mu \geq 0 \text{ for all } \mu, \sum_{\mu=1}^{m} \lambda_\mu = 1\}. \tag{1.4.1}$$

Lemma 1.4.2. (Knaster, Kuratowski, Mazurkiewicz): *Let* X *be a nonempty subset of* V*. Suppose that to each* $x \in X$*, there is assigned a nonempty compact subset* $\psi(x)$ *of* V*, such that for every finite subset* $\{x_1, \dots, x_m\}$ *of* X

$$C(x_1, \dots, x_m) \subset \bigcup_{\mu=1}^{m} \psi(x_\mu). \tag{1.4.2}$$

Then

$$\bigcap_{x \in X} \psi(x) \neq \emptyset. \tag{1.4.3}$$

In particular, (1.4.2) entails

$$x \in \psi(x) \text{ for all } x \in X, \tag{1.4.4}$$

and this may be helpful to understand the subsequent applications of Lemma 1.4.2.

Remark: In fact, it suffices that one of the sets $\psi(x)$ be compact and the others closed. This follows because one can then replace $\psi(x)$ by the compact set $\psi(x) \cap \psi(x_0)$ which is nonempty as a consequence of (1.4.2).

Proof. It suffices to treat the case where X is finite. The reason is the following general property of compact sets: When K is compact and $A_j, j \in \mathbb{N}$, are closed and when every intersection of finitely many sets A_ν with K is nonempty, then also $\bigcap_{\nu \in \mathbb{N}} A_\nu \cap K \neq \emptyset$. (This follows simply by taking a sequence $(y_j)_{j \in \mathbb{N}}$ such that $y_j \in \bigcap_{\nu=1}^{j} A_\nu \cap K$. Since K is compact, a subsequence of (y_j) converges to some $y \in K$ which then by construction is contained in every A_ν, because these sets are closed.)
By a small perturbation, if necessary, we may then assume that the finitely many points $x_1, \dots, x_{n+1}$ of X are linearly independent, that is, they span a topological n-simplex Σ in V. For $k = 1, 2 \dots$, we then consider simplicial decompositions of Σ such that the diameter of all simplices goes to zero as $k \to \infty$. Because of (1.4.2), we may label the vertices of the simplices in each decomposition in such a manner that a vertex with label i is contained in $\psi(x_i)$. By Lemma 1.4.1, we then find a Sperner simplex, that is, one with all $n+1$ labels. By construction, this means that each of the sets $\psi(x_i)$ contains at least one of the vertices $p_1^k, \dots, p_{n+1}^k$ of that simplex. When we let $k \to \infty$, after selection of a subsequence, the vertices $p_1^k, \dots, p_{n+1}^k$ converge to a single point p because the diameter of the simplex goes to 0. Since the $\psi(x_i)$ are closed, p then is contained in all of them. Therefore, their intersection is nonempty. □

Even though we shall not need the statement itself in the sequel, we now state and derive the famous Brouwer fixed point theorem.

Theorem 1.4.1. (Brouwer): *Every continuous map $f : M \to M$ of a topological space homeomorphic to the n-dimensional simplex Σ^n has a fixed point.*

Proof. It suffices to treat the case where M is the simplex Σ^n. We ket X be the set of the vertices of Σ^n, denoted by $x_1, \dots, x_{n+1}$. Every $x \in \Sigma^n$ then can be represented *uniquely* as a convex combination of these vertices

$$x = \sum_{\mu=1}^{n+1} \lambda_\mu(x) x_\mu \text{ with } \lambda_\mu(x) \geq 0 \text{ for all } \mu \text{ and } \sum_{\mu=1}^{n+1} \lambda_\mu(x) = 1. \tag{1.4.5}$$

We put

$$\psi(x_\mu) := \{x \in \Sigma^n : \lambda_\mu(f(x)) \leq \lambda_\mu(x)\}. \tag{1.4.6}$$

These sets are closed, by continuity of $f(x)$ and the $\lambda_\mu(x)$ and bounded, hence compact. They also satisfy (1.4.2), because as $0 \leq \lambda_\mu(x), \lambda_\mu(f(x)) \leq 1$ and $\sum \lambda_\mu(x) = 1 = \sum \lambda_\mu(f(x))$, at least one of the $\lambda_\mu(f(x))$ must be $\leq \lambda_\mu(x)$. Therefore, we may apply Lemma 1.4.2 to obtain some $x \in \bigcap_{\mu=1}^{n+1} \psi(x_\mu)$. This means that

$$\lambda_\mu(f(x)) \leq \lambda_\mu(x) \text{ for all } \mu. \tag{1.4.7}$$

Since these coefficients sum to 1 on both sides of (1.4.7), they then must all be equal, $\lambda_\mu(f(x)) = \lambda_\mu(x)$ for all μ. This, however, implies $f(x) = x$, by the uniqueness of the representation (1.4.5). Thus, x is the desired fixed point. □

Theorem 1.3.2 can also be proved along these lines; we leave out the details. In fact, there are many results that are all equivalent to Brouwer's fixed point theorem, and we now turn to another such result, the inequality of Fan. Subsequently, we shall derive the Nash equilibrium theorem from that result.

Theorem 1.4.2. (Fan): *Let X be a compact, convex, nonempty subset of a topological vector space. Let $f : X \times X \to \mathbb{R}$ satisfy:*

1. *Concavity in the first argument:*

$$f(\sum_{\mu=1}^{m} \lambda_\mu x_\mu, y) \geq \sum_{\mu=1}^{m} \lambda_\mu f(x_\mu, y) \text{ for all } \lambda_\mu \geq 0 \text{ with } \sum_{\mu=1}^{m} \lambda_\mu = 1 \quad (1.4.8)$$

   for all $x_\mu, y \in X$.

2. *Lower semicontinuity in the second argument:*

$$f(x,y) \leq \liminf_{n\to\infty} f(x, y_n) \quad (1.4.9)$$

   whenever the y_n converge to $y \in X$.

Then

$$\min_{y\in X} \sup_{x\in X} f(x,y) \leq \sup f(x,x). \quad (1.4.10)$$

Proof. We put $M := \sup_{x\in X} f(x,x)$ and

$$\psi(x) := \{y \in X : f(x,y) \leq M\}. \quad (1.4.11)$$

$\psi(x)$ is closed by Assumption 2, hence compact, because contained in the compact set X. We now claim

$$C(x_1, \ldots, x_m) \subset \bigcup_{\mu=1}^{m} \psi(x_\mu) \quad (1.4.12)$$

for any $x_1, \ldots, x_m \in X$. If not, there exists some convex combination $x = \sum_{\mu=1}^m \lambda_\mu x_\mu$ with $x \notin \bigcup_{\mu=1}^m \psi(x_\mu)$. This means that $f(x_\mu, x) > M$ for all μ. Assumption 1 then implies $f(x,x) > M$, contradicting the definition of M, however. Therefore, (1.4.12) holds. We may then apply Lemma 1.4.2 to find some $y \in \bigcap_x \psi(x)$, that is, $f(x,y) \leq M = \sup_x f(x,x)$ which is (1.4.10). □

We can now prove the Nash Equilibrium Theorem 1.3.1, in fact in a setting that is more general than that discussed in Section 1.3. We have n players (instead of 2), and the strategy space of each player is some compact, convex, nonempty set K_i in some topological vector space (instead of being simply a simplex Σ_i). The pay-off function

$$\pi_i : K_1 \times \cdots \times K_n \to \mathbb{R} \quad (1.4.13)$$

of the ith player is only required to be continuous in all variables and concave in the strategies of i, that is,

$$\pi_i(x_1,\ldots,x_{i-1},\sum_\mu \lambda_\mu \xi_\mu, x_{i+1},\ldots,x_n) \geq \sum_{\mu=1}^m \lambda_\mu \pi_i(x_1,\ldots,x_{i-1},\xi_\mu,x_{i+1},\ldots,x_n) \tag{1.4.14}$$

for all $\xi_\mu \in K_i$. In (1.3.3), we had, more restrictively, required the extension of the utility function to mixtures to be a linear extension of that on pure actions. Here, in contrast, the utility derived from a mixture of pure actions could be higher than the corresponding convex combinations of the utilities of those pure actions.

Theorem 1.4.3. (Nash): *Under the assumptions just specified, there exists a Nash equilibrium, that is, some x^* with $x_i^* \in K_i$ for every i that satisfies*

$$\pi_i(x_1^*,\ldots,x_n^*) = \max_{x_i \in K_i} \pi_i(x_1^*,\ldots,x_{i-1}^*,x_i,x_{i+1}^*,\ldots,x_n^*). \tag{1.4.15}$$

Proof. For $x, y \in X := K_1 \times \cdots \times K_n$, we put

$$f(x,y) := \sum_{i=1}^n (\pi_i(y_1,\ldots,y_{i-1},x_i,y_{i+1},\ldots,y_n) - \pi_i(y)). \tag{1.4.16}$$

f then is concave in the first and continuous in the second argument and satifies $f(x,x) = 0$. By Theorem 1.4.2, there exists some $x^* \in X$ with

$$f(x,x^*) \leq 0 \text{ for all } x \in X. \tag{1.4.17}$$

This means, when applying it to $x = (x_1^*,\ldots,x_{i-1}^*,x_i,x_{i+1}^*,\ldots,x_n^*)$

$$\pi_i(x_1^*,\ldots,x_n^*) \geq \pi_i(x_1^*,\ldots,x_{i-1}^*,x_i,x_{i+1}^*,\ldots,x_n^*) \text{ for all } i, \tag{1.4.18}$$

because for that particular choice of x, all terms in (1.4.16) with $j \neq i$ are zero. (1.4.18) is (1.4.15). □

1.5 Quantal response equilibria and another proof of the Nash equilibrium theorem

In fact, the topological aspects behind the Nash equilibrium theorem are simpler than the preceding may suggest and a geometrically more intuitive proof can be found. In order to understand the topological situation, let us again look at the case of two players. This case already reveals the geometric picture, and what we shall describe in this simplest situation readily generalizes to the general case.

If the response correspondance from (1.3.1), here recalled as

$$\begin{aligned} \rho_i : \Sigma_{-i} &\rightrightarrows \Sigma_i \\ \sigma_{-i} &\mapsto \{\sigma_i \in \Sigma_i : \pi_i(\sigma_i,\sigma_{-i}) \geq \pi_i(\tau_i,\sigma_{-i}) \text{ for all } \tau_i \in \Sigma_i\}, \end{aligned} \tag{1.5.1}$$

were single-valued, it could just be depicted as a graph on Σ_{-i} with values in Σ_i. Likewise, the corresponding map ρ_{-i} for $-i$ would then be a graph over Σ_i with values in Σ_{-i}. It then seems geometrically clear that any two such graphs need to intersect somewhere. Just consider the case where each players has only two pure strategies. Then Σ_i and Σ_{-i} both reduce to the unit interval $[0, 1]$. And the product $\Sigma = \Sigma_i \times \Sigma_{-i}$ then is the unit square $[0, 1] \times [0, 1]$. ρ_i then is a map from the second to the first factor, whereas ρ_{-i} maps the first to the second factor. And clearly, their graphs then have to intersect. Such an intersection point of the graphs, however, is a fixed point for the best-response map ρ of (1.3.12), hence a Nash equilibrium, by Lemma 1.3.1. This principle holds in any dimension. A graph from Σ_{-i} to Σ_i and one from Σ_i to Σ_{-i} need to intersect. And the principle also generalizes to more than two players. For a player i, $-i$ then stands for the collection of all her opponents, that is, the set of all other players in the game. Likewise, Σ_{-i} then is the product of the strategy simplices of those players, i.e.,

$$\Sigma_{-i} := \prod_{j \neq i} \Sigma_j. \tag{1.5.2}$$

The best response correspondance, or as we assume for the moment, map, then is defined as before in (1.5.1). It simply gives the best response to the combination of the strategies of all the opponents. When there are n players, we have the total strategy space

$$\Sigma := \prod_{j=1}^{n} \Sigma_j \tag{1.5.3}$$

of dimension $\sum_j (m_j - 1)$ (where m_j is the number of pure actions of j), and for each i, we have the best response map $\rho_i : \Sigma_{-i} \to \Sigma_i$. The graph of ρ_i has dimension equal to the dimension of its domain Σ_{-i}, that is, $\sum_{j \neq i}(m_j - 1)$. Its codimension is therefore simply $m_i - 1$. Since the sum of all these codimensions equals the dimension of Σ, these graphs should then intersect in finitely many points. This is a general principle in algebraic topology, see e.g. [19]. Of course, we have to make sure that this finite number of intersection points is not 0. In this section, we shall provide a homotopy argument to see this. This argument will at the same time address the slight technical problem that in general, in the Nash situation, the ρ_i are only correspondances, but not maps, because the best response need not be unique. The argument will consist in approximating such a correspondance by a genuine map.

This approximation will be of independent interest. In fact, the idea is to abandon the requirement of strict rationality. Players will no longer be able to play deterministically, but only stochastically. Therefore, they can make errors, but it is assumed that the probability of a choice of strategy should depend on the utility it yields. Strategies yielding higher utilities should be chosen with higher probability.

When the concept of rationality of game theory is abandoned or at least relaxed, we need some other rule for transforming the utilities of the players into action choices. Here, we shall utilize a particular rule for determining the probabilities of i from her utility values, the concept of a quantal response equilibrium (QRE) as introduced by McKelvey and Palfrey in [14]. As the discerning reader will realize, only certain continuity properties and limiting behaviors of this rule will be needed in the sequel, but it may be helpful to work with some definite rule, and the QRE is particularly simple. The QRE rule, then, stipulates that

$$p_i^\alpha = P_i^\alpha(p_{-i}^{\boldsymbol{\gamma}}; \lambda) := \frac{1}{Z_i(\lambda)} \exp(\lambda \sum_{\boldsymbol{\gamma}} \pi_i(\alpha, \boldsymbol{\gamma}) p_{-i}^{\boldsymbol{\gamma}}) \tag{1.5.4}$$

for each player i, where $\boldsymbol{\gamma}$ stands for the collection of strategies of all the opponents. $0 \le \lambda \le \infty$ is a parameter, to be utilized below, and

$$Z_i(\lambda) := \sum_\delta \exp(\lambda \sum_{\boldsymbol{\gamma}} \pi_i(\delta, \boldsymbol{\gamma}) p_{-i}^{\boldsymbol{\gamma}}) \tag{1.5.5}$$

is the normalization factor that ensures $\sum_\alpha p_i^\alpha = 1$ for each i. Because of this normalization, the range from which player i can choose her probabilities is the $(m_i - 1)$-dimensional simplex Σ_i.
We note that, for $\lambda < \infty$, the probabilities in (1.5.4) are all strictly positive, that is,

$$P_i^\alpha(p_{-i}^{\boldsymbol{\gamma}}; \lambda) > 0 \text{ for all } i, \alpha, \boldsymbol{\gamma}. \tag{1.5.6}$$

Thus, any strategy will be chosen with some positive probability. When λ is large, suboptimal strategies will only be chosen with rather small probabilities, however. λ may be taken to reflect the degree of rationality of the players. In fact, we could even give each player i an individual such parameter λ_i.

In any case, this means that the map

$$P_i^\alpha(.; \lambda); \Sigma_{-i} \to \Sigma_i \tag{1.5.7}$$

always maps Σ_{-i} to the interior of Σ_i.
An equilibrium is achieved when (1.5.4) holds for all players i simultaneously. Geometrically, this means that we look for a point of intersection of the graphs of the functions $P_i^\alpha(p_{-i}^{\boldsymbol{\gamma}})$, $i = 1, \dots, n$. Since each such graph is of codimension $m_i - 1$ in the space Σ of (1.5.3) of dimension $\sum_j (m_j - 1)$, these graphs generically intersect in a finite collection of points. Here is a simple intuitive argument to see this: Each such graph comes from a function defined on the unit quadrant in some subvector space of dimension $m_i - 1$ in the vector space of dimension $\sum_j (m_j - 1)$. These subvector space intersect in a single point, the coordinate origin. Then also a collection of graphs over those subvector spaces should typically intersect in one or several points.
At this point, however, this collection of intersection points is possibly empty. We shall show that the algebraic intersection number is always 1, for each value of λ, so that there always has to be at least one intersection point. Moreover, we shall show that for $\lambda \to \infty$, such intersection points converge to Nash equilibria of the original game, thereby showing the existence of such Nash equilibria.

We make the following observations. For elementary properties of algebraic intersection numbers, we refer to [19] or any other good textbook on algebraic topology.

1. For $\lambda = 0$, (1.5.4) becomes
$$p_i^\alpha = \frac{1}{m_i} \text{ for all } i, \alpha. \tag{1.5.8}$$
Therefore, with orientations appropriately chosen, the algebraic intersection number for the graphs at $\lambda = 0$ is 1.

2. For $0 \le \lambda < \infty$, an intersection cannot take place at any boundary point of Σ, as observed above as a consequence of (1.5.6), the P_i do not map to boundary points.

3. Since the graphs of $P_i^\alpha(p_{-i}^{\boldsymbol{\gamma}}; \lambda)$ depend continuously on λ and since by 2., no intersection point can disappear at the boundary, the algebraic intersection number of the graphs is 1 for all $0 \le \lambda < \infty$. In particular, there is always at least one intersection point of those graphs in the interior of Σ.

Thus, we have shown that for any $\lambda < \infty$, there always exists a QRE. We now consider the limits of the QRE graphs (1.5.4) and their intersection points for $\lambda \to \infty$, in order to see that these QREs always have limits which are Nash equilibria.

4. If for a given collection $p_{-i}^{\boldsymbol{\gamma}}$, there is a unique α_0 with
$$\alpha_0 = \text{argmax}_\alpha \sum_{\boldsymbol{\gamma}} \pi_i(\alpha, \boldsymbol{\gamma}) p_{-i}^{\boldsymbol{\gamma}}, \tag{1.5.9}$$
then
$$\lim_{\lambda \to \infty} P_i^{\alpha_0}(p_{-i}^{\boldsymbol{\gamma}}; \lambda) = 1 \text{ and } \lim_{\lambda \to \infty} P_i^{\beta}(p_{-i}^{\boldsymbol{\gamma}}; \lambda) = 0 \text{ for } \beta \neq \alpha_0. \tag{1.5.10}$$
This follows directly from (1.5.4), (1.5.5).

5. The set of tuples $(p_{-i}^{\boldsymbol{\gamma}})$ for which there exists more than one α' with
$$\alpha' = \text{argmax}_\alpha \sum_{\boldsymbol{\gamma}} \pi_i(\alpha, \boldsymbol{\gamma}) p_{-i}^{\boldsymbol{\gamma}} \tag{1.5.11}$$
is a union of hyperplanes in Σ_{-i} because the functions $\sum_{\boldsymbol{\gamma}} \pi_i(\alpha, \boldsymbol{\gamma}) p_{-i}^{\boldsymbol{\gamma}}$ are linear wr.t. the $p_{-i}^{\boldsymbol{\gamma}}$. In fact, if there exist two such maximizing values α_1, α_2, then we have the equation $\sum_{\boldsymbol{\gamma}}(\pi_i(\alpha_1, \boldsymbol{\gamma}) - \pi_i(\alpha_2, \boldsymbol{\gamma})) p_{-i}^{\boldsymbol{\gamma}} = 0$, and if we ignore the trivial case where this is satisfied for all $p_{-i}^{\boldsymbol{\gamma}}$, the solution set is of codimension 1, that is, a hyperplane. We call these hyperplanes singular. When crossing such a singular hyperplane, the maximizing α'

in (1.5.11) changes, because the expression $\sum_\gamma(\pi_i(\alpha_1,\gamma)-\pi_i(\alpha_2,\gamma))p^\gamma_{-i}$ changes its sign upon crossing. In particular, on the two sides of such a hyperplane, the maximizing α_0 in (1.5.9) is different. Thus, also the α_0 with (1.5.10) changes.
On such a singular hyperplane, that is when the opponents play a singular p^γ_{-i}, player i is indifferent between the different α' satisfying (1.5.11) as they all yield the same utility.

6. In particular, when we restrict the functions P^α_i to a line L (a one-dimensional subspace) in Σ_{-i} that is transversal to all the singular hyperplanes, we find some α_0 for which $\lim_{\lambda\to\infty} P^{\alpha_0}_i(p^\gamma_{-i};\lambda)$ jumps from 0 to 1 at the intersection π of L with such a singular hyperplane.

7. Therefore, the entire line $(\pi, 0 \le p^{\alpha_0}_i \le 1)$ is contained in the pointwise limit set of the graphs of the functions $P^{\alpha_0}_i(p^\gamma_{-i};\lambda)$ for $\lambda\to\infty$.

8. Thus, the limit set of the graphs of the $P^\alpha_i(p^\gamma_{-i};\lambda)$ for $\lambda\to\infty$ consists either solely of the set $\Sigma_{-i}\times(p^\alpha_i=0)$ or of $\Sigma_{-i}\times(p^\alpha_i=1)$ or of some regions where $p^\alpha_i=0$ and some where $p^\alpha_i=1$ connected by pieces of hyperplanes in Σ_{-i} times sets $0\le p^\alpha_i\le 1$ for those α for which the probabilities jump across the hyperplane.

9. Intersection points of the graphs of the functions $P^{\alpha_0}_i(p^\gamma_{-i};\lambda)$ for the different players i converge to intersection points of their pointwise limit sets. Some or all of these limit intersection points may lie in the boundary of Σ. It is also possible, however, that some of them lie in the connected pieces described in 8. where some or all of the p^α_i are undetermined.

We now see the *proof* of the Nash theorem: By 8., the limits of the QRE graphs yield the best response sets for the players. By 9., the intersection of these best response sets is not empty. Since by definition, an intersection point of the best response sets of all players is a Nash equilibrium, the existence of such an equilibrium follows. □

Let us consider some **examples** in details, to make the idea of the preceding topological argument clear.

1. We start with a simple version of the battle-of-the-sexes game,

2,1	0,0
0,0	1,2

, (1.5.12)

i being the row and $-i$ being the column player. The first move of each player, that is, up for i and left for $-i$, will be denoted by $+$, the second on, down for i and right for $-i$, by $-$. An appropriately indexed p will again stand for the corresponding probabilities.
We note that this game is symmetric in the sense that it remains invariant if we simultaneously interchange the two players and the labels of the move, as the pay-off of i for some combination of moves equals the pay-off

of $-i$ when the opposite moves are selected. This symmetry will then extend to the QRE situation.

The Nash equilibria can be easily understood in the following geometric manner. In general, the expected pay-offs of i and $-i$ are given by

$$(\pi_i(x_i^+,x_{-i}^+)p_{-i}^+ + \pi_i(x_i^+,x_{-i}^-)p_{-i}^-)p_i^+ + (\pi_i(x_i^-,x_{-i}^+)p_{-i}^+ + \pi_i(x_i^-,x_{-i}^-)p_{-i}^-)p_i^- \tag{1.5.13}$$

and

$$(\pi_{-i}(x_{-i}^+,x_i^+)p_i^+ + \pi_{-i}(x_{-i}^+,x_i^-)p_i^-)p_{-i}^+ + (\pi_{-i}(x_{-i}^-,x_i^+)p_i^+ + \pi_{-i}(x_{-i}^-,x_i^-)p_i^-)p_{-i}^-. \tag{1.5.14}$$

Since $p_i^- = 1 - p_i^+$, for each fixed probabilities of her opponent, the expected pay-off of i is a linear function of her p_i^+ which is constrained to the unit interval $[0,1]$. When this linear function has a negative slope, the maximum is achieved for $p_i^+ = 0$, when the slope is positive, at $p_i^+ = 1$, and when the slope is 0, she is indifferent. In the present example, the slope is positive for $p_{-i}^+ > 1/3$, negative for $p_{-i}^+ < 1/3$, and 0 for $p_{-i}^+ = 1/3$. Since the game is symmetric, the corresponding holds for $-i$. Therefore, the Nash equilibria are given by

$$\begin{aligned} p_i^+ = 1 = p_{-i}^+, &\quad \text{the right endpoints of two lines with positive slope} \\ p_i^+ = 0 = p_{-i}^+, &\quad \text{the left endpoints of two lines with negative slope} \\ p_i^+ = 2/3, p_{-i}^+ = 1/3, &\quad \text{the intersection of the two lines with 0 slope.} \end{aligned} \tag{1.5.15}$$

At the mixed equilibrium, the expected pay-off for each player is 2/3, according to (1.5.13), which is smaller than the pay-offs at the pure equilibria which are 2|1 and 1|2.

We consider p_i^+ and p_{-i}^+ as coordinates, ranging from 0 to 1, of course. By (1.5.13), noting $p_i^- = 1 - p_i^+$ and $p_{-i}^- = 1 - p_{-i}^+$, for each value of p_{-i}^+, we find either a single value of p_i^+ or the line $0 \le p_i^+ \le 1$ of those values that maximize the pay-off of i. The collection of these maximizing sets when p_{-i}^+ ranges from 0 to 1 is a connected set that connects the line $p_{-i}^+ = 0$ with the line $p_{-i}^+ = 1$. Similarly, the collection of the maximizing values of p_{-i}^+ when p_i^+ ranges from 0 to 1 connects the line $p_i^+ = 0$ with the line $p_i^+ = 0$. Therefore these two maximizing sets need to meet at a boundary point or intersect at least once. Since the intersections correspond to the Nash equilibria, this geometrically demonstrates the existence of a Nash equilibria.

The same kind of reasoning also works for QRE equilibria. For each value of p_{-i}^+, we find the optimal p_i^+, and therefore, when p_{-i}^+ varies, this collection of optimal values of p_i^+ connects the line $p_{-i}^+ = 0$ with the line $p_{-i}^+ = 1$, and analogously for $-i$, and the intersections of these curves (or their higher dimensional analogues in the general case) then yield the QREs. In fact, when the rationality parameter λ goes to 0, the limiting

response curves are the lines $p_i^+ = \frac{1}{2}$ and $p_{-i}^+ = \frac{1}{2}$, resp., which intersect at the common value $\frac{1}{2}$. From this, the existence of intersections for positive parameter values can then also be deduced from a homotopy argument. This is the basic idea of the proof of the Nash theorem as presented in this section.
According to (1.5.4), (1.5.5), the QRE is given by

$$p_i^+ = \frac{\exp(\lambda 2 p_{-i}^+)}{\exp(\lambda 2 p_{-i}^+)+\exp(\lambda(1-p_{-i}^+))} =: f_i(p_{-i}^+) \tag{1.5.16}$$

$$p_{-i}^+ = \frac{\exp(\lambda p_i^+)}{\exp(\lambda 2(1-p_i^+))+\exp(\lambda p_i^+)} =: f_{-i}(p_i^+). \tag{1.5.17}$$

The function f_i is the QRE analogue of the best response correspondance ρ_i. We note that both these functions, f_i and f_{-i}, are strictly increasing. We have the symmetry

$$f_i(p) + f_{-i}(1-p) = 1 \text{ for all } p. \tag{1.5.18}$$

Inserting one of the equations of (1.5.16) into the other then yields a fixed point equation for p_i^+ or p_{-i}^+. Equivalently, we look for the intersections of the two graphs resulting from (1.5.16), that is for p_i^+ as a function of p_{-i}^+ and for p_{-i}^+ as a function of p_i^+.
For large enough λ, the two graphs then intersect thrice. One intersection point is symmetric, that is, $p_{-i}^+ = 1 - p_i^+$ and unstable, whereas the two other ones are nonsymmetric and stable, but symmetric to each other; one of them is near $(1,1)$, the other near $(0,0)$ for large λ. In contrast, for small λ, we have the same situation as for $\lambda = 0$, that is, a single intersection. In fact, it is not hard to show that a pitchfork bifurcation creates three from one intersection points as λ increases.

2. We now consider the game with pay-off table (1.5.19)

1,1	2,1
1,2	0,0

.

Here, i prefers the first move, $+$, except when $-i$ plays $+$, in which case she is indifferent. Likewise, $-i$ prefers $+$, except if i plays $+$, in which case she is indifferent. Thus, since both of them prefer their first move, there is a tendency to end up at $(1,1)$ even though this is not Pareto optimal, in the sense that one player could get more without the other losing anything. This effect will now become clearer when we analyze the QREs.
For finite λ, analogously as in (1.5.16), (1.5.18), we put

$$p_i^+ = \frac{\exp(\lambda(2-p_{-i}^+))}{\exp(\lambda(2-p_{-i}^+))+\exp(\lambda p_{-i}^+)} =: f_i^\lambda(p_{-i}^+) \tag{1.5.20}$$

$$p_{-i}^+ = \frac{\exp(\lambda(2-p_i^+))}{\exp(\lambda(2-p_i^+))+\exp(\lambda p_i^+)} =: f_{-i}^\lambda(p_i^+). \tag{1.5.21}$$

Then $f_i(p)$ is a decreasing function with

$$f_i^\lambda(1) = \frac{1}{2}, \quad f_i^\lambda(0) > \frac{1}{2}, \tag{1.5.22}$$

and the same holds for f_{-i}^λ. In fact,

$$\lim_{\lambda\to\infty} f_i^\lambda(s) = 1 \text{ for } 0 \le s < 1. \tag{1.5.23}$$

Therefore, the intersection of their graphs, that is, the QRE, has to be contained in the upper quadrant, that is, it has to occur for

$$\frac{1}{2} < p_i^+ < 1, \quad \frac{1}{2} < p_{-i}^+ < 1. \tag{1.5.24}$$

This then also constrains the possible limits of QREs for $\lambda \to \infty$ to that region. In particular, those two Nash equilibria that are strict for at least one of the players, that is $+,-$ and $-,+$ are not limits of QREs. In particular, from (1.5.23), we see that the limit is $+,+$ which is not strict as either player could increase the other's pay-off while keeping her/his own. Of course, the reason why the equilibria $+,-$ or $-,+$ do not occur as limits of QREs is that at such a equilibrium, in contrast to $+,+$, the player playing $-$ bears the risk that the opponent might make a mistake and also play $-$ in place of $+$. Now, in the QRE situation, for finite λ, the players are not completely rational and play the wrong move $-$ with a certain positive probability. Thus, they do make mistakes with a certain nonvanishing probability, and therefore, players should guard themselves against mistakes of their opponents and play $+$ with a higher probability than $-$, and in fact with probability going to 1 as λ tends to infinity.
This example thus shows that not all Nash equilibria have to be limits of QREs.

3. We now wish to analyze the dependence on the value of the parameter λ more closely.

   We recall the matching pennies game (1.2.6)

1,-1	-1,1
-1,1	1,-1

. (1.5.25)

According to (1.5.4), (1.5.5), as in (1.5.16), (1.5.18), the QRE is given by

$$p_i^+ = \frac{\exp(\lambda p_{-i}^+ - \lambda(1-p_{-i}^+))}{\exp(\lambda p_{-i}^+ - \lambda(1-p_{-i}^+)) + \exp(\lambda(1-p_{-i}^+) - \lambda p_{-i}^+)} =: f_i(p_{-i}^+) \tag{1.5.26}$$

$$p_{-i}^+ = \frac{\exp(\lambda(1-p_i^+) - \lambda p_i^+)}{\exp(\lambda(1-p_i^+) - \lambda p_i^+) + \exp(\lambda p_i^+ - \lambda(1-p_i^+))} =: f_{-i}(p_i^+). \tag{1.5.27}$$

With

$$\phi_\lambda(p) := \exp(\lambda(2p-1)), \tag{1.5.28}$$

this can be written as

$$p_i^+ = \frac{\phi_\lambda^2(p_{-i}^+)}{\phi_\lambda^2(p_{-i}^+)+1} \tag{1.5.29}$$

$$p_{-i}^+ = \frac{1}{\phi_\lambda^2(p_i^+)+1}. \tag{1.5.30}$$

The unique QRE is then given by

$$p_i^+ = p_{-i}^+ = \frac{1}{2}. \tag{1.5.31}$$

This is the unique fixed point of the iteration of (1.5.26), (1.5.27). That is, we start with some value of p_{-i}^+, say, insert this value into (1.5.26) to get the corresponding $p_i^+ = f_i(p_{-i}^+)$, then in turn insert that value into (1.5.27) to compute the next iteration value of p_{-i}^+, and so on, then the only stationary pair of values of this iteration is given by (1.5.31). Equivalently, this is a fixed point of both $f_{-i} \circ f_i$ and $f_i \circ f_{-i}$.

In order to check for stability, we compute the derivatives of $f_i(p_{-i}^+)$ and $f_{-i}(p_i^+)$ at this fixed point; they are

$$\begin{aligned} f_i'(p_{-i}^+) &= \frac{2\lambda\phi_\lambda^2}{(\phi_\lambda^2+1)^2} &= \frac{\lambda}{2} \\ f_{-i}'(p_i^+) &= -\frac{2\lambda\phi_\lambda^2}{(\phi_\lambda^2+1)^2} &= -\frac{\lambda}{2} \quad \text{at } p_i^+ = p_{-i}^+ = \frac{1}{2}. \end{aligned}$$

These derivatives are smaller than 1 in absolute value for $\lambda < 2$, but larger for $\lambda > 2$. At $\lambda = 2$, they become 1 and -1, resp. Therefore (see [10] for an explanation of the basic bifurcation phenomena), at this value of λ, the fixed point loses its stability, and for $\lambda > 2$, we obtain a stable orbit of period 2 for the iteration. That is, exploiting also the symmetry between the two players, there are values $p^1 < \frac{1}{2}, p^2 > \frac{1}{2}$ such that

$$f_i(p^1) = p^1, \quad f_{-i}(p^1) = p^2, \quad f_i(p^2) = p^2, \quad f_{-i}(p^2) = p^1. \tag{1.5.32}$$

For $\lambda \to \infty$, p^1 tends to 0, p^2 to 1. This simply reflects the fact that when $-i$ chooses the first move, so should i, but when i chooses the first move, $-i$ should use the second one, whereas i should respond to the second move of $-i$ by her second move, but in turn i should respond to i's second move by his first move. For $\lambda < 2$, however, the fixed point is stable, and both players would happily converge to that fixed point under iteration. While the preceding analysis only establishes the local stability of the fixed point under small perturbations, in fact, in this example, it is not hard to check that it is globally stable, that is, the claim of the preceding sentence is justified.

4. When we change the game to the symmetric

1,1	-1,-1
-1,-1	1,1

, (1.5.33)

a coordination game like the battle of sexes, we have instead for the QRE

$$p_i^+ = \frac{\exp(\lambda p_{-i}^+ - \lambda(1-p_{-i}^+))}{\exp(\lambda p_{-i}^+ - \lambda(1-p_{-i}^+)) + \exp(\lambda(1-p_{-i}^+) - \lambda p_{-i}^+)} =: f_i(p_{-i}^+) \quad (1.5.34)$$

$$p_{-i}^+ = \frac{\exp(\lambda p_i^+ - \lambda(1-p_i^+))}{\exp(\lambda p_i^+ - \lambda(1-p_i^+)) + \exp(\lambda(1-p_i^+) - \lambda p_i^+)} =: f_{-i}(p_i^+), \quad (1.5.35)$$

that is, the two response functions are identical. Again

$$p_i^+ = p_{-i}^+ = \frac{1}{2} \quad (1.5.36)$$

is a fixed point. With (1.5.28), we then have

$$\begin{aligned} f_i'(p_{-i}^+) &= \tfrac{2\lambda\phi_\lambda^2}{(\phi_\lambda^2+1)^2} &= \frac{\lambda}{2} \\ f_{-i}'(p_i^+) &= \tfrac{2\lambda\phi_\lambda^2}{(\phi_\lambda^2+1)^2} &= \frac{\lambda}{2} \quad \text{at } p_i^+ = p_{-i}^+ = \frac{1}{2}. \end{aligned}$$

Now, for $\lambda = 2$, both derivatives become 1 at this fixed point. The fixed point again loses its stability, but this time through a pitchfork bifurcation. That means that for $\lambda > 2$, two further fixed points $p^1 < \frac{1}{2}, p^2 > \frac{1}{2}$ emerge which are both stable, while the fixed point at $\frac{1}{2}$ is unstable. Thus, for an iteration as in the preceding example, when the starting value is $< \frac{1}{2}$, an iteration as in the preceding example will converge to p^1, and when it is $> \frac{1}{2}$ to p^2. For $\lambda \to \infty$, these two stable fixed points converge to the two pure Nash equilibria where both players take the same move. This is, of course, the same phenomenon as in the first example.

We now turn to the game theoretical interpretation of the concept of a QRE and the mathematical findings of our examples.

We recall that game theory models agents that can choose between different options or moves by taking into account the effects of the moves of the other players in a fully rational manner. A Nash equilibrium (NE))in such a game is a selection of players' moves in such a manner that neither of them can increase their pay-offs by unilateral deviation from that choice of moves. In a pure NE, each player plays a single move, whereas in a mixed NE, players can play random mixtures of moves with fixed probabilities. As we have seen, each game possesses at least one NE, and generically only finitely many. Players are assumed fully rational in the sense that they can determine their best moves in mutual anticipation of their opponents' actions. This leads to the question of how to model the empirically observed deviations from full rationality of real humans in game theoretic situations. In fact, there are several different reasons why the behavior of humans in such situations may not be or may not seem or appear to be completely rational.
The concept of a quantal response equilibrium (QRE) of McKelvey and Palfrey[14] represents an attempt to address this issue within a formal model. Here, player

i selects her move probabilities according to (1.5.4) on the basis of the expected pay-offs given the move probabilities of her opponents. The distribution occurring in (1.5.4) is called a Gibbs distribution in statistical mechanics, and the parameter λ would correspond to an inverse temperature in the jargon of that field. (Of course, such an interpretation makes no sense in our context, but we nevertheless will use that terminology.)

Her opponents do the same, and when the resulting probabilities match, in the sense that her move probabilities induce precisely those move probabilities for her opponents that enter into the Gibbs distribution that determines hers, and this is true for all players, then they are at a QRE.

Of course, this is analogous to the concept of a NE which requires that a player's move leads to precisely those opponent moves against which it is a best response. Therefore, as we have formally shown, when the parameter λ of the Gibbs distributions that can be interpreted as expressing the degree of rationality of the players tends to infinity the situation approaches the one of the original rational game, and the QREs converge to NEs.
Both empirically and theoretically, there are different sources of non-rationality of players, and in fact, different interpretations of the deviations from rationality in such quantal response games have been offered in [14, 15]. Let us list some possible interpretations:

1. Players are irrational in the sense that they do not make optimal choices of their actions given their pay-offs, but with the resulting probability distribution of their moves depending on the pay-offs, in the sense that very bad mistakes (in terms of pay-offs) are less probable than milder ones. This has some basis in mathematical psychology, under the label "probabilistic choice". This is the original interpretation of [14].

2. The model is not about individual players, but about collections of players. Each individual player may well be fully rational, but fluctuations in the pay-offs lead to a distribution of moves. An external observer then has to accept the resulting probability distribution as the basis of his models. This is the situation of econometrics as treated, for instance, by McFadden [13].

3. The players face stochastic perturbations of their pay-offs (with noise of a certain type, assumed in standard econometrics and leading to the particular form of the quantal response functions discussed here). They know their own pay-off distributions exactly, but do not have the precise information about the pay-off distributions of the other players. This interpretation is proposed in [15].

There also arises the question about the meaning and the interpretation of the iterations discussed in the last examples for the stability of the fixed points. We could have

1. An iterated game. That is, the two players take turns in selecting their moves on the basis of the previous round of the game. The meaning of the probabilities involved then is not so clear, however; we could have

   (a) Subjective probabilities, that is, beliefs of each player about the probabilities employed by her opponent, or

   (b) Relative frequencies of moves from all the previous rounds. In that case, however, the iteration should proceed somewhat differently. Each player plays a concrete move according to his current probability distribution, and on this basis, his relative move frequencies are updated. This, however, would lead to a different iteration than that considered above. A similar type of stability analysis nevertheless applies.

2. A population game. That is, each index i represents a population of players, and for each instance of the game, one representative of each population is chosen for playing. The game is played in parallel with many such instances so that each individual has the opportunity to play. In each round, the players are newly assembled from their populations so that each individual usually plays against different opponents in different rounds. The probabilities involved then are the relative frequencies for the moves within each population. These frequencies will then change from round to round, and the adaptation rule (1.5.4) concerns these frequencies in the population.

3. A cognitive process. That is, player i thinks that the opponent will play with his probabilities p^1_{-i}, and she then selects her move probabilities $p^0_i = f_i(p^1_{-i})$ according to (1.5.4). Iteration then means that she assumes that the opponent had chosen his probabilities in turn via $p^1_{-i} = f_{-i}(p^2_i)$ for some putative probability p^2_i of hers which in turn would have been the response $f_i(p^3_{-i})$ to some putative probability p^3_{-i} of $-i$, and so on.

This section is based on joint work with Nils Bertschinger, Eckehard Olbrich and David Wolpert. The proof of the Nash theorem presented in this section is taken from the unpublished manuscript [3]. For applications of these ideas, see [23].

Chapter 2

Epistemic aspects

2.1 Cycles

We consider once more the matching-pennies game (1.2.6),

	b_1	b_2
a_1	1,-1	-1,1
a_2	-1,1	1,-1

. (2.1.1)

Here, when B knows that A plays a_1, he would play b_2, in which case, however, A would switch to a_2 which would induce B to play b_1 which would let A play a_1, and then the cycle would repeat itself. We represent this by the diagram

$$\begin{array}{ccc} 1,-1 & \longrightarrow & -1,1 \\ \uparrow & & \downarrow \\ -1,1 & \longleftarrow & 1,-1 \end{array} \quad , \qquad (2.1.2)$$

or in terms of strategy combinations,

$$\begin{array}{ccc} a_1, b_1 & \longrightarrow & a_1, b_2 \\ \uparrow & & \downarrow \\ a_2, b_1 & \longleftarrow & a_2, b_2 \end{array} \quad . \qquad (2.1.3)$$

These diagrams represent the reactions of the players to the actions of their opponents. We may, however, also carry out an epistemic analysis. Thus, A would play a_1 if she believed that B played b_1 which, assuming that B is rational, must be caused by B believing that A played a_2 which in turn must come from A's belief, at this stage of the belief iteration, that B played b_2 which in turn comes from B's belief that A played a_1, and the cycle repeats itself. That is, epistemically, the arrows should go in the opposite directions,

$$\begin{array}{ccc} a_1, b_1 & \longleftarrow & a_1, b_2 \\ \downarrow & & \uparrow \\ a_2, b_1 & \longrightarrow & a_2, b_2 \end{array} \quad . \tag{2.1.4}$$

2.2 Spaces of probability measures

We need to recall some mathematical concepts. While we shall try to explain the meaning of those concepts, we do not provide all the technical proofs, referring to [11] or [1] instead. The main result of interest for us is Theorem 2.2.1, and readers familiar with the mathematical concepts can directly turn there.

Definition 2.2.1. A topological space $(X, \mathcal{O}(X))$ is defined by a collection $\mathcal{O}(X)$ of subsets, called open, of some set X satisfying

(i) $\emptyset, X \in \mathcal{O}(X)$.

(ii) If $A, B \in \mathcal{O}(X)$, then also $A \cap B \in \mathcal{O}(X)$.

(iii) For any collection $(A_i)_{i \in I} \subset \mathcal{O}(X)$, also $\bigcup_{i \in I} A_i \in \mathcal{O}(X)$.

The complements of open sets are called closed.

We shall often simply write X instead of $(X, \mathcal{O}(X))$ for a topological space.

Of course, one could take for $\mathcal{O}(X)$ simply the collection of all subsets of X. However, in most cases of interest, the open sets should be distinguished by particular properties in order to make the concept useful. Likewise, the opposite case where one lets $\mathcal{O}(X)$ solely consist of X itself and $\emptyset$ is typically not of much interest or use. The next concept will provide us with better examples of topological spaces.

Definition 2.2.2. A metric space (X, d) is a set X with a function $d : X \times X \to \mathbb{R}^+$ satisfying for all $x, y, z \in X$

(i) $d(x, y) > 0$ iff $x \neq y$ (positive definiteness)

(ii) $d(x, y) = d(y, x)$ (symmetry)

(iii) $d(x, z) \leq d(x, y) + d(y, z)$ (triangle inequality).

We say that a sequence $(x_n)_{n \in \mathbb{N}}$ in a metric space (X, d) converges to the point $x \in X$, in symbols $x_n \to x$, if

$$d(x_n, x) \to 0 \text{ for } n \to \infty. \tag{2.2.5}$$

Such a point x is then called the limit of the sequence (x_n), $x = \lim_{n\to\infty} x_n$. (Observe that such a limit is necessarily unique.)
We say that the metric space (X, d) is complete if every Cauchy sequence (x_n) (that is, for every $\varepsilon > 0$, there exists some $N \in \mathbb{N}$ with $d(x_m, x_n) < \varepsilon$ whenever $m, n > N$) possesses a limit.
We say that the metric space (X, d) is compact if every sequence posseses a convergent subsequence.
(X, d) is separable if it carries a dense subsequence x_n, that is, for every $x \in X, \varepsilon > 0$, there exists some $n \in \mathbb{N}$ with $d(x, x_n) < \varepsilon$.

We shall often simply write X instead of (X, d) for a metric space.
A metric d defines a topology on a set X by calling the *distance balls*

$$U(x, r) := \{y \in X : d(x, y) < r\} \tag{2.2.6}$$

for all $x \in X, r \geq 0$ open and letting $\mathcal{O}(X)$ consisting of all finite intersections and arbitrary unions of such balls. (We also say that the open balls generate the topology.)
Subsequently, we shall also need the closed balls

$$B(x, r) := \{y \in X : d(x, y) \leq r\}. \tag{2.2.7}$$

Of course, one needs to verify that the sets $B(x, r)$ are closed, indeed, but this is not too hard. We observe that a set consisting of a single point x is closed, as it is given by the ball $B(x, 0)$.

Definition 2.2.3. A topological space $(X, \mathcal{O}(X))$ is called metrizable if there exists a metric on X that generates the topology, that is, the open balls of d generate the collection $\mathcal{O}(X)$ of open sets. It is called completely metrizable if its topology can be generated by a complete metric.
A Polish shace is a completely metrizable, separable, topological space.

For instance, on the real line $\mathbb{R}$, we have the Euclidean metric $d(x, y) = |x - y|$, and this generates its standard topology.
We note that a metric generating a given topology in general is not unique. Nevertheless, when we shall work with a metrizable space, we shall usually implicitly choose some metric generating its topology, so long as our constructions and results will not depend on that particular choice. Put differently, the reason is that metric spaces enjoy stronger properties than general topological spaces, and therefore, when a space is metrizable, we can appeal to those properties of metric spaces.

When the metrizable topological space X is separable, for instance Polish, with a dense sequence (x_n), then the balls $U(x_n, r)$ with *rational* r already generate the topology. Such a space whose topology is generated by a countable family of open sets is called second countable.

Definition 2.2.4. A function $f : X \to \mathbb{R}$ on a topological space $(X, \mathcal{O}(X))$ is called continuous if the preimage of any open subset of $\mathbb{R}$ is an open subset of X.

It is easy to check that a function $f : X \to \mathbb{R}$ on a metric space (X, d) is continuous iff for every $x \in X, \varepsilon > 0$ there exists some $\delta > 0$ with $|f(x) - f(y)| < \varepsilon$ whenever $d(x, y) < \delta$.

Definition 2.2.5. We denote by $C_b(X)$ the space of bounded continuous function on the topological space $(X, \mathcal{O}(X))$. We provide this space with the norm

$$\|f\|_\infty := \sup_{x \in X} |f(x)|. \tag{2.2.8}$$

The sequence $(f_n)_{n\in\mathbb{N}} \subset C_b(X)$ converges uniformly to $f \in C_b(X)$ iff

$$\lim_{n\to\infty} \|f_n - f\|_\infty = 0. \tag{2.2.9}$$

There also exists amother, weaker, notion of convergence. (f_n) converges pointwise to the function f iff

$$\lim_{n\to\infty} f_n(x) = f(x) \text{ for all } x \in X. \tag{2.2.10}$$

Whereas the uniform limit of a sequence of continuous functions is continuous again, this need not be so for pointwise limits. In fact, for instance, we may consider characteristic functions of sets B,

$$\chi_B(y) := \begin{cases} 1 \text{ if } y \in B \\ 0 \text{ if } y \notin B. \end{cases} \tag{2.2.11}$$

Such functions are usually not continuous, as they jump from 0 to 1 at the boundary of B. The characteristic function of a closed ball $\chi_{B(x,r)}$ then is the decreasing (pointwise) limit of a sequence of continuous functions,

$$f_n(y) := \begin{cases} 1 \text{ if } y \in B(x, r) \\ 1 - nd(y, B(x, r)) \text{ if } d(y, B(x, r)) \leq \frac{1}{n} \\ 0 \text{ else.} \end{cases} \tag{2.2.12}$$

Similarly, the characteristic function of an open ball is the increasing limit of a sequence of continuous functions. This offers the advantage that it often suffices to check certain properties for continuous functions and then pass to a pointwise limit to extend them to characteristic functions of open or closed balls, and then more generally to open or closed sets in a metric space, as those sets are generated by the balls. And when we can handle characteristic functions, we can also handle step functions, that is, functions, that assume only finitely many values, each of them on some measurable subset of our space, as defined in the next definition. This is the basis of Lebesgue integration theory.

In particular, we can use a construction as in (2.2.12) to show that for any two points x_1, x_2 in a metric space X, there exists a continuous function with $f(x_1) = 1, f(x_2) = 0$. From this, we can derive

Lemma 2.2.1. *In a metric space X, a sequence (x_n) converges to a point x iff*

$$f(x_n) \to f(x) \text{ for all continuous functions } f. \tag{2.2.13}$$

Definition 2.2.6. A measurable space $(X, \mathcal{B})$ is a set X equipped with a σ-algebra $\mathcal{B}$, a set of subsets of X satisfying:

(i) $X \in \mathcal{B}$.

(ii) If $B \in \mathcal{B}$, then so is $X \backslash B$.

(iii) If $B_n \in \mathcal{B}$ for all $n \in \mathbb{N}$, then so is $\bigcup\limits_{n \in \mathbb{N}} B_n$.

The members of $\mathcal{B}$ then are called measurable sets.

The preceding properties imply:

(iv) $\emptyset \in \mathcal{B}$.

(v) If $B_1, ..., B_m \in \mathcal{B}$, then so is $\bigcap\limits_{j=1}^{m} B_j$.

In order to obtain a σ-algebra, we can start with any collection of subsets of X and close it up under complements and countable unions. For a topological space $(X, \mathcal{O}(X))$, the collection of open sets in general is not closed under complements, but we can thus take the smallest σ-algebra containing all open subsets of X. The sets in this σ-algebra are called Borel sets. Since the closed sets are the complements of the open ones, this Borel sigma algebra then contains both the open and the closed sets. And since this Borel sigma algebra is generated by either the open or the closed sets, it suffices to check properties on one of those subclasses of sets, for instance when we shall consider probability measures below.
So far, however, one might be naturally inclined to simply use the entire Boolean algebra of all subsets of X as a σ-algebra. The reason why one works with smaller σ-algebras than that one is that when the σ-algebra is too large, it becomes too restrictive to satisfy the properties in the next definition, the really important one.

Definition 2.2.7. A probability measure on $(X, \mathcal{B})$ is a function

$$\mu : \mathcal{B} \to [0, 1]$$

satisfying:

(i) $\mu(\bigcup\limits_{n \in \mathbb{N}} B_n) = \sum\limits_{n \in \mathbb{N}} \mu(B_n)$, if $B_i \cap B_j = \emptyset$ for all $i \neq j$, i.e., if the sets B_n are pairwise disjoint,

(ii) $\mu(X) = 1$.

A triple $(X, \mathcal{B}, \mu)$ with the preceding properties is called a probability measure space. When $\mathcal{B}$ is the Borel σ-algebra, we speak of a Borel probability measure.

One important point is that the additivity in (i) is required only for *countable* families whereas in (iii) of Def. 2.2.1, arbitrary families were considered. This may lead to certain technical difficulties. These difficulties, however, are avoided for second countable spaces, that is, spaces with a countable basis of their topology, like separable metric spaces as observed above, after Def. 2.2.3. For instance,

Lemma 2.2.2. *Any measure μ on a second countable topological space X possesses a unique support, that is, a closed set $\operatorname{supp}\mu \subset X$ such that*

(i) $\mu(X\backslash \operatorname{supp}\mu) = 0$,

(ii) $\mu(U) > 0$ *whenever U is open with $U \cap \operatorname{supp}\mu \neq \emptyset$.*

Proof. Let

$$V := \bigcup \{U \text{ open}, \ \mu(U) = 0\}, \tag{2.2.14}$$

which, since X is second countable, is a countable union of open sets of measure 0, hence by (i) of Def. 2.2.7,

$$\mu(V) = 0. \tag{2.2.15}$$

We put $M := X\backslash V$. Then M is closed, and for any open U with

$$U \cap M \neq \emptyset, \tag{2.2.16}$$

we have $\mu(U) > 0$, as otherwise by (2.2.14), $U \subset V$, contradicting (2.2.15), (2.2.16). Thus, $M = \operatorname{supp}\mu$. (In fact, we have $\mu(U \cap \operatorname{supp}\mu) > 0$, as $U = (U \cap \operatorname{supp}\mu) \cup (U \cap V)$, and the second component has measure 0.)

□

More generally, we could consider (Borel) measures μ, that is, $\mu : \mathcal{B} \to \mathbb{R}^+ \cup \infty$, dropping the requirement $\mu(X) = 1$. On a (probability) measure space, we can develop an integration theory. First, one defines the integral of a step function χ (i.e., for a collection of disjoint measurable sets $B_i, i = 1, \dots, n$ with $X = \bigcup_i B_i$, $\chi \equiv c_i$ is constant on each B_i), one defines the integral in the obvious way,

$$\int \chi d\mu := \sum_i c_i \mu(B_i). \tag{2.2.17}$$

One then looks at the class of functions f that can be approximated by a sequence of step functions in such a way that the corresponding integrals converge to a value that then is defined as the integral $\int f d\mu$. We omit the details, referring to, e.g., [11]. In particular, for a Borel probability measure on a metric space, all continuous functions are integrable.

As in [1], we now consider a metrizable topological space X and the set $\mathcal{P}(X)$ of all Borel probability measures on it.

Definition 2.2.8. A sequence $(\mu_n)_{n\in\mathbb{N}}$ of Borel probability measures on X converges in the weak*-sense to $\mu \in \mathcal{P}(X)$, in symbols

$$\mu_n \overset{\star}{\rightharpoonup} \mu, \text{ or simply } \mu_n \rightharpoonup \mu,$$

iff for every $f \in C_b(X)$,

$$\int f d\mu_n \to \int f d\mu \text{ for } n \to \infty. \tag{2.2.18}$$

We have

Lemma 2.2.3.

$$\mu_n \overset{\star}{\rightharpoonup} \mu \tag{2.2.19}$$

iff

$$\limsup_n \mu_n(A) \le \mu(A) \textit{ for every closed } A, \tag{2.2.20}$$

or equivalently,

$$\liminf_n \mu_n(U) \ge \mu(U) \textit{ for every open } U. \tag{2.2.21}$$

We sketch the

Proof. As usual, for verifying (2.2.20) or (2.2.21) – which are clearly equivalent, it suffices to consider distance balls. As we have seen above, the characteristic function χ_A of a closed distance ball is a decreasing limit of continuous functions f_ν. Apply on one hand (2.2.19) to each f_ν to get

$$\int f_\nu d\mu = \lim_{n\to\infty} \int f_\nu d\mu_n \ge \limsup \mu_n(A)$$

and use on the other hand Lebesgue's dominated convergence theorem to get

$$\mu(A) = \int \chi_A d\mu = \lim_{\nu\to\infty} \int f_\nu d\mu,$$

hence (2.2.20).

For the other direction, approximate a continuous function by step functions of closed sets to obtain $\int f d\mu \ge \limsup_n \int f d\mu_n$, and apply this then also to $-f$ to get equality in (2.2.19). □

Corollary 2.2.1. *If, for a sequence of Borel probability measures μ_n,*

$$\mu_n \overset{\star}{\rightharpoonup} \mu, \tag{2.2.22}$$

then μ is again a Borel probability measure. In other words, $\mathcal{P}(X)$ is closed under weak *- convergence.

Proof. From (2.2.20), (2.2.21), it follows that $\mu(X) = 1$ as the μ_n satisfy this property. Thus, μ is a probability measure. □

There is an obvious class of probability measures, the Dirac measures $\delta_x, x \in X$. They are also called point masses. Such a measure is characterized by the property

$$\delta_x(A) = \begin{cases} 1 & \text{if } x \in A \\ 0 & \text{if } x \notin A \end{cases} \tag{2.2.23}$$

or equivalently,

$$\delta_x(f) = \int f d\delta_x = f(x) \text{ for all } f \in C_b(X). \tag{2.2.24}$$

Therefore,

$$\delta_{x_n} \rightharpoonup \delta_x \tag{2.2.25}$$

iff

$$f(x_n) \to f(x) \text{ for all } f \in C_b(X). \tag{2.2.26}$$

Now, for a metric space, the latter is equivalent to

$$x_n \to x, \tag{2.2.27}$$

as one readily checks, for instance by considering continuous functions $f_{\nu,x}, \nu \in \mathbb{N}$, of the form

$$f_{\nu,x}(y) = \begin{cases} 1 - \nu d(x,y) & \text{if } \nu d(x,y) \le 1 \\ 0 & \text{else.} \end{cases} \tag{2.2.28}$$

Therefore,

$$\begin{array}{ccc} X & \to & \mathcal{P}(X) \\ x & \mapsto & \delta_x \end{array} \tag{2.2.29}$$

identifies X with a topological subspace of $\mathcal{P}(X)$.

Lemma 2.2.4. *When X is separable, X is closed as a topological subspace of $\mathcal{P}(X)$.*

Proof. Let

$$\delta_{x_n} \rightharpoonup \mu \in \mathcal{P}(X). \tag{2.2.30}$$

We need to show that $\mu = \delta_x$ for some $x \in X$. By Lemma 2.2.2, μ has a support $\operatorname{supp} \mu$, and since $\mu(X) = 1$, this support is nonempty. By utilizing the above functions $f_{\nu,x}$ (2.2.28), however, it is easy to see that when $x \in \operatorname{supp} \mu$, then $x_n \to x$. Hence, as we have observed that this is equivalent to (2.2.25), $\mu = \delta_x$. □

We wish to utilize this embedding of X into $\mathcal{P}(X)$ in order to better understand the structure of the latter space. We first observe

Lemma 2.2.5. *$\mathcal{P}(X)$ is convex, and the $\delta_x, x \in X$ are extreme points of this convex space, that is, they cannot be represented as nontrivial convex combinations of other probability measures.*

Proof. Obviously, when $\mu, \nu \in \mathcal{P}(X)$ and $0 \le t \le 1$, then also $t\mu + (1-t)\nu \in \mathcal{P}(X)$. Thus, $\mathcal{P}(X)$ is convex. Also, it is clear that the δ_x are extreme points. □

Lemma 2.2.6. *When X is a separable metrizable topological space, then conversely every extreme point of $\mathcal{P}(X)$ is of the form δ_x for some $x \in X$.*

We sketch the

Proof. Let μ be extreme in $\mathcal{P}(X)$. Again, we use Lemma 2.2.2, that is, the existence of $\operatorname{supp}\mu$. When, argueing by contradiction, $\operatorname{supp}\mu$ should contain more than one point, say x and y, then one represents μ as a convex combination of two measures one of which contains x in its support, and the other y. □

We may then also consider the convex hull of X in $\mathcal{P}(X)$, that is, the set of all finite convex combinations of Dirac measures δ_x. This is nothing but the set of all probability measures with finite support, as any such measure can be written as a convex combination of Dirac measures. With the help of the Hahn-Banach Theorem, one may show that this convex hull of the Dirac measures is dense in $\mathcal{P}(X)$. Thus, curiously, while we have seen in Lemma 2.2.4 that X is closed in $\mathcal{P}(X)$, its convex hull is not closed, unless X is finite, because there exist probability measure that do not have finite support. $\mathcal{P}(X)$ itself is closed again under weak * - convergence by Corollary 2.2.1.

We can now formulate the result that is of main interest for our purposes.

Theorem 2.2.1. *The metrizable topological space X is compact iff $\mathcal{P}(X)$ is compact and metrizable. Also, X is a Polish space, i.e., completely metrizable and separable, iff $\mathcal{P}(X)$ is.*

That is, the set of Borel probability measures on a compact X is again compact. Therefore, we can iterate the construction and consider $\mathcal{P}(\mathcal{P}(X))$ etc and obtain a family of compact metrizable spaces.

In some sense, this is a surprising result because $\mathcal{P}(X)$ is a much larger space than X. For instance, when X is finite, for instance $X = \{x_1, x_2\}$, then X is no longer finite, but a finite dimensional space, $\mathcal{P}(X) = \{t\delta_{x_1} + (1-t)\delta_{x_2}, 0 \le t \le 1\} = [0,1]$ in this example. $\mathcal{P}(\mathcal{P}(X))$ then already is an infinite dimensional space. The reason why these iterated spaces will nevertheless all stay compact is that weak * - convergence is, as the name indicates, a rather weak form of convergence. That is, for a sequence it is quite easy to converge weakly; essentially, this convergence behaves as in the finite dimensional situation even if $\mathcal{P}(X)$ is infinite dimensional.

We cannot provide a complete *proof* of Theorem 2.2.1 here. We only describe some of the main points and refer to [1] for details. It might be helpful to identify the key steps.

1. By Lemma 2.2.1, $\mathcal{P}(X)$ is weak * - closed. Thus, when below, we construct weak * - limits of probability measures, these limits will be automatically contained in $\mathcal{P}(X)$.

2. For the compactness of $\mathcal{P}(X)$, the key result is the Theorem of Banach-Alaoglu. In order to understand its formulation, we need to introduce the norm
$$\|\mu\| := \sup\{\mu(f) : f \in C_b(X), \|f\|_\infty \le 1\}. \tag{2.2.31}$$
For every $\mu \in \mathcal{P}(X)$, we then have $\|\mu\| = \mu(1) = 1$, where 1 is the function that is identically 1 on all of X.

   Theorem 2.2.2. *Any sequence ν_n of (not necessarily probability) measures with uniformly bounded norms, i.e., $\|\nu_n\| \le K$ for some constant K and all n contains a weak * - convergent subsequence. In other words, any weak * - closed and norm-bounded set of measures is weak * - compact.*

   This result by itself does not refer to any particular structure of X. What it needs is only that $C_b(X)$ is a vector space with a norm.

3. When X is compact, the situation simplifies, however. In that case, all continuous functions are automatically bounded, and we can work simply with the space $C(X)$ of continuous functions on X.

4. Also, when X is in addition metrizable, then $C(X)$ is separable. This follows from the Stone-Weierstrass Theorem. The latter result says that an algebra of real-valued continuous functions on a compact topological space X is dense in $C(X)$ if it contains the constant function 1 and separates the points of X, that is, for any two distinct points $x, y \in X$, we can find a continuous function f with $f(x) \neq f(y)$. When X is metrizable and separable, we can take a dense sequence $(x_m)_{m\in\mathbb{N}}$ and the functions $f_0(x) = 1$ and $f_m(x) := d(x, x_m)$ for $m \in \mathbb{N}$. These functions then separate the points, and we can apply the Stone-Weierstrass theorem to the algebra generated by finite products of such functions. This then implies, in particular, that $C(X)$ is separable.

5. When $C(X)$ is separable, Theorem 2.2.2 is due to Banach and can be proved by a diagonal sequence argument. That is, we take a dense sequence $(g_k)_{k\in\mathbb{N}}$ of continuous functions. We then first find a subsequence $(\nu_{1,n})_{n\in\mathbb{N}}$ of ν_n for which the sequence $\nu_{1,n}(g_1)$ of bounded real numbers converges. Having iteratively found a sequence $\nu_{k,n}$ for which $\nu_{k,n}(g_j)$ converges for $j = 1, \dots, k$ as $n \to \infty$, we find a subsequence $\nu_{k+1,n}$ of $\nu_{k,n}$ for which in addition $\nu_{k+1,n}(g_{k+1})$ converges. For the diagonal sequence $\nu_{n,n}$, then $\nu_{n,n}(g_k)$ converges for every k as $n \to \infty$. Since g_k is dense in $C(X)$, therefore $\nu_{n,n}(g)$ converges for every $g \in C(X)$, and we obtain a weak * - limit ν. This sketches the proof of Theorem 2.2.2 in the separable case.

6. Next, when $C(X)$ is separable, we take a dense sequence (g_k) as before and define a metric on $\mathcal{P}(X)$ via

$$d(\mu,\nu) := \sum_k \frac{1}{2^k}|\mu(g_k) - \nu(g_k)|. \tag{2.2.32}$$

   One then checks that the identity map is continuous as a map from $\mathcal{P}(X)$ with its weak * - topology to the metric space $(\mathcal{P}(X), d)$. Since $\mathcal{P}(X)$ is weak * - compact by Theorem 2.2.2, the identity then has to be a homeomorphism. Therefore, the metric d induces the weak * - topology. Thus, $\mathcal{P}(X)$ is metrizable. (The metric (2.2.32) itself is not very useful; it only serves the auxiliary role of showing the metrizability of $\mathcal{P}(X)$.)

7. Summarizing the preceding steps, when X is compact and metrizable, then $C(X)$ is separable, and then $\mathcal{P}(X)$ with its weak * - topology is compact and metrizable.

8. In turn, these conditions on X are also necessary. Indeed, when $\mathcal{P}(X)$ is compact and metrizable, hence also separable, then first, by the embedding (2.2.29), X is a topological subspace of $\mathcal{P}(X)$, hence in particular separable itself. By Lemma 2.2.4, it is then closed in $\mathcal{P}(X)$, hence also compact and metrizable because the latter is.

9. Thus, we have handled the first part of Theorem 2.2.1. This is the important part for our applications in game theory. For the second part, which is more technical, we refer to [1].

2.3 Some examples

We start with some examples inspired by [17]. We have two companies, A and B (where these might be acronyms for their CEOs, Alice and Bob) that are contemplating to invest into one of three possible markets, m_1, m_2 or m_3. In Game 1, when only one company invests into market m_i, it will get a return of i. When both of them invest into the same market, however, either of them will get a return of 0, as the market is too small to accommodate two competitors. In Game 2, the pay-off of a single investor is the same as in Game 1, but when both of them go into the same market, they both face a loss of -2. Finally, in Game 3, A is an innovator, but B is a copier. That is, A develops a new product, and with that product, it can develop a new market, m_1, m_2 or m_3, but its production costs are relatively high. B, in contrast, does not invent, but specializing in the cheap production of products it copies from A. Thus, in Game 3, for A the pay-offs are as in Game 1, but for B, the pay-off is 4 when it can go into the same market as A, because it can then benefit from the marketing efforts of A in that market.

In tabular form, the pay-offs for both A and B in Game 1 are

m_3	m_2	m_1	both in same market
3	2	1	0

. (2.3.1)

In Game 2, the pay-offs for both A and B are

m_3	m_2	m_1	both in same market
3	2	1	-2

. (2.3.2)

In Game 3, the pay-offs for A are as in Game 1, i.e.,

m_3	m_2	m_1	both in same market
3	2	1	0

, (2.3.3)

whereas those for B are

m_3	m_2	m_1	both in same market
3	2	1	4

. (2.3.4)

In each of these games, neither company knows what the other will do, but it will try to anticipate the action of the competitor so as to avoid being in the same market, except for B in Game 3 which in contrast will try to copy A.

So, let's try to figure out how the players might think and act in a rational manner. Obviously, in Game 1, when either company believes that the competitor will go into m_3 or m_2, it should go into the opposite market. Moreover, in this game, for neither player it will be rational to go into m_1. Here "not rational" is informally understood in the sense that there will be a strategy with a higher expected pay-off. Indeed, let us assume, for instance, that A thinsk that B will go into m_3 with probability p_3 and into m_2 with probability p_2, with $p_3 + p_2 \leq 1$. Then A's expected pay-off from going into m_1 is ≤ 1 ($= 1$ if $p_3 + p_2 = 1$), whereas the expected pay-off for m_3 or m_2 is

$$0p_3 + 3(1 - p_3) \text{ or } 0p_2 + 2(1 - p_2), \tag{2.3.5}$$

and since $p_3 + p_2 \leq 1$, at least one of these numbers is > 1. Thus, A always has a better strategy than m_1. This is no longer so for Game 2. Here, the corresponding pay-offs would be

$$-2p_3 + 3(1 - p_3) = 3 - 5p_3 \text{ or } -2p_2 + 2(1 - p_2) = 2 - 4p_2, \tag{2.3.6}$$

and for $p_3 > 2/5$, the former is < 1, while for $p_2 > 1/4$, the latter becomes < 1. Thus, if A for instance believes that with probability .7, B will go into m_3 and with probability .3 into m_2, A should choose m_1.

However, so far, this analysis is clearly incomplete because A does not yet take into account what B believes. Let us return to Game 1. We assign to A two possible *types*, t_A^1, t_A^2, and likewise t_B^1, t_B^2 to B. t_A^1 (t_A^2) means that A believes that B will play m_3 (m_2) and is of type t_B^2 (t_B^1). Analogously, t_B^1 (t_B^2) means that B believes that A will play m_3 (m_2) and is of type t_A^2 (t_A^1). Thus, the type t_A^1 means that A believes that B, being of type t_B^2, will play m_2 and in turn believes that A will play m_3 and that in turn that A, being of type t_A^1 according to B's belief at this stage, believes that B will play m_2, and so on. And all this is completely rational in the sense that at each stage, either player plays the best response against the putative action of the competitor.

In Game 3, these types would no longer correspond to rational behavior. Here, we might introduce rational types as follows. τ_A^1 (τ_A^2) means that A believes that B will play m_3 (m_2) and is of type τ_B^1 (τ_B^2). And τ_B^1 (τ_B^2) means that B believes that A will play m_3 (m_2) and is of type τ_A^2 (τ_A^1). Thus, the type τ_A^1 means that A believes that B, being of type τ_B^1, will play m_3 and in turn believes that A will play m_3 and that in turn that A, being of type t_A^2 according to B's belief at this stage, believes that B will play m_2, and so on. And again, all this is completely rational in the sense that at each stage, either player plays the best response against the putative action of the competitor.

For Game 2, we keep the same types as in Game 1, but add the type t_A^3, incorporating A's belief that with probability .7, B will play m_3 and is of type t_B^2, whereas with probability .3, B will choose m_2 and is of type t_B^1. Thus, with probability .7, A will also think that B thinks that A will play m_2, and we then get into the same cycle as before. We can also represent this in a table for the types of A

	m_3, t_B^2	m_2, t_B^1
t_A^1	1	0
t_A^2	0	1
t_A^3	.7	.3

, (2.3.7)

whereas for B, we have

	m_3, t_A^2	m_2, t_A^1
t_B^1	1	0
t_B^2	0	1

. (2.3.8)

By alternating between these two tables, we can then reconstruct the iterated beliefs of the two players.

Of course, according to this scheme, we can introduce further types for A and B and check their rationality. For instance, we could introduce the type t_B^4 that believes that A will play m_3 and is of type t_A^1. In that case, B would believe that A is not rational as A should not play m_3 when believing that B would also play m_3.

2.4 Types (Belief hierarchies)

In this section, we present the theory of types and belief hierarchies, as developed by Brandenburger-Dekel [4] and others in the framework proposed by Harsanyi [8]. Other sources of the theory that have been useful for our presentation can be found in [18, 17, 12], for instance.

We shall work with two players, i and $-i$. i will be referred to by the pronoun "she", $-i$ by "he", even though the situation is entirely symmetric between them.

The players form beliefs about each other, but this then needs to iterated, that is, each player then also has to form beliefs about the beliefs of the opponent, and so on. Thus, there are two aspects, the possible contents of those beliefs, and the iterative aspect. For the content, we introduce the uncertainty domains S_i and S_{-i}; S_{-i} contains what i is uncertain about regarding $-i$. One may assume, as is frequently done, that $S_i = S_{-i} =: S$, i.e., that there is a common domain of uncertainty, even though the two players may be uncertain about different aspects, parts, or elements of S. The iterative part will be captured by the notion of a type. Putting it succinctly and somewhat paradoxically, a type will be a probability distribution over types of the opponent. And to add the content, it will also involve a probability distribution over the opponent's domain of uncertainty. In order to make the iterative construction possible, we assume either that all spaces involved, the S_i already introduced and the T_i to come, are compact metrizable, or that they are Polish spaces. By Theorem 2.2.1, then, all spaces of probability distributions to be iteratively introduced will then be of the same class.

Definition 2.4.1. For i, we assume to have a type space T_i together with a map

$$p_i : T_i \to \mathcal{P}(S_{-i} \times T_{-i}). \tag{2.4.1}$$

A type $t_i \in T_i$ then associates to i a probability distribution $p_i(\cdot, \cdot|t_i)$ (often written as $p_i(t_i)$ in shorthand notation) over the domain of uncertainty and the type space of the opponent.

The tupels (T_i, p_i) will be restricted in the sequel to satisfy certain consistency and rationality conditions.

More precisely, for every t_i, we have a probability distribution

$$\begin{aligned} p_i(\cdot, \cdot|t_i) &\in \mathcal{P}(S_{-i} \times T_{-i}) \\ (s_{-i}, t_{-i}) &\mapsto p_i(s_{-i}, t_{-i}|t_i) \text{ for every } s_{-i} \in S_{-i}, t_{-i} \in T_{-i}. \end{aligned} \tag{2.4.2}$$

I hope that the shorthand notation employed in the Definition and subsequently will not give rise to problems.

p_i as a probability distribution on $S_{-i} \times T_{-i}$ by marginalization gives rise to a probability distribution on S_{-i}

$$q_i := p_{i,S_{-i}} \tag{2.4.3}$$

by integrating out the variable $t_{-i} \in T_{-i}$ for every s_{-i}, t_i. (Formally, we have $q_i(s_{-i}|t_i) = \int p_i(s_{-i}, t_{-i}|t_i)dt_{-i}$, that is, we integrate out t_{-i} w.r.t. the probability measure $p_i(s_{-i}, \cdot|t_i)$.) $q_i(\cdot|t_i)$ represents that belief that type t_i entertains about the distributions of the values in the uncertainty domain S_{-i}. This is the first order belief of t_i.

Similarly, we can marginalize on the type space T_{-i} to obtain

$$\pi_i := p_{i,T_{-i}}. \tag{2.4.4}$$

We can now iterate the types and construct belief hierarchies. In order to obtain i's second order beliefs, that is, what the type t_i of i believes that $-i$ believes about i's actions, we have

$$\int q_{-i}(s_i|\tau_{-i})\pi_i(\tau_{-i}|t_i))d\tau_{-i}. \tag{2.4.5}$$

In words: t_i entertains a belief $\pi_i(\cdot|t_i)$ about the distribution of types τ_{-i} of $-i$, and each such type τ_{-i} in turn has a distribution $q_{-i}(\cdot|\tau_{-i})$ over the possible states s_i of i. By integrating out the types τ_{-i}, we then get the second order belief distribution of i about s_i.

Similarly, we obtain the third order belief distribution, what type t_i believes that $-i$ believes that i believes about the actions of $-i$, as

$$\int\int q_i(s_{-i}|\tau_i)\pi_{-i}(\tau_i|\tau_{-i})\pi_i(\tau_{-i}|t_i))d\tau_i d\tau_{-i}. \tag{2.4.6}$$

Thus, we are also integrating out the types τ_i of i for which each type τ_{-i} has a distribution $\pi_{-i}(\cdot|\tau_{-i})$.

It turns out that the structure will become clearer when we go to a higher level of abstraction and generality. More precisely, we shall consider the space of all possible types and also that of all possible belief hierarchies. Here, the belief hierarchy induced by a type t_i should encode all the beliefs of t_i of all orders k simultaneously. We shall follow the presentation in [12]. In particular, we shall assume that $S_i = S_{-i} =: S$. Since we might have replaced S_i and S_{-i} by their product $S_i \times S_{-i}$, this does not restrict the generality of our considerations.

Definition 2.4.2. The space of belief hierarchies of player i of order k is inductively defined via

$$\begin{aligned}
X^1_{-i} &:= \mathcal{P}(S) \text{ and for } k \geq 2 && (2.4.7)\\
X^k_{-i} &:= X^{k-1}_{-i} \times \mathcal{P}(S \times X^{k-1}_i) && (2.4.8)\\
&= X^1_{-i} \times \mathcal{P}(S \times X^1_i) \times \mathcal{P}(S \times X^2_i) \times \dots \mathcal{P}(S \times X^{k-1}_i) &&\\
&= \mathcal{P}(S) \times \mathcal{P}(S \times X^1_i) \times \dots \mathcal{P}(S \times X^{k-1}_i) \text{ by (2.4.7)} && (2.4.9)
\end{aligned}$$

X^k_{-i} is the space of kth order beliefs of player i about both the states in S and the opponent $-i$.

Note that X_i^{k-1}, the space of $(k-1)$st order beliefs of player $-i$ contains as a factor X_{-i}^{k-2}, the space of $(k-2)$nd order beliefs of player i, and so on. This means that player i entertains first a belief about the states in S, then about $S \times X_i^1$, that is, what $-i$ believes about the states and the first-order beliefs of i, then about the states and the beliefs of $-i$ concerning what i believes about the states and the beliefs of $-i$, and so on. Or more plainly, you ask first about the probabilities of the states, then what she thinks about those probabilities, then what she thinks what I think about those probabilities, and so on.
Importantly, an element of $\mathcal{P}(S \times X_i^k)$ need not be a product distribution. For instance, i may believe that $-i$ knows the true state in S. As a concrete example, $S = \{0, 1\}$ might just consist of two possible states, and i might assign probability .4 to the combination of 0 and the belief of $-i$ that i believes in either possibility with probability .5, and i might assign probability .6 to the possibility that the state is 1 and that $-i$ believes that i thinks that the state is 0.
In fact, since we have convinced ourselves that we may assume that $S_i = S_{-i} = S$, the situation for the two players is symmetric, and we may drop the indices i and $-i$ of the X^k. Thus, we now simply consider

$$X^1 := \mathcal{P}(S) \text{ and for } k \geq 2 \tag{2.4.10}$$
$$X^k := X^{k-1} \times \mathcal{P}(S \times X^{k-1}) \tag{2.4.11}$$
$$= X^1 \times \mathcal{P}(S \times X^1) \times \mathcal{P}(S \times X^2) \times \ldots \mathcal{P}(S \times X^{k-1})$$
$$= \mathcal{P}(S) \times \mathcal{P}(S \times X^1) \times \ldots \mathcal{P}(S \times X^{k-1}) \text{ by (2.4.10)} \tag{2.4.12}$$

as the common belief spaces of the two players. Of course, the beliefs of the two players will in general be different from each other, that is, the two players will choose different elements from these belief spaces.

With this simplification, it is also natural to formulate belief sequences over external events. Let us consider an example. We have two states, 0 meaning that the stock will crash within the next month, 1 that it won't. Let i believe with probability .8 that it will crash and that $-i$ believes with probability .5 that it will, and believe with probability .2 that it won't, but that $-i$ believes with probability .6 that it will. From, this we can infer that i believes with probability .8 that a crash will come and with probability $.8 \times .5 + .2 \times .6 = .52$ that $-i$ believes that the market will crash. In order to go further, we must also assign to i beliefs about the beliefs of $-i$ concerning i's beliefs. Thus, let us assume now that i believes with probability .8 that it will crash and that $-i$ believes with probability .5 that it will and that i will believe in the crash with probability .2. And let i believe with probability .2 that it won't, but that $-i$ believes with probability .6 that it will and that $-i$ furthermore is certain, according to this particular belief of i, that i believes in a crash. Then, in addition to the preceding, i believes with probability $.8 \times .2 + .2 \times 1 = .36$ that $-i$ believes that she will believe in a crash. And we can iterate this now, that is, assign to $-i$ also proabilities for i's beliefs about his belief in a crash, and then compute the probability that i believes that $-i$ believes that she believes that he will believe in a crash. Instead of such iterations, we can again obtain a shortcut by considering types. For instance, i may have the types t_i^1, where she believes with probability .8 in the crash and with probability .5 that $-i$ is of type t_{-i}^1 and with probability .5 of type t_{-i}^2, or t_i^2, where she believes with probability 1 in the crash and that $-i$ is of type t_{-i}^1. And in turn, let t_{-i}^1 mean that $-i$ believes with probability .6 in the crash and with probability 1 that i is of type t_i^1, and let t_{-i}^2 mean that $-i$ believes with probability .2 in the crash and that with probability .5, i is of type t_i^1. Then, when i is of type t_i^1, she believes with probability .8 in the crash, with probability $.5 \times .6 + .5 \times .2 = .4$ that $-i$ believes in the crash, with probability $.5 \times 1 \times .8 + .5(.5 \times .8 + .5 \times 1) = .85$ that $-i$ believes that i believes in the crash, and so on.

We first observe the following direct consequence of Theorem 2.2.1.

Lemma 2.4.1. *If S is a compact metrizable space, then so is each X^k.*

So, henceforth, we shall always assume that S is such a space. This is unproblematic, as in our examples and applications, S usually is finite.

According to (2.4.12), an element of X^k comprises beliefs about $S, S \times X^1, \dots, X^{k-2}, X^{k-1}$, and an element of X^{k-1} in turn comprises beliefs about $S, S \times X^1, \dots, X^{k-2}$. In principle, these beliefs may contradict each other. That is, the belief of an element in X^k about $S \times X^{k-2}$ might be different from the belief that the resulting belief about $S \times X^{k-1}$ induces on $S \times X^{k-2}$. This is, however, not what we want when we consider the iterative beliefs induced by types. When i has a joint belief about S and the types of $-i$, then, when she marginalizes over S, that is, integrate over the types of $-i$, this should reduce to her original belief about S. This should not be confused with the fact that her belief about the states in S and the types of $-i$ need not be a product belief. She could well believe that every state in S is linked to one particular type of $-i$. That is, slightly varying our above example, let us assume that $S = \{0, 1\}$ and i assigns nonzero probabilities only to the two types t_{-i}^0, t_{-i}^1, with the probabilities $p(0, t_{-i}^0) = .4, p(1, t_{-i}^1) = .6, p(0, t_{-i}^1) = 0 = p(1, t_{-i}^0)$. Then $p_S(0) = .4, p_S(1) = .6$, but of course the latter can also be achieved by other joint probability distributions, for instance by the product distribution $p(0, t_{-i}^0) = .2 = p(1, t_{-i}^1), p(1, t_{-i}^0) = .3 = p(1, t_{-i}^1)$. However, if she has a probability distribution $p(\cdot)$ on S, then coherence would require that this distribution coincides with the marginal $p_S(\cdot)$ of the probability distribution $p(\cdot, \cdot)$ on $S \times \mathcal{P}(S)$. Or put differently, when i has a joint distribution over states and opponent types, then marginalizing over the states should yield her distribution for the states. If not, she should simply replace the latter by the former.
Note that coherence does not mean that the beliefs of different orders coincide. For instance, the probabilities that i assigns to the actions of $-i$ may well be different from those that she thinks that $-i$ believes that i herself has. In other words, it is perfectly possible that i believes that $-i$ has incorrect beliefs about her own beliefs.

We now formulate this coherence condition in general

Definition 2.4.3. A family $q_k \in X^k$ is coherent if

$$\begin{aligned} q^k &= (q^{k-1}, r^{k-1}) \text{ with} \\ r^{k-1}_{S\times X^{k-2}} &= r^{k-2} \text{ where} \\ q^{k-1} &= (q^{k-2}, r^{k-2}); \end{aligned} \tag{2.4.13}$$

here, $r^{k-1} \in \mathcal{P}(S \times X^{k-1}) = \mathcal{P}(S \times X^{k-2} \times \mathcal{P}(S \times X^{k-2}))$.

Putting the marginal condition (2.4.13) differently, we have

$$q^k = (q^1; r^1, \dots, r^{k-1}), \tag{2.4.14}$$

where q^1 represents i's belief about S and $r^j \in \mathcal{P}(S \times X^j)$ is a probability distribution over S and the beliefs of order $\leq j$ of the opponent $-i$.
Thus, the coherence condition means that the marginal distribution of r^1 over S be q^1 and that the marginal distribution of r^k over $S \times X^j$ is r^j, whenever $1 < j < k$.
With (2.4.14), we can also express the coherence condition by the requirement that the projection of q^k onto X^j be q^j, for $1 \leq j < k$.

We have seen above that two families of types T_i, T_{-i} generate coherent systems of beliefs. It is now an important insight that, conversely, every such coherent system of beliefs is generated by types of the participants. Since this is somewhat technical, we need to refer for some details to [12], but we shall try to explain the principle.

Definition 2.4.4. The universal type space T consists of all coherent families $q_k \in X^k$, that is, all families that are of the iterative form $q^k = (q^{k-1}, r^{k-1})$ and satisfy (2.4.13).

The fundamental result then is

Theorem 2.4.1. *T is a compact space, and*

$$T = \mathcal{P}(S \times T). \tag{2.4.15}$$

Here, "=" means that the two spaces are homeomorphic to each other, that is, there exists a continuous bijective map between the two compact spaces T and $\mathcal{P}(S \times T)$ (the latter space is compact by Theorem 2.2.1 if the former is).

This result tells us that a type generates a probability distribution over (states and) types. Thus, this formalizes the iterative principle described above. When we have a type in the sense of the beginning of this section, then by recursion it generates a coherent belief sequence, hence a member of the universal type space, and conversely. The proof of Theorem 2.4.1 then formalizes this idea.

Proof. For the proof, we shall construct a map

$$h : \mathcal{P}(S \times T) \to T \tag{2.4.16}$$

and show that it is a homeomorphism, that is, bijective and continuous in both directions. Thus, let $q \in \mathcal{P}(S \times T)$. Let q^1 be its marginal distribution on S, and let r^k be its marginal distribution on $S \times X^k$, and put $q^{k+1} = (q^k, r^k)$ for $k = 1, 2, \dots$. Then, because T only contains coherent belief hierarchies, the projection of q^k onto X^j for $1 \leq j < k$ is q^j. Thus, from $q \in \mathcal{P}(S \times T)$, we have constructed an element of T. This defines our map h, and we now need to check its properties.

First, h is continuous w.r.t. the weak-$\star$ topologies, because if q_n converges to some q w.r.t. to this topology, then so do all its projections q_n^k, with limits the corresponding projections q^k of q.

The construction of the inverse of h is the most interesting part of the proof. This means that we can reconstruct q from $h(q) = (q^1, q^2, \dots)$. The idea is to disentangle the coherent hierarchy of beliefs and apply the Kolmogorov extension theorem. To see the principle, we first simplify the situation. The family $(q^1, q^2, \dots)$ gives us a family of probability distributions $q_{n_1,\dots,n_k}$ on S^k with

$$q_{n_1,\dots,n_{j-1},n_{j+1},\dots,n_k}(A_1 \times \dots A_{j-1} \times A_{j+1} \times \dots A_k) = q_{n_1,\dots,n_{j-1},n_j,n_{j+1},\dots,n_k}(A_1 \times \dots A_{j-1} \times S \times A_{j+1} \times \dots A_k) \tag{2.4.17}$$

for all $1 \leq n_1 < n_2 \dots n_k, 1 \leq j \leq k$ and all Borel subsets A_i of S, $i = 1, \dots, k$. This is simply the coherence condition (called "consistency" in the theory of stochastic processes), that we can marginalize over intermediate stages. The Kolmogorov extension theorem then says that this family of distributions is generated by a stochastic process, that is, a probability measure p on the sequence space $S^{\mathbb{N}}$ and a map

$$\xi_n : S^{\mathbb{N}} \to S \tag{2.4.18}$$

with

$$q_{n_1,\dots,n_k}(A_1 \times \dots A_k) = p(\xi_{n_1} \in A_1, \dots, \xi_{n_k} \in A_k) \tag{2.4.19}$$

for all Borel subsets $A_i \subset S$. The probability measure p, a so-called cylinder measure, is in fact characterized by (2.4.19). A similar version holds with the spaces X^k, but we leave that to the reader. This probability measure is, of course, the object that we want. That is, from the sequence of marginals, i.e., an element of T, we have reconstructed the probability measure, that is, an element of $\mathcal{P}(S \times T)$ that generates this sequence of marginals. Thus, we have constructed the inverse of h. Also, this shows that h is surjective.

We can now conclude the proof by appealing to the general fact, Lemma 2.4.2 below, that a continuous and bijective map from a compact space to a Hausdorff space is a homeomorphism, noting that our spaces are metrizable, hence Hausdorff.

□

Lemma 2.4.2. *A continuous bijective map h from a compact topological space to a Hausdorff space is a homeomorphism.*

We sketch the

Proof. We need to show that the inverse h^{-1} is continuous. A map between topological spaces is continuous iff the preimages of open sets are open.

Now, any closed subset of a compact space is compact. Since h is continuous, it maps compact sets to compact sets. In turn, compact subsets of Hausdorff spaces are closed. (Here, a topological space X is called a Hausdorff space if any $x \neq y \in X$ have disjoint open neighborhoods. For a compact subset K of a Hausdorff space Y, its complement is open, and hence K itself is closed. In fact, for $x \notin K$, and every $y \in K$, by the Hausdorff property, there exist disjoint open neighborhoods U_y of y and V_y of X. Since K is compact, it is covered by finitely many of them, say $U_{y_1}, \dots, U_{y_k}$. But then $\bigcap_{i=1,\dots,k} V_{y_i}$ is an open neighborhood of x disjoint from K, showing the openness of the complement of K.)

Thus, h maps closed sets to closed sets. But then it also maps open sets to open sets. Equivalently, for its inverse h^{-1}, preimages of open sets are open, that is, it is continuous. □

Definition 2.4.5. The type t_i of i believes that $-i$ is rational iff $\operatorname{supp} q_i^2(t_i) \subset S_{-i} \times \mathcal{P}(S_i)$ contains only pairs (s_{-i}, μ_i) for which s_{-i} is a best response to μ_i. And the type t_i of i believes that both $-i$ is rational and believes that i is rational in turn iff $\operatorname{supp} q_i^3(t_i) \subset (S_{-i} \times \mathcal{P}S_i) \times \mathcal{P}(S_i \times \mathcal{P}(S_{-i}))$ (recall (2.4.10), (2.4.12)) contains only pairs $((s_{-i}, \mu_i), \nu_i)$ where s_{-i} is a best response to μ_i as before and ν_i is supported on pairs (s_i, μ_{-i}) where s_i is a best response to μ_{-i}. And so on for higher orders of belief in rationality.

2.5 Bayesian equilibria

In this section, we shall mostly follow [7]. After having formalized beliefs and types, we now wish to apply this to game theory and determine equilibria in games with incomplete information. The setting will, however, be somewhat simpler than in the preceding section as we shall assume that the uncertainty structure of the game is common knowledge among the players. That is, the values of certain variables may be unknown to some or all players, but they will be distributed according to some distribution p. This distribution will be assumed to be known to all players, and every player knows that this distribution is known, knows that everybody knows that it is known, and so on. In such a situation, we speak of common knowledge of p. Then, some of the players may have private knowledge, for instance know their type. Every other player then knows that, say, player j has private knowledge. The other players then do not know j's type, but they know that j knows his type, and they also know the distribution from which j's type is drawn.

We consider a game where player i has possible types $t_i^\alpha \in T_i$ and a pure strategy space S_i. The types are distributed according to some commonly known distribution p. W.l.o.g., each type is assumed to occur with positive probability.

(Types with zero probability can thus be ignored; we do not discuss here the somewhat subtle issue whether it makes a difference whether some possibility is unknown and unconceivable to the players, that is, they are unaware of it, or in principle conceived, but simply assigned zero probability, that is, they consider this possibility as impossible to occur.) The utility u_i of player i may depend on her own action, those of her opponents, as well as one the values of the types,

$$u_i(s_i, s_{-i}, (t_i, t_{-i})). \tag{2.5.20}$$

Definition 2.5.1. We expand the game by assigning to player i the strategy set $S_i^{T_i}$, the set of all maps $T_i \to S_i$. That is, player i can choose some strategy for each type of hers.

A Bayesian equilibrium of a game of incomplete information is a Nash equilibrium of this expanded game.

Of course, the existence of a Bayesian equilibrium in such a game of incomplete information is a direct consequence of the Nash equilibrium theorem. It was Harsanyi's idea to formulate games of incomplete information in such a manner that Nash's theorem becomes applicable.

Again, an equilibrium could be pure or mixed. A family of mappings $s_j : T_j \to S_j$ for all players j is a pure Bayesian equilibrium iff

$$s_i(.) = \mathrm{argmax}_{\sigma_i \in S_i^{T_i}} \sum_{t_i} \sum_{t_{-i}} p(t_i, t_{-i}) u_i(\sigma_i(t_i), s_{-i}(t_{-i}), (t_i, t_{-i})) \text{ for all } i. \tag{2.5.21}$$

Here, i maximizes her expected utility for every type of hers. Note that the probabilities $p(t_i, t_{-i})$ of the type constellations need not be products, that is, the types of the different players could be correlated. For instance, if the types of the players are "sick" or "healthy", then in the case of major disease, many players will be sick simultaneously.

We can also express (2.5.21) by saying that for each type t_i, player i chooses her action in such a way that her utility is maximized given the distribution of types and the resulting actions of her opponents. Equivalently, we have

$$s_i(t_i) = \mathrm{argmax}_{\sigma_i \in S_i} \sum_{t_{-i}} p(t_{-i}|t_i) u_i(\sigma_i, s_{-i}(t_{-i}), (t_i, t_{-i})), \tag{2.5.22}$$

that is, for each type t_i, i maximizes her utility given the conditional distribution of the types of her opponents. (When players have infinitely many possible types, the sums in (2.5.22) are replaced by integrals.)

Here, higher order beliefs of the players do not enter. All that is relevant is the belief about the distribution of the opponents' types and their type-dependent strategies. Of course, the point of the previous section was that such a belief about the distribution of types can result from a coherent higher order belief structure.

In the terminology of game theory, we are considering here an "ex ante" situation. That means that the predictions about the opponents are made independently of each player's own type. In contrast, an "interim" situation would be one where the predictions a player makes about her opponents depend on her own type.

Let us now use these conceptual tools to analyze some **examples:**

1. This is a version of the public goods game where each player i has the choice of contributing at a cost c_i or not contributing, and where both players earn 1 unit if at least one of them contributes.

	Contribute	Don't
Contribute	$1-c_A, 1-c_b$	$1-c_A, 1$
Don't	$1, 1-c_B$	0,0

(2.5.23)

Thus, each player prefers to let the other contribute.

We assume that the cost c_i is only known to i herself, but both players know that both costs c are randomly drawn from the same continuous distribution with a strictly increasing cumulative distribution function $P(c) = p(\text{cost} \le c)$ on an interval $[\underline{c}, \overline{c}]$.

The cost c_i then becomes i's type. We thus consider functions $s_i : [\underline{c}, \overline{c}] \to \{0,1\}$ where 1 means "contribute", 0 "don't contribute". Let (s_A, s_B) be a Bayesian equilibrium, and let $\pi_i := p(s_i(c_i) = 1)$ be the probability that i contributes at this equilibrium. Thus, when i contributes, her expected pay-off is $1 - c_i$ whereas if she doesn't, it is π_{-i}. Thus, she should contribute, that is, $s_i(c_i) = 1$, when

$$c_i < 1 - \pi_{-i} =: c_i^*. \tag{2.5.24}$$

Therefore, her probability of contributing satisfies

$$\pi_i = p(\underline{c} \le c_i \le c_i^*) = P(c_i^*). \tag{2.5.25}$$

From (2.5.24), (2.5.25)

$$c_i^* = 1 - P(c_{-i}^*), \tag{2.5.26}$$

and hence

$$c_i^* = 1 - P(1 - P(c_i^*)). \tag{2.5.27}$$

We first consider the case where (2.5.27) has a single solution c^*. For instance, when p is the uniform distribution on the interval $[0, m]$, i.e., $P(c) = \frac{c}{m}$ for $0 \le c \le m$, then $c^* = \frac{m}{m+1}$. In particular, the optimal strategy of a player depends not only on his cost, but on the distribution of possible costs. This comes from the uncertainty about the opponent's costs. The optimal threshold is an increasing function of m, the upper cost bound. When the other player faces a potentially higher cost, he is less likely to contribute, and therefore, one should increase one's own threshold of contributing to be sure to reap the benefit.

If, on the other hand, $\underline{c} > 1 - P(1)$, then there can be no solution $c^* = c_i^* = c_{-i}^*$ of (2.5.26), and there are instead two assymmetric equilibria where one player always contributes as long as her cost is below 1 and the other never contributes.

2. In order to prepare for the next game, we first consider a monopoly situation where a single player M decides about his output s to maximize his utility function

$$u(s) = s(1 - s). \tag{2.5.28}$$

The optimal value then is $s = \frac{1}{2}$. We now play this game with two players A and B (a Cournot duopoly), and now assume that they have to share the market so that the profit of player i with type t_i playing $s_i \in \mathbb{R}^+$ is given by

$$u_i = s_i(t_i - s_i - s_{-i}) \tag{2.5.29}$$

We assume that player A's type is fixed to be $t_A = 1$, whereas player B's type follows a probability distribution $p(t_B)$, commonly known to both players. Thus, both players know player A's type, but player B's type is only known to B himself. The expected pay-off u_A of A is

$$\int s_A(1 - s_A - s_B(t))p(t)dt, \tag{2.5.30}$$

which, using $\int p(t)dt = 1$, is maximized by

$$s_A = \frac{1}{2} - \frac{1}{2}\int s_B(t)p(t)dt. \tag{2.5.31}$$

Since B knows that, he can insert this value s_A into (2.5.29) and for each type t_B then maximize u_B as given by (2.5.29) and therefore choose

$$s_B(t_B) = \frac{t_B - s_A}{2}. \tag{2.5.32}$$

Inserting (2.5.32) in turn into (2.5.31) yields

$$s_A = \frac{2}{3} - \frac{1}{3}\int tp(t)dt, \tag{2.5.33}$$

and reinserting this into (2.5.32) finally gives

$$s_B(t_B) = \frac{t_B}{2} - \frac{1}{3} + \frac{1}{6}\int tp(t)dt. \tag{2.5.34}$$

In particular, in the trivial case where B only has the unique type $t_B = 1$, that is, where $p(t)$ is the Dirac delta distribution δ_1, this becomes

$$s_A = s_B = \frac{1}{3}. \tag{2.5.35}$$

As a side remark, when we compare this with the monopoly player of (2.5.28), the total output $s = s_i + s_{-i}$ now is $\frac{2}{3}$ which is larger than the monopoly output of $\frac{1}{2}$. Therefore, the market price will be lower, and more will be sold at lower price; that is, competition of producers is good for the consumers according to this stylized model. The sum of the utilities $u_i + u_{-i}$ is the same as the utility u, but now neither player decides about the total output, but each can fix only her/his individual contribution to that output.

Returning to our analysis, nevertheless, from the epistemic perspective, there are some subtleties here. A could think that B will always follow her value s_A and choose s_B according to (2.5.32). Anticipating that, A could insert this value into (2.5.30) before maximizing w.r.t. s_A. That is, A would then maximize

$$\int s_A(1 - s_A - \frac{t - s_A}{2})p(t)dt = \int s_A(1 - \frac{s_A}{2} - \frac{t}{2})p(t)dt \tag{2.5.36}$$

w.r.t. s_A, leading to

$$s_A = 1 - \int \frac{t}{2}p(t)dt, \tag{2.5.37}$$

which in turn leads to

$$s_B(t_B) = \frac{t_B - 1}{2} + \int \frac{t}{4}p(t)dt. \tag{2.5.38}$$

In particular, in the trivial symmetric case $p(t) = \delta_1(t)$, we get the asymmetric result

$$s_A = \frac{1}{2}, \quad s_B = \frac{1}{4}. \tag{2.5.39}$$

The pay-offs then are also unequal,

$$u_A = \frac{1}{8}, \quad u_B = \frac{1}{16}, \tag{2.5.40}$$

whereas in the symmetric case (2.5.35), each player would get $u = \frac{1}{9}$. Thus, in the asymmetric situation, A can slightly increase her pay-off at the expense of B. One possible interpretation of this is a so-called leader-follower game. In such a game, A, the leader, can determine and announce her action first, and B, the follower, then is forced to react to it. The increase in pay-off that A can achieve in this situation when compared to the case of independent simultaneous actions is called the leader's advantage.

Of course, this is not an epistemic interpretation, and the question remains under which epistemic conditions, that is, players forming beliefs about each other, such an outcome is possible.

The equilibrium concept of Definition 2.5.1 did not refer to higher order beliefs. In fact, it is difficult to formulate general equilibrium concepts in such situations, because at each stage of the iteration of the belief hierarchy, the beliefs might be different. For instance, i might believe something that is different from what she thinks that $-i$ believes that she believes. The following simplifying concept rules this out. (We formulate it again for two players only for simplicity. In the general case, one needs to include beliefs about the opponents of i about each other and the player i to arbitrary order.)

Definition 2.5.2. A combination (q_i, q_{-i}) of beliefs of players generates the simple belief hierarchy where each player i has the belief q_i about the opponent $-i$'s actions, believes that he has the belief q_{-i}, that he believes that i has the belief q_i, etc.

Note that this condition is much stronger and more restrictive than the above coherence condition of Definition 2.4.3. There, it was only required that at each stage of the belief hierarchy, the beliefs at each level are the same. Here, in contrast, it is required that the beliefs at different levels always coincide. In particular, it is excluded that a player ascribes false beliefs to an opponent.

In the presence of (mutually known) utility functions, such a simple belief hierarchy then encodes a common belief in rationality, or equivalently, is a Nash equilibrium, if each player i's beliefs only assign positive probabilities to actions of $-i$ that are optimal for $-i$ under his beliefs q_{-i}. The proof of the Nash equilibrium then shows that for games with finitely many actions available to each player, there always exists some such Nash equilibrium.

For general coherent belief hierarchies, equilibrium concepts and results become more subtle. The basic idea of rationality tells us that at no stage at the belief hierarchy, any player should be assumed to give positive probability to a suboptimal action. Putting this around, we can then eliminate such suboptimal actions from consideration, and iterating this elimination can considerably constrain the available options and in some cases even lead to a unique equilibrium. In technical language, one speaks of the iterated elimination of weakly dominated strategies. Here, a strategy is weakly dominant if there exists some other strategy that is never worse, but strictly better in at least one condition, that is, for one action combination of the opponents.

Bibliography

[1] C.Aliprantis, K.Border, Infinite dimensional analysis, Springer, [3]2006

[2] J.-P.Aubin, Optima and equilibria, Springer, 1993

[3] N.Bertschinger, J. Jost, E. Olbrich, D. Wolpert, A simple topological proof of the Nash equilibrium theorem of game theory, unpublished

[4] A.Brandenburger, E. Dekel, Hierarchies of Beliefs and Common Knowledge, Journal of Economic Theory 59, 189-198, 1993

[5] H.Gintis, Game theory evolving, Princeton Univ.Press, 2000

[6] D.Fudenberg, D.Levine, The theory of learning in games, MIT Press, 1998

[7] D.Fudenberg, J.Tirole, Game theory, MIT Press, 1991

[8] J.Harsanyi, Games of Incomplete Information Played by Bayesian Players. Parts I, II, III, Management Science 14, 159-182, 320-334, 486-502, 1967-68

[9] J.Hofbauer and K.Sigmund, Evolutionary games and population dynamics, Cambr.Univ.Press, 1998

[10] J.Jost, Dynamical systems. Examples of complex behavior, Springer, 2005

[11] J.Jost, Postmodern analysis, Springer, [3]2005

[12] M.Maschler, E.Solan, S.Zamir, Game theory, Cambridge Univ.Press, 2013

[13] D.McFadden, Quantal choice analysis: A survey, Ann.Econ.Soc.Measurement 5, 363-390, 1976

[14] R.McKelvey, T.Palfrey, Quantal response equilibria for normal form games, Games Econ.Beh.10, 6-38,1995

[15] R.McKelvey, T.Palfrey, A statistical theory of equilibrium in games, Japan.Econ.Review 47, 186-209, 1996

[16] R.Myerson, Game theory. Analysis of conflict, Harvard Univ.Press, 1991

[17] A.Perea, Epistemic game theory. Reasoning and choice, Cambr.Univ.Press, 2012

[18] M.Siniscalchi, Epistemic game theory: Beliefs and types

[19] R. Stöcker, H.Zieschang, Algebraische Topologie, Teubner, ^{2}1994

[20] F.Vega-Redondo, Economics and the theory of games, Cambr.Univ.Press, 2003

[21] J.Weibull, Evolutionary game theory, MIT Press, 1995

[22] D.Wolpert, M.Harré, It can be smart to be dumb, Preprint

[23] D.Wolpert, M.Harré, N.Bertschinger, E.Olbrich, J.Jost, Hysteresis effects of changing parameters of noncooperative games, PRE

[24] P.Young, Individual strategy and social structure, Princeton Univ.Press, 1998

[25] E.Zeidler, Nonlinear functional analysis and its applications, Vol.IV, Springer, 1987