Lecture Notes on Game Theory

Tilman Börgers

April 12, 2024

Contents:

Topic 1: Nash Equilibria In Pure Actions

A game is a triple $(I, (A_i)_{i\in I}, (u_i)_{i\in I})$, where:

- $I = \{1, 2, \ldots, n\}$ is a finite set of players;
- for every $i \in I$ the set A_i is the set of actions of players i;
- for every $i \in I$ the function $u_i : \bigtimes_{i\in I} A_i \to \mathbb{R}$ is player i's utility function.

For the moment, the functions u_i represent ordinal preferences over $\bigtimes_{i\in I} A_i$.

Note that we are allowing arbitrary externalities.

Given a game $(I, (A_i)_{i\in I}, (u_i)_{i\in I})$ a Nash equilibrium of this game is an n-tuple of actions $(a_i^*)_{i\in I}$ such that for every player $i \in I$ and for every action $a_i \in A_i$:

$$u_i(a^*) \geqslant u_i(a_i, a_{-i}^*),$$

where (a_i, a_{-i}^*) is the list of actions that is identical to a^* except that we have replaced a_i^* by a_i.

Interpretation: A Nash equilibrium represents a rest point of a learning process in which each player chooses the optimal action assuming that all other players choose the same actions as in the previous period.

Agenda for Topic 1:

1. Nash equilibria are fixed points of the best reply correspondence.
2. Finding Nash equilibria by determining fixed points of the best reply correspondence.
3. Are Nash equilibria Pareto efficient?

4. When do Nash equilibria exist?

5. When are Nash equilibria unique?

6. How do Nash equilibria change when the game changes?

1. Nash equilibria are fixed points of the best reply correspondence.

For every player $i \in I$ define a best reply correspondence $BR_i : \times_{i \in I} A_i \twoheadrightarrow A_i$ by setting for every $a \in \underset{i \in I}{\times} A_i$:

$$BR_i(a) = \{a_i \in A_i | u_i(a_i, a_{-i}) \geqslant u_i(a'_i, a_{-i}) \text{ for every } a'_i \in A_i\}.$$

The best reply correspondence $BR : \times_{i \in I} A_i \twoheadrightarrow \times_{i \in I} A_i$ is defined by setting for every $a \in \underset{i \in I}{\times} A_i$:

$$BR(a) = \underset{i \in I}{\times} BR_i(a).$$

By definition: a^* is a Nash equilibrium if and only if $a^* \in BR(a^*)$, i.e. if and only if a^* is a fixed point of BR.

2. Finding Nash equilibria by determining the best reply correspondence.

Example 1: $I = \{1, 2\}$, $A_i = \mathbb{R}_+$, and $u_i(a_1, a_2) = a_i \max\{10 - a_1 - a_2, 0\} - c_i a_i$, where for each $i = 1, 2$ c_i is a strictly positive constant. We assume $c_i < 10$ for $i = 1, 2$.

Let's determine player i's best reply correspondence. Let $j \neq i$. Obviously:

- If $a_j \geqslant 10 - c_i$, then $BR_i(a) = \{0\}$.

Suppose $a_j < 10 - c_i$. Then the first order condition for player i's maximization problem is:

$$10 - 2a_i - a_j - c_i = 0 \Leftrightarrow a_i = \frac{10 - a_j - c_i}{2},$$

and this is obviously all the unique maximizer.

- If $a_j < 10 - c_i$, then $BR_i(a) = \left\{\frac{10-a_j-c_i}{2}\right\}$.

Let's look for fixed points of BR.

When is there a fixed point such that $a_1 = 0$? The best reply of player 2 to $a_1 = 0$ is: $a_2 = \frac{10-c_2}{2}$. Player 1's best reply is $a_1 = 0$ if:

$$\frac{10 - c_2}{2} \geqslant 10 - c_1 \Leftrightarrow c_1 \geqslant \frac{10 + c_2}{2}.$$

Therefore, if $c_1 \geqslant \frac{10+c_2}{2}$, then there is a Nash equilibrium where $a_1 = 0$ and $a_2 = \frac{10-c_2}{2}$.

When is there a fixed point such that $a_2 = 0$? Similar reasoning as above leads to the conclusion that, if $c_2 \geqslant \frac{10+c_1}{2}$, then there is a Nash equilibrium where $a_2 = 0$ and $a_1 = \frac{10-c_1}{2}$.

When is there a fixed point such that $a_1 > 0$ and $a_2 > 0$? Both players have to satisfy the first order condition:

$$\begin{aligned} a_1 &= \frac{10 - a_2 - c_1}{2} \\ a_2 &= \frac{10 - a_1 - c_2}{2}. \end{aligned}$$

Subtracting yields:

$$a_1 - a_2 = c_2 - c_1.$$

Substituting into the first order condition for a_1 yields:

$$a_1 = \frac{10 - a_1 + (c_2 - c_1) - c_1}{2} \Leftrightarrow a_1 = \frac{10 + c_2 - 2c_1}{3}.$$

Analogously we get for a_2:

$$a_2 = \frac{10 + c_1 - 2c_2}{3}.$$

Such a Nash equilibrium exists if both a_i are positive, i.e. if:

$$c_1 < \frac{10 + c_2}{2} \quad \text{and} \quad c_2 < \frac{10 + c_1}{2}.$$

3. Are Nash equilibria Pareto efficient?

An action profile $a \in A$ is Pareto efficient if there is no other action profile $a' \in A$ such that: $u_i(a') \geqslant u_i(a)$ for all $i \in I$, and $u_i(a') > u_i(a)$ for at least one $i \in I$.

Consider the previous example. In an interior Nash equilibrium, we have:

$$\frac{\partial u_1}{\partial a_1} = 0 \quad \text{and} \quad \frac{\partial u_2}{\partial a_2} = 0.$$

But notice that we also have:

$$\frac{\partial u_1}{\partial a_2} = -a_1 < 0 \quad \text{and} \quad \frac{\partial u_2}{\partial a_1} = -a_2 < 0.$$

Therefore, when we lower a_1 and a_2 by a little bit, both players' utility will go up, and therefore the Nash equilibrium is not Pareto efficient.

Here is an illustration. Suppose $c_1 = c_2 = 6$. Then the unique Nash equilibrium is: $a_1 = a_2 = 4/3$, and therefore, each player's utility is: $u_i(a_1, a_2) = (4/3)(10-8/3)-6(4/3) = 16/9$. But, if players chose: $a_1' = a_2' = 1$, then their utilities would be: $u_i(a_1, a_2) = 1 \cdot 8 - 6 \cdot 1 = 2 > 16/9$.

Intuitively, the reason why Nash equilibria are not Pareto efficient is that players don't take externalities into account when they make their choices.

The most famous game with a Nash equilibrium that is not Pareto efficient is the Prisoner's Dilemma:

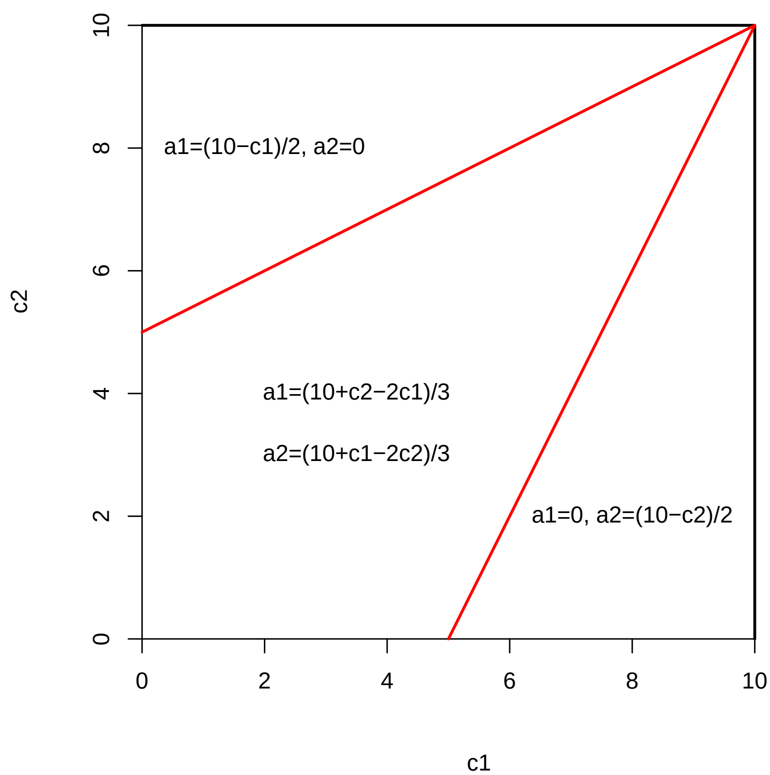

Figure 1: Nash equilibria in Example 1 as a function of the parameters c_1 and c_2

	C	D
C	3,3	0,4
D	4,0	1,1

The following example is known as "Pigou's Example," after Arthur Pigou. There are 10 drivers who commute every day from city A to city B. They can choose one of two routes: r_1 or r_2. The time it takes to travel on route 1 is equal to 10 units of time. The units of time that it takes to travel on route 2 is equal to the number of drivers who use it: n_2. Drivers seek to minimize their travel time. This game has the following Nash equilibrium: all drivers use route 2. But if, say, 5 drivers used route 1 and 5 drivers used route 2, the second set of drivers would only need 5 units of time. This example predates the Prisoners' Dilemma.

4. When do Nash equilibria exist?

The most famous game with no Nash equilibrium (in pure actions) is "Matching Pennies:"

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

When does $BR : \bigtimes_{i \in I} A_i \twoheadrightarrow \bigtimes_{i \in I} A_i$ have a fixed point? By Kakutani's Fixed Point Theorem it will have a fixed point if for every player $i \in I$ the set A_i is a non-empty, compact, convex subset of some finite dimensional Euclidean space, and BR_i is non-empty and convex-valued and upper hemi-continuous. The following result provides assumptions that guarantee this:

Proposition 1. *If for every player $i \in I$:*

1. *the set A_i is a non-empty, compact, and convex subset of $\mathbb{R}^{m_i}$ for some $m_i \in \mathbb{N}$;*
2. *the function u_i is continuous in a;*
3. *the function u_i is quasi-concave in a_i;*

then at least one Nash equilibrium exists.

Here is an alternative result that does not assume that the action sets are convex, nor that the utility functions are quasi-concave:

Proposition 2. *If for every player $i \in I$:*

1. *the set A_i is a non-empty and compact subset of $\mathbb{R}$;*

2. *the function u_i is continuous in a;*

3. *the function u_i has increasing differences, i.e. if: $i \in I$, $a_i, a_i' \in A_i$, $a_i \geqslant a_i'$, for every $j \in I$ with $j \neq i$: $a_j, a_j' \in A_j$ and $a_j \geqslant a_j'$ then:*

$$u_i(a_i, a_{-i}) - u_i(a_i', a_{-i}) \geqslant u_i(a_i, a_{-i}') - u_i(a_i', a_{-i}'),$$

then at least one Nash equilibrium exists.

This result is a consequence of the Tarski fixed point theorem.

5. When are Nash equilibria unique?

Of course, in general Nash equilibria need not be unique. Consider this example:

	L	R
T	1,1	0,0
B	0,0	2,2

This game has two Nash equilibria: (T, L) and (B, R). (I am happy to defend the claim that (T, L) is a perfectly reasonable Nash equilibrium.)

Some results on the uniqueness of Nash equilibria use concepts that we shall only introduce later.

One could derive uniqueness results using the contraction mapping fixed point theorem (Banach fixed point theorem). But I don't know of interesting assumptions on the primitives of a game that imply that the best replies are unique, and that the best reply function is a contraction mapping.

There are other uniqueness results in the literature, but they are not used very often.

A classic result is in: J. B. Rosen, Existence and Uniqueness of Equilibrium Points for Concave N-Person Games, *Econometrica*, 33 (1965), 520-534.

A related result is Theorem 2.6 in: James W. Friedman, *Game Theory With Applications to Economics*, Oxford University Press, 1986. This result is based on D. Gale and H. Nikaido, The Jacobian Matrix and Global Univalence of Mappings, *Mathematische Annalen* 159 (1965), 81-93.

Future versions of these notes will hopefully include simple versions of these results.

6. How Do Nash Equilibria Change When The Game Changes?

Consider the following two player game:

	L	C
T	1,0	1,1
B	0,-2	0,-1

This game has a unique Nash equilibrium: (T, C). Both players receive utility 1 in this Nash equilibrium.

Now suppose that player 2 had an additional strategy, R, and that the game were:

	L	C	R
T	1,0	1,1	0,2
B	0,-2	0,-1	1,0

This game has a unique Nash equilibrium: (B, R). Player 1 receives utility 1 in this Nash equilibrium, and player 2 receives utility 0. Thus, we have given player 2 an additional strategy, but as a consequence player 2's utility dropped. Having more actions need not be to a player's advantage.

It is also not always to a player's advantage if we increase that player's payoff. Consider the following two player game:

	L	R
T	0,1	1,2
B	1,-1	0,0

This game has a unique Nash equilibrium: (T, R). Player 2's utility is 2 in this Nash equilibrium.

Now suppose we increased player 2's utility for (T, L) from 1 to 3, and also her utility for (B, L) from -1 to 1. Then we obtain the game:

	L	R
T	0,3	1,2
B	1,1	0,0

This game has a unique Nash equilibrium: (B, L). Player 2's utility in this new Nash equilibrium is 1. Thus it is lower than it was before.

The one statement that we can make regarding the comparative statics of Nash equilibria is that under certain conditions the Nash equilibrium correspondence is upper hemicontinuous.

Suppose each player's utility function is a function of some parameter $x \in \mathbb{R}$: $u_i : A \times \mathbb{R} \to \mathbb{R}$.

Proposition 3. *If for every player $i \in I$:*

1. *the set A_i is a non-empty and compact subset of $\mathbb{R}^{m_i}$ for some $m_i \in \mathbb{N}$;*

2. *the function u_i is continuous in a and x:*

Then the correspondence $\mathcal{E}$ that assigns to every $x \in \mathbb{R}$ the set $\mathcal{E}(x)$ of Nash equilibria of the game with parameter x is upper hemicontinuous.

Appendix for Topic 1

Proof of Proposition 1

We mentioned in Section 1 of Topic 1 that Nash equilibria are fixed points of the best reply correspondence BR. Therefore, Nash equilibria exist if and only if this correspondence has at least one fixed point. Kakutani's fixed point theorem gives sufficient conditions under which a correspondence has at least one fixed point.

Kakutani's Fixed Point Theorem: Let X be a non-empty, compact and convex subset of $\mathbb{R}^n$ for some $n \in \mathbb{N}$ and suppose that $f : X \twoheadrightarrow X$ is a correspondence that maps every $x \in X$ into a non-empty, compact and convex subset $f(x)$ of X. Suppose also that the graph of X, that is the set $\{(x, x') \in X^2 | x' \in f(x)\}$ is a closed subset of X^2. Then f has at least one fixed point $x^* \in X$, that is, there is at least one $x^* \in X$ such that:

$$x^* \in f(x^*).$$

The domain and co-domain of the best reply correspondence BR is $\times_{i \in I} A_i$. Proposition 1 assumes that for every player $i \in I$ the set A_i is a non-empty, compact, and convex subset of $\mathbb{R}_i^m$ for some $m_i \in \mathbb{N}$. But this immediately implies that $\times_{i \in I} A_i$ is a non-empty, compact, and convex subset of $\mathbb{R}^n$ where $n = \sum_{i \in I} m_i$. Thus, the assumptions of Kakutani's fixed point theorem for the domain and co-domain of the correspondence under consideration are satisfied.

Next, we have to show that for every $a \in \times_{i \in I} A_i$ the set $BR(a)$ is a non-empty, compact, and convex subset of $\times_{i \in I} A_i$. This is the case if and only if for every player i the set $BR_i(a)$ is a non-empty, compact, and convex subset of $\times_{i \in I} A_i$. The non-emptiness, that is, the existence of an optimal action of player i given any profile a_{-i} of other player's action, follows from the Weierstrass theorem according to which any continuous function on a compact domain attains a maximum. To show that the set of maximizers is compact we note that it is bounded, because the set A_i is bounded. We postpone the proof that it is closed until the last step of this proof. That last step implies the closeness of the set of maximizers for fixed a_{-i} if, in the notation of that proof, we consider the constant sequence $a_{-i}^{\nu} = a_{-i}$ for every $\nu \in \mathbb{N}$.

We also need to show that the set of maximizers is convex. This follows from the quasi-concavity of each player's utility function in their own action a_i. Recall that quasi-concavity of an arbitrary function $f : X \to \mathbb{R}$ means that for any constant $k \in \mathbb{R}$ the set of arguments x of f such that the value $f(x)$ is equal to k or larger than k, that is the set $\{x \in X | f(x) \geqslant k\}$, is convex. Now fix any a_{-i}, and let a_i^* be any action that maximizes $u_i(a_i, a_{-i})$. We noted in the previous paragraph that at least one such a_i exists. Denote the utility that corresponds to a_i^* by u_i^*, that is, $u_i^* = u_i(a_i^*, a_{-i})$. The set of best responses $BR_i(a_{-i})$ is then the set $\{a_i | u_i(a_i, a_{-i}) \geqslant u_i^*\}$. (Of course, the $\geqslant$ sign will be for all elements of this set an $=$ sign, because otherwise u_i^* would not be the largest achievable utility.) But this set is a convex set according to the definition of quasi-concavity, if we apply this definition to the constant $k = u_i^*$.

The final condition of Kakutani's Fixed Point Theorem is that the set $\{(x, x') \in X^2 | x' \in f(x)\}$ is a closed subset of X^2. This set is often referred to as the "graph" of the correspondence f. To verify that BR satisfies this condition, it is sufficient to verify that every correspondence BR_i has a closed graph. We shall verify that the limit of every convergent sequence of points in the graph of BR_i is also in the graph of BR_i. Let $\nu \in \mathbb{N}$ be the index of our sequence. What we have to show is this. If for every $\nu \in \mathbb{N}$:

$$a_{-i}^\nu \in A_{-i} \text{ and } a_i^\nu \in BR_i(a_{-i}^\nu),$$

and if the sequences $\left(a_{-i}^\nu\right)_{\nu \in \mathbb{N}}$ and $\left(a_i^\nu\right)_{\nu \in \mathbb{N}}$ converge to a_{-i}^∞ and a_i^∞ respectively, then the limit is in the graph of BR_i, which means:

$$a_{-i}^\infty \in A_{-i} \text{ and } a_i^\infty \in BR_i\left(a_{-i}^\infty\right).$$

Now $a_{-i}^\infty \in A_{-i}$ because A_{-i} is by assumption compact, hence closed. To prove that $a_i^\infty \in BR_i\left(a_{-i}^\infty\right)$, note that $a_i^\nu \in BR_i(a_{-i}^\nu)$ for every $\nu \in \mathbb{N}$ means:

$$u_i\left(a_i^\nu, a_{-i}^\nu\right) \geqslant u_i\left(a_i, a_{-i}^\nu\right) \text{ for all } a_i \in A_i$$

for all $\nu \in \mathbb{N}$. Because u_i is continuous, the sequences on the left and on the right hand sides of these inequalities exist, and we have:

$$\lim_{\nu \to \infty} u_i\left(a_i^\nu, a_{-i}^\nu\right) \geqslant \lim_{\nu \to \infty} u_i\left(a_i, a_{-i}^\nu\right) \text{ for all } a_i \in A_i.$$

Again by the continuity of u_i this means:

$$u_i\left(a_i^{\infty}, a_{-i}^{\infty}\right) \geqslant u_i\left(a_i, a_{-i}^{\infty}\right) \text{ for all } a_i \in A_i,$$

i.e.

$$a_i^{\infty} \in BR_i\left(a_{-i}^{\infty}\right).$$

This completes the proof.

Sketch of the Proof of Proposition 2

I shall only make some remarks on the ideas on which the proof of Proposition 2 is based. I shall not give the complete proof.

Because Nash equilibria are fixed points of the best reply correspondence, it is useful to ask what the assumptions of Proposition 2 imply for these correspondences. A simple observation is:

Claim 1: For every player i, and for every $a_{-i} \in A_{-i}$ the set $BR_i(a_{-i})$ is non-empty and closed.

This follows from the same argument that we used in the previous section of this Appendix.

Next, define for every i and every $a_{-i} \in A_{-i}$:

$$\bar{a}_i(a_{-i}) = \max BR_i(a_{-i}).$$

Note that Claim 1 implies that this is well-defined.

Claim 2: For any player i and any action profiles $a_{-i}, \hat{a}_{-i} \in A_{-i}$ such that $a_j \geqslant \hat{a}_j$ for all $j \neq i$ we have:

$$\bar{a}_i(a_{-i}) \geqslant \bar{a}_i(\hat{a}_{-i}).$$

This claim says that the upper bound of a player's best response set is increasing in the other players' actions.

Proof of Claim 2: The proof is indirect. Suppose under the assumptions of Claim 2 we had:

$$\bar{a}_i(a_{-i}) < \bar{a}_i(\hat{a}_{-i}).$$

Because $\bar{a}_i(\hat{a}_{-i})$ is a best response to $\hat{a}_{-i}$ we have:

$$u_i(\bar{a}_i(\hat{a}_{-i}), \hat{a}_{-i}) - u_i(\bar{a}_i(a_{-i}), \hat{a}_{-i}) \geqslant 0.$$

By assumption 3 of Proposition 2 we have:

$$u_i(\bar{a}_i(\hat{a}_{-i}), a_{-i}) - u_i(\bar{a}_i(a_{-i}), a_{-i}) \geqslant u_i(\bar{a}_i(\hat{a}_{-i}), \hat{a}_{-i}) - u_i(\bar{a}_i(a_{-i}), \hat{a}_{-i}),$$

and therefore:

$$u_i(\bar{a}_i(\hat{a}_{-i}), a_{-i}) - u_i(\bar{a}_i(a_{-i}), a_{-i}) \geqslant 0,$$

which contradicts the fact that by construction $\bar{a}_i(a_{-i})$ is the largest best response to a_{-i}.

Q.E.D.

Instead of looking for a fixed point of the correspondence BR_i, we can now look for a fixed point of the function $\bar{a}_i$. The only property of the function $\bar{a}_i$ that we have found so far is that it is increasing. Is this enough to obtain a fixed point?

Let us simplify the setting a little bit, and let us consider an abstract, single-dimensional setting. Let us show that monotonicity may be enough to obtain the existence of a fixed point.

Claim 3: Let $f : [0,1] \to [0,1]$ be a monotonically increasing function, that is: $x \geqslant y \Rightarrow f(x) \geqslant f(y)$. Then f has at least one fixed point, that is, $f(x) = x$ for at least one $x \in [0,1]$.

Note that we have not assumed in Claim 3 that f is continuous. When f is assumed to be continuous, the existence of a fixed point follows by Brouwer's fixed point theorem even if f is not monotone. I found the following proof of Claim 3 on *Stackexchange*.

Proof of Claim 3: If $f(0) = 0$, then $x = 0$ is a fixed point. So, let's focus in the remainder on the case: $f(0) > 0$. Similarly, if $f(1) = 1$, then $x = 1$ is a fixed point. So, let's focus in the remainder on the case: $f(1) < 1$.

Define $\ell_1 = 1$ and $h_1 = 1$. We have just assumed that $f(\ell_1) > \ell_1$, and $f(h_1) < 1$. We shall now define a recursive algorithm that generates a sequence of numbers ℓ_n and h_n for $n \in \mathbb{N}$ with $n \geqslant 2$. For $n \geqslant 2$, let $a_n = \frac{\ell_{n-1}+h_{n-1}}{2}$. If $f(a_n) = a_n$, stop the algorithm, because you have found a fixed point: $x = a_n$. But otherwise, if $f(a_{n-1}) > a_{n-1}$ set $\ell_n = a_{n-1}$ and $h_n = h_{n-1}$, and if $f(a_{n-1}) < a_{n-1}$ set $\ell_n = \ell_{n-1}$ and $h_n = a_{n-1}$. If this algorithm stops in finite time, we have found a fixed point. So, let's focus on the case that the algorithm does not stop in finite time.

Observe that for every $n \in \mathbb{N}$ we have: $f(\ell_n) > \ell_n$ and $f(h_n) < h_n$. Also observe that the sequence $(\ell_n)_{n\in\mathbb{N}}$ is non-decreasing, and the sequence $(h_n)_{n\in\mathbb{N}}$ is non-increasing. The two sequences therefore both converge. Let their respective limits be ℓ^* and h^*. Because the distance $h_n - \ell_n$ halves in each step, we must have: $\ell^* = h^*$. We now claim that $f(\ell^*) = \ell^*$, that is, $x = \ell^*$ is a fixed point of f. By proving this we shall complete the proof of Claim 2.

Suppose $f(\ell^*) > \ell^*$. Then, for sufficiently large n, we have: $h_n < f(\ell^*)$. Also, by construction, $f(h_n) < h_n$. Thus, for sufficiently large n, we have: $f(h_n) < f(\ell^*)$. But, by construction, $h_n \geqslant \ell^*$. Thus, $f(h_n) < f(\ell^*)$ contradicts the assumption that f is monotonically increasing. Thus, $f(\ell^*) > \ell^*$ leads to a contradiction. Similarly, $f(\ell^*) < f(\ell^*)$ leads to a contradiction. Therefore, $f(\ell^*) = \ell^*$.

Q.E.D.

Tarski's fixed point theorem is a multi-dimensional generalization of Claim 3. It is this generalization that we omit here.

Topic 2: Nash Equilibria In Mixed Actions

We shall now introduce the possibility that players choose their actions randomly. We call this "randomization." We refer to a probability distribution over actions as a "mixed action." We begin, by briefly considering mixed actions in a setting with just one decision maker.

1. Randomization in One Person Decision Problems

- Consider a decision maker who has to choose one action a from a finite set of actions A.
- Suppose each action $a \in A$ is risky, and leads to a probability distribution over possible outcomes.
- Suppose the decision maker satisfies the von Neumann Morgenstern axioms, and evaluates actions by calculating the corresponding expected utility.
- Let the expected utility of action $a \in A$ be denoted by $U(a)$.
- A rational choice for the decision maker is then any action a^* such that $U(a^*) \geqslant U(a)$ for all $a \in A$.
- Might this decision maker do even better by choosing the action randomly?

Let us denote by $\Delta(A)$ the set of all probability distributions over A. Thus, an element α of $\Delta(A)$ is a function $\alpha : A \to [0, 1]$ such that:

$$\sum_{a \in A} \alpha(a) = 1.$$

The elements of $\Delta(A)$ are called "mixed actions." What is the decision maker's expected utility if choosing probability distribution α? By the "law of iterated expectations" it is:

$$U(\alpha) = \sum_{a \in A} \alpha(a) U(a).$$

Note that this can never be more than $U(a^*)$. This is because each term $U(a)$ that appears in the sum on the right hand side is no more than $U(a^*)$. Moreover, obviously, the expected utility from choosing α is equal to $U(a^*)$ if and only if α assigns probability 1 to a^*, or, if there are multiple actions in A that are optimal, if and only if α assigns probability 1 only to actions in A that are optimal. This, fairly obvious, observation can be written as follows:

Claim: $U(\alpha^*) \geqslant U(\alpha)$ for all $\alpha \in \Delta(A)$ if and only if:

$$\{a \in A | \alpha(a) > 0\} \subseteq \{a \in A | U(a) \geqslant U(a') \text{ for all } a' \in A.\}$$

The set on the left hand side is known as the "support" of α^*. We have thus found that α^* is optimal if and only if its support is a subset of the set of all actions in A that maximize expected utility in A.

There is thus never any strict benefit to randomization. Moreover, if it is optimally to choose a random action, then a player could obtain exactly the same expected utility by not randomizing, and instead picking deterministically one of the actions to which the optimal randomization assigns positive probability. It is therefore redundant to introduce randomization in a single agent setting where the agent is assumed to follow the von Neumann Morgenstern axioms.

2. Nash Equilibria in Mixed Actions

Now let us return to the setting of game theory. Let us recall our definition of a game. To make it easy to define probability distributions over actions, I shall now assume that the set of actions is finite.

A game is a triple $(I, (A_i)_{i \in I}, (u_i)_{i \in I})$, where:

- $I = \{1, 2, \ldots, n\}$ is a finite set of players;
- for every $i \in I$ the finite set A_i is the set of actions of players i;
- for every $i \in I$ the function $u_i : \bigtimes_{i \in I} A_i \to \mathbb{R}$ is player i's von Neumann Morgenstern utility function.

A mixed action of player i is a probability distribution α_i over A_i. The set of all mixed actions is denoted by $\Delta(A_i)$.

We do not distinguish between an action $a_i \in A_i$ and the mixed action α_i that puts probability 1 on a_i.

Let us denote by α a list of mixed actions $(\alpha_1, \alpha_2, \ldots, \alpha_n)$. We shall write (α_{-i}) for a list of mixed actions of all players other than i. We shall write (α_i, α_{-i}) for the list of mixed actions of all players that we obtain if we complete the list α_{-i} by inserting mixed action α_i for player i.

It is important to note how we calculate expected utilities. If α is a list of mixed actions, then the expected utility of player i is:

$$U_i(\alpha) = \sum_{(a_1,a_2,\ldots,a_n)\in A} u_i\left(a_1, a_2, \ldots, a_n\right) \cdot \alpha_1(a_1) \cdot \alpha_2(a_2) \cdot \ldots \cdot \alpha_n(a_n).$$

Note that, on the right hand side, we multiply the probabilities of the different actions. This means that we are assuming that different players' randomizations are stochastically independent.

A Nash equilibrium in mixed actions of a game is an n-tuple of mixed actions α^* such that for every player $i \in I$ and for every mixed action $\alpha_i \in \Delta(A_i)$:

$$U_i(\alpha^*) \geqslant U_i(\alpha_i, \alpha^*_{-i}).$$

Because we identify pure actions a_i with the mixed action α_i that puts probability 1 on a_i, we can say that every Nash equilibrium in pure actions is also a Nash equilibrium in mixed actions.

Let's consider an example. Recall "Matching Pennies:"

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

We saw earlier that this game has no Nash equilibrium in pure actions. But it has a Nash equilibrium in mixed actions:

$$\alpha_i(H) = \alpha_i(T) = 1/2 \text{ for } i = 1, 2.$$

Let us verify that this is a Nash equilibrium. For each player, the expected utility from playing H is 0, and also the expected utility from playing T is zero. The Claim from the previous section implies that each player i maximizes expected utility by playing H, or playing T, or playing any probability distribution over H and T. Thus, in particular, it is optimal to choose H and T with probability 1/2 each.

We invoked the Claim from the previous section when discussing Matching Pennies. In fact, we can use the Claim from the previous section to infer the following general result:

Proposition 1. *α^* is a Nash equilibrium if and only if for every player i:*

$$\{a_i \in A_i | \alpha_i^*(a_i) > 0\} \subseteq \{a_i \in A_i | U_i(a_i, \alpha_{-i}^*) \geqslant U(a_i', \alpha_{-i}^*) \textit{ for all } a_i' \in A_i.\}$$

Note that this implies that in a Nash equilibrium in mixed actions each player is indifferent between all pure actions to which the player's mixed action assigns positive probability.

But: note that for a list of mixed actions to be a Nash equilibrium each player has to play the particular mixed action assigned to them. They must not switch to one of the other mixed actions that are just as good. Consider Matching Pennies. If player 1 chooses H with probability 2/3 and T with probability 1/3, whereas player 2 chooses both actions with probability 1/2, we no longer have a Nash equilibrium, even though player 1's expected utility has not change. The reason is that player 2 is no longer indifferent between H and T, but strictly prefers T, and therefore player 2's best reply is to choose T with probability 1. Thus, player 1 has to choose each of the two actions with probability 1/2 not because that is in his own interest, but because that is needed to keep player 2 indifferent between H and T.

I'll give one more example of a Nash equilibrium in mixed actions in this section. Recall from the previous Topic the coordination game:

	L	R
T	1,1	0,0
B	0,0	2,2

This game has two Nash equilibria in pure actions: (T, L) and (B, R). Once we allow mixed actions, there is another Nash equilibrium: player 1 chooses T with probability $2/3$, and B with probability $1/3$. Player 2 chooses L with probability $2/3$, and R with probability $1/3$. Each player has expected utility $2/3$, less than in either of the two Nash equilibria in pure actions. Note that each player chooses the "less attractive" action (T for player 1, L for player 2) with higher probability than the "more attractive" action. The reason players choose these mixtures is not their own interest, but this is needed to keep the other player indifferent between her two actions.

3. Existence of Nash Equilibria in Mixed Actions

Unlike in Topic 1, we will state only one result on the existence of Nash equilibria in mixed actions (and no result on the uniqueness of such equilibria).

Proposition 2. *Every game in which each player's action set A_i is non-empty and finite has at least one Nash equilibrium in mixed actions.*

This proof follows directly from the first of the two existence results that we stated in Topic 1. To see this, identify the mixed action set $\Delta(A_i)$ of each player with their action set A_i in Proposition 1 of Topic 1. Obviously, $\Delta(A_i)$ is non-empty, compact, and convex. Next, we identify the expected utility function U_i introduced in the previous section with the utility function u_i of player i in Proposition 1. This function is obviously continuous. Moreover, note that it is linear in a player's own mixed action. Therefore, it is quasi-concave.

4. How do Nash Equilibria in Mixed Actions Change When the Game Changes?

Let's consider one example to illustrate how Nash equilibria in mixed actions change when payoffs change. Consider the following game:

	L	R
T	1,1	0,0
B	0,0	1+a,1+a

and assume $a \geqslant 0$. In addition to the two pure action Nash equilibria (T, L) and (B, R), this game has a Nash equilibrium in mixed actions. In this equilibrium firm 1 chooses T with probability p where p satisfies:

$$p \cdot 1 = (1 - p) \cdot (1 + a) \Leftrightarrow p = \frac{1 + a}{2 + a}.$$

Firm 2 chooses L with the same probability. Note that this probability increases as a increases. Thus, as B or R apparently become more attractive, they are chosen in equilibrium with lower probabilities. This is because the probability with which B or R is chosen is pinned down by the requirement that the other player is indifferent between her/his two strategies.

5. Further Examples

Example 1

	L	C	R
T	0,3	0,0	3,0
M	0,0	2,2	0,0
B	3,0	0,0	0,3

Let's write $p(T), p(M), p(B)$ for player 1's mixed action and $q(L), q(C), q(R)$ for player 2's mixed action.

Is there a Nash equilibrium in pure strategies? You may check that the only such equilibrium is: $p(M) = q(C) = 1$.

Is there a Nash equilibrium in which player 1 chooses only two out of her three strategies with positive probability? If player 1 chooses T and M with positive probability, then player 2 will never choose R. For player 1 to play T and M with positive probability, she needs to be indifferent between T and M, and, if player 2 never chooses R, this only happens if player 2 chooses L with probability 1. But then the best response of player 1 is B, not T or M Therefore, there is no such Nash equilibrium.

Is there a Nash equilibrium in which player 1 chooses M and B, both with positive probability, but not T? Then player 2 will never choose L with positive probability. For player 1 to be indifferent between M and B, we then need player 2 to choose R with probability 1. But then the best response of player 1 is T, not M or B. Therefore, there is no such Nash equilibrium.

Is there a Nash equilibrium in which player 1 chooses T and B, both with positive probability? Then player 2 will not choose C with positive probability. To make player 1 indifferent between M, player 2 has to chooses L and T with probability 0.5: $q(L) = q(R) = 0.5>$ To make player 2 indifferent between L and R, player 1 has to choose T and B with probability 0.5: $p(T) = p(B) = 0.5$. This is indeed a Nash equilibrium. No player can gain by deviating to the action that they play with probability 0.

Is there a Nash equilibrium in which player 1 chooses T, M and B with positive probability? To make player 1 indifferent between T, M and B, we have to have: $3q(R) = 2q(C) = 3q(L)$. Combining this with the equation: $q(L) + q(C) + q(R) = 1$, we get: $q(L) = 2/7$, $q(C) = 3/7$, $q(R) = 2/7$. Because now player 2 now chooses all actions with positive probability, a symmetric argument implies that: $p(T) = 2/7$, $p(M) = 3/7$, $p(B) = 2/7$. Because both players are indifferent between all their actions, we have indeed found a third Nash equilibrium.

Example 2: A Price Setting Game

Two firms $i = 1, 2$ sell identical products. Each firm has a stock of $\bar{q}$ units. Each firm chooses the price $p_i \geqslant 0$ at which it wants to offer its stock. Demand is given by:

$$d(p) = \begin{cases} 1 & \text{if } 0 \leqslant p \leqslant 1; \\ \\ 0 & \text{if } p > 1. \end{cases}$$

Firm i's sales are $s_i(p_1,p_2)=0$ if $p_i>1$. Otherwise, the sales are given by:

$$s_i(p_1,p_2)=\begin{cases}\bar{q} & \text{if } p_i<p_j;\\ \\ 1/2 & \text{if } p_i=p_j;\\ \\ 1-\bar{q} & \text{if } p_i>p_j.\end{cases}$$

where $j\neq i$. Firms seek to maximize revenues: $p_i s_i(p_1,p_2)$. Thus we assume that production costs are zero.

We shall assume:

$$\bar{q}<1<2\bar{q},$$

so that each firm individually does not have enough capacity to serve all customers, but together both firms have enough capacity to serve all customers.

We begin with the following observation:

Claim 1: This game has no Nash equilibrium in pure strategies.

Proof: Suppose both firms chose the same price: $p_1=p_2=p$. If $p>1$, then each firm can gain by deviating and choosing a price below 1. If $0<p\leqslant 1$, then each firm can gain by deviating and choosing a price just below p, that is, $p_i=p-\varepsilon$ where $\varepsilon>0$ and very small. Finally, if $p=0$ each firm can gain by raising its price above zero but below 1.

Now suppose the firms did not choose the same price. Without loss of generality: $p_1<p_2$. If $p_2\leqslant 1$, then firm 1 can gain by raising p_1 to a higher price that is still less than p_2. If $p_2>1$ and $p_1>0$, then firm 2 can gain by lowering its price below 1.

Q.E.D.

We next study Nash equilibria in mixed strategies. We focus on symmetric mixed strategy Nash equilibria, that is, mixed strategy Nash equilibria in which both firms choose the same mixed strategy. A mixed strategy is now a probability distribution F on $\mathbb{R}_+$. We shall focus on equilibria in which F has a support that is equal to, or a subset of, $[0,1]$. We shall also focus on equilibria where F has a density f. Let f be such that

$$f(p)>0 \text{ if and only if } p\in[\underline{p},\bar{p}]$$

where $0 \leqslant \underline{p} < \bar{p} \leqslant 1$.

To find equilibrium distributions F we shall appeal to the characterization of mixed strategy Nash equilibria that we proved earlier, although earlier we only proved it for the case that the strategy sets are finite whereas in this example they are infinite. This is a gap in the proof that we shall not try to fill. We shall thus pretend we had shown that necessary and sufficient conditions for a symmetric Nash equilibrium were:

1. firms are indifferent between all prices in $[\underline{p}, \bar{p}]$.

2. firms do not strictly prefer prices outside of $[\underline{p}, \bar{p}]$.

We shall now seek to find distributions for which this is true. We begin by showing:

Claim 2a: We must have $\bar{p} = 1$.

Proof: If a firm chooses $\bar{p}$ it knows it will sell $1 - \bar{q}$. But, if $\bar{p} < 1$, then each firm would strictly prefer charing price 1 over charing price $\bar{p}$. The second of the necessary and sufficient conditions would be violated.

Q.E.D.

Claim 2b: We must have: $F(p) = \frac{\bar{q}p-(1-\bar{q})}{(2\bar{q}-1)p}$ for all $p \in [\underline{p}, 1]$.

Proof: The expected profit from setting price $p \in [\underline{p}, 1]$ is:

$$F(p)(1-\bar{q})p + (1-F(p))\bar{q}p.$$

The first of the necessary and sufficient conditions says that this has to equal for all p the expected profit from setting price $p = 1$. This profit is: $1 - \bar{q}$. Solving for $F(p)$ yields the claim.

Q.E.D.

Claim 2c: We must have: $\underline{p} = \frac{1-\bar{q}}{\bar{q}}$.

Proof: We can determine $\bar{p}$ because we must have: $F(\underline{p}) = 0$, and because we already have a formula for $F(p)$. If we plug in $\bar{p}$ and solve for $\bar{p}$, we get the formula in the claim.

Q.E.D.

We can now summarize what we have found in:

Claim 3: There is a unique symmetric mixed strategy Nash equilibrium. It satisfies:

$$F(p) = \frac{\bar{q}p - (1 - \bar{q})}{(2\bar{q} - 1)p}$$

for all $p \in \left[\frac{1-\bar{q}}{\bar{q}}, 1\right]$.

Proof: Among the necessary and sufficient conditions that we need to check the only one that we have not yet verified is that no firm has an incentive to choose a price below $\underline{p}$. But choosing such a price would yield the same quantity as choosing $\underline{p}$, but at a lower price, and would therefore lower sales.

Q.E.D.

The following graph shows the density of the equilibrium distribution for different values of $\bar{q}$.

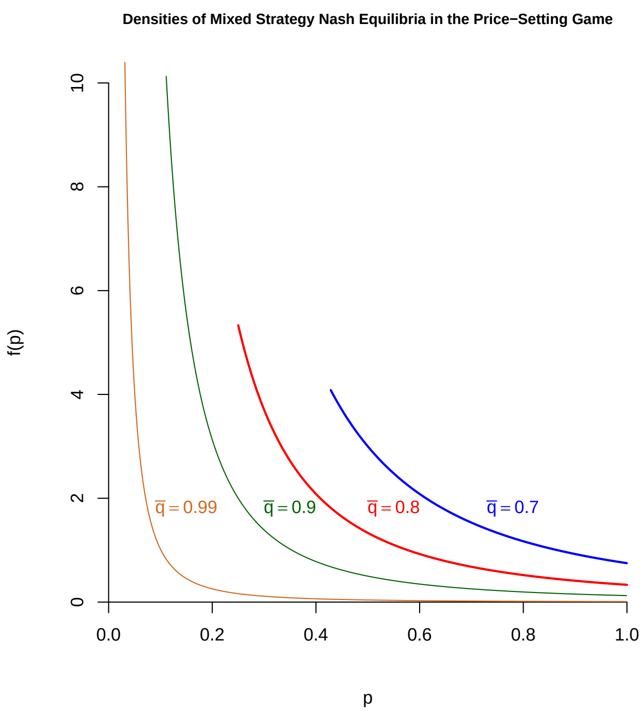

Note that as $\bar{q} \to 1$ for every p we have that $F(p)$ converges to 1. Thus, the distribution of prices converges to the point mass at price zero.

Topic 3: Rationalizability

This topic is about an alternative approach to predicting how players play a game, using the concept of "rationalizability." Whereas Nash equilibrium is best interpreted as a "stable convention" for playing a game, rationalizability seeks to describe the behavior of a player who is rational, whose beliefs attach probability 1 to the true game and also probability 1 to the event that the other players are also rational, and indeed the game, and the rationality of all players, are common probability 1 beliefs, in the sense that every player has beliefs that attach probability 1 to the true game and the rationality of all players, and every player believes with probability 1 that all other players' beliefs attach probability 1 to the true game and the rationality of all players, etc.

1. Rationality

Let's start by thinking about what rationality in a game might mean. We shall adopt a Bayesian definition of rationally. A player is rational when the player has a belief about the other players' actions (that is, a probability measure on A_{-i}), and chooses an action that maximizes his or her expected utility given this belief. It is convenient for this section, in fact, to attribute the property of rationality not to players, but to their actions.

Because we are considering probability measures on the action sets, let's restrict attention to games where all action sets are finite, as we did in the previous Topic. This allows us to bypass any measure-theoretic issues.

Definition 1. *An action a_i of player i is "rational" if there exists a probability measure μ_i on A_{-i} such that a_i satisfies:*

$$\sum_{a_{-i} \in A_{-i}} u_i(a_i, a_{-i})\mu_i(a_{-i}) \geqslant \sum_{a_{-i} \in A_{-i}} u_i(a_i', a_{-i})\mu_i(a_{-i}) \quad \textit{for all } a_i' \in A_i.$$

Here, μ_i represents player i's beliefs about the other players' behavior. Note that our definition of rationality allows player i to believe whatever he/she wants about other players' behavior. Therefore, there will often be many "rational" actions of player i.

Consider again the Prisoner's Dilemma:

	C	D
C	3,3	0,4
D	4,0	1,1

Note that action C of player 1 is not rational. This is because player 1's action D yields higher payoffs regardless of what player 2 does. We say that action D "strictly dominates" action C. Clearly, a strictly dominated action cannot be rational. Thus, there cannot be any belief of player 1 such that C maximizes expected utility when player 1 holds that belief. What about the converse: Is every action that is not strictly dominated rational?

The answer to this is "yes," but only if we are very careful about how we define "strict dominance." Consider the game below, where I only show player 1's utility.

	L	R
T	3	0
M	0	3
B	1	1

Let us find the rational actions of player 1. Let player 1's beliefs be that μ is the probability of L and that $1-\mu$ is the probability of R, where $\mu \in [0,1]$. Consider first the choice of player 1 assuming that player 1 only considers T and M. The expected utility from T is then 3μ and the expected utility from M is: $3(1-\mu)$. Note that $\max\{3\mu, 3(1-\mu)\}$ is minimized when $\mu = 0.5$. In that case, both actions give expected utility 1.5. If $\mu \neq 0.5$, either T or M gives expected utility higher than 1.5. Note that this implies that B is never optimal. In other words, B is not rational. T and M, however, are both rational.

Observe that B is not strictly dominated by any pure action. But it is strictly dominated by the mixed action that places probability 0.5 on T and M. This mixed action guarantees expected utility 1.5, regardless of what player 2 does. What this example therefore illustrates is that, when talking about strict dominance, we not only have to consider strict dominance by pure actions, but also by mixed actions. We have the following definition:

Definition 2. *A mixed action α_i of player i strictly dominates a pure action a_i of player i if:*

$$U_i(\alpha_i, a_{-i}) > u_i(a_i, a_{-i}) \quad \textit{for all } a_{-i} \in A_{-i}.$$

With this definition, we have:

Proposition 1. *An action a_i of player i is rational if and only if it is not strictly dominated by a mixed action.*

The proof of Proposition 1 is worth studying. It uses the separating hyperplane theorem. I provide it in the appendix to this topic.

The idea of Proposition 1 is that it is often much easier to figure out which actions are strictly dominated than to figure out which actions are rational.

2. Rationalizability

Players who are rational don't play strictly dominated actions. If every player believes with probability 1 that the other players are rational, then their beliefs cannot attach any positive probability to the strictly dominated actions of the other player. They must choose an action that is rational in the reduced game that one obtains when one just drops the strictly dominated actions of the other player. By Proposition 1, they have to choose in this reduced game a action that is not strictly dominated.

Iterating this argument we obtain that, if the game and rationality are commonly believed with probability 1, then players can only choose actions that survive iterated elimination of strictly dominated actions.

Definition 3. *An action a_i of player i is "rationalizable" if it survives the process of iterated elimination of strictly dominated actions.*

This is the solution concept that we introduce in this Topic. It is appropriate when players know the game, and believe with probability 1 that the other players are rational, and in fact these statements are common probability 1 belief.

Let us consider an example:

	L	C	R
T	0,1	2,3	3,1
M	1,3	0,1	4,2
B	0,0	1,0	3,3

For player 1 B is strictly dominated by a mixed action that places positive probability on both T and M, but that places more probability on T than on M. After we eliminate B for player 1, R for player 2 is strictly dominated by a action that places strictly positive probability on both L and C, but more probability on L than on C. In the remaining game, no action is strictly dominated. Thus, T and M are rationalizable for player 1, and L and C are rationalizable for player 2.

A few remarks:

- In the process that we have described we eliminate in each step all strictly dominated actions. You might alternatively try out eliminating in each step only one strictly dominated action of some player. These two processes of elimination lead to the same conclusion, as one can show. Indeed any process in which you eliminate at each stage only strictly dominated actions leads to the same conclusion, as long as you don't stop when there are still some strictly dominated actions.

- Do rationalizable actions exist? For the class of games in which all action sets are finite, which is the class that we are considering here, the answer is trivial: Yes, of course they exist. This is because it is not possible that at some stage you eliminate *all* actions of a player. That does not make sense.

- The problem with rationalizability is typically that there are lots of rationalizable actions. The concept often does not have much predictive power. (This claim deserves to be considered in much more detail. Unfortunately, we cannot do that here.)

How do rationalizability and Nash equilibrium relate to each other?

Proposition 2. *Suppose* $\alpha^* = (\alpha_1^*, \alpha_2^*, \ldots, \alpha_n^*)$ *is a Nash equilibrium, let* $i \in I$*, and suppose* $a_i \in A_i$ *satisfies:* $\alpha_i^*(a_i) > 0$*. Then* a_i *is rationalizable.*

Proof. Because a_i is in the support of α_i^*, it is a best response to α_{-i}^*. This means it is a best response to a belief about the other players' choices, and therefore it survives the first step of iterated elimination of strictly dominated actions. Indeed, all actions that are in the support of any players' equilibrium actions survives the first step. We obtain a reduced game, and in this reduced game α^* is still a Nash equilibrium. Iterating the argument, we find that actions that are in the support of a mixed action Nash equilibrium survive all steps of iterated elimination of strictly dominated actions. This is what Proposition 2 claims. □

I state the next result in a slightly informal language.

Proposition 3. *Suppose for every player i the set A_i^* is the set of rationalizable actions, and consider the reduced game in which players' action sets A_i are replaced by A_i^*, but their utility functions are not changed. The set of Nash equilibria of this reduced game is the same as the set of Nash equilibria of the original game.*

Proof. I first show that every Nash equilibrium of the original game is also a Nash equilibrium of the reduced game. By Proposition 3 every strategy available in the original game is also available in the reduced game. Thus, in particular, the Nash equilibrium strategies are still available. Moreover, because in the original game no deviation was profitable, and the reduced game has fewer strategies than the original game, no deviation is profitable in the reduced game, and we have a Nash equilibrium of the reduced game.

Now consider a Nash equilibrium of the reduced game. I want to show that it is also a Nash equilibrium of the original game. For this I have to show that it is not profitable to deviate to a strategy that is available in the original game but not in the reduced game. But because such strategies were eliminated during the process of iterated elimination of strictly dominated strategies, they are strictly dominated by a strategy in the reduced game. (This claim needs a proof, but it is a very simple proof, and therefore I omit it.) And thus, if a deviation in the original game were profitable, there would also be a profitable deviation in the reduced game, contradicting the assumption that we are considering a Nash equilibrium of the reduced game. □

3. Other Concepts of Rationalizability and Dominance

An alternative to the concept of strict dominance is the concept of weak dominance:

Definition 4. *A mixed action α_i of player i weakly dominates a pure action a_i of player i if:*

$$U_i(\alpha_i, a_{-i}) \geqslant u_i(a_i, a_{-i}) \quad \text{for all } a_{-i} \in A_{-i}$$

and

$$U_i(\alpha_i, a_{-i}) > u_i(a_i, a_{-i}) \quad \text{for at least one } a_{-i} \in A_{-i}.$$

We might ask whether this concept has a similar characterization in terms of rationality as strict dominance has, according to Proposition 1. This is indeed the case, as we show now.

Definition 5. *An action a_i of player i is "cautiously rational" if there exists a probability measure μ_i on A_{-i} which has full support, i.e. $\mu_i(a_{-i}) > 0$ for every $a_{-i} \in A_{-i}$, and such that a_i satisfies:*

$$\sum_{a_{-i} \in A_{-i}} u_i(a_i, a_{-i}) \mu_i(a_{-i}) \geqslant \sum_{a_{-i} \in A_{-i}} u_i(a_i', a_{-i}) \mu_i(a_{-i}) \quad \text{for all } a_i' \in A_i.$$

Proposition 4. *An action a_i of player i is cautiously rational if and only if it is not weakly dominated by a mixed action.*

We omit the proof of Proposition 3. The proof is also an application of a separating hyperplane theorem.

We might also consider the iterated deletion of weakly dominated actions. This procedure is known to be problematic. The following example shows that the order in which weakly dominated actions are eliminated might matter for the conclusion:

	L	R
T	2,2	1,2
B	2,1	0,0

If we eliminate B first, then L and R remain for player 2. If we eliminate R first, then T and B remain for player 1. Or we could eliminate R and B at the same time. In each case we arrive at different predictions about which actions players might choose and about which utilities they might get. A deeper problem with the iterated deletion of weakly dominated actions is that it is difficult to derive the concept from primitive assumptions about what players believe about each other, what they believe about others' beliefs, etc.

Iterated deletion of weakly dominated actions might eliminate some Nash equilibria. In the previous example, (T, R) is a Nash equilibrium, but it might be eliminated by iterated elimination of weakly dominated actions, depending on the order in which weakly dominated actions are eliminated. Another example is this:

	L	R
T	1,1	0,0
B	0,0	0,0

(B, R) is a Nash equilibrium, but iterated elimination of weakly dominated actions will always eliminate it.

Another concept of rationality requires actions to be best responses to beliefs that are product measures, reflecting the independence of other players' actions. In games with more than two players, this may be more restrictive than the concept of rationality that we introduced in Definition 1. This concept of rationality then has no characterization in terms of dominance. When the term "rationalizability" was first introduced, it referred to the iterated elimination of actions that are not best responses to product measure beliefs. Today, the meaning seems to shift, and rationalizability is simply identified with actions that survive iterated elimination of strictly dominated actions.

The following definitions are also important:

Definition 6. *An action a_i of player i is called strictly dominant if for all $a'_i \in A_i$ with $a'_i \neq a_i$:*

$$u_i(a_i, a_{-i}) > u_i(a'_i, a_{-i}) \quad \textit{for all } a_{-i} \in A_{-i}.$$

Definition 7. *An action a_i of player i is called weakly dominant if for all $a_i' \in A_i$ with $a_i' \neq a_i$:*

$$u_i(a_i, a_{-i}) \geqslant u_i(a_i', a_{-i}) \quad \textit{for all } a_{-i} \in A_{-i}.$$

and

$$u_i(a_i, a_{-i}) > u_i(a_i', a_{-i}) \quad \textit{for at least one } a_{-i} \in A_{-i}.$$

These concepts have the following characterizations:

Proposition 5. *An action a_i is strictly dominant if and only if it is the only rational action of player i.*

Proposition 6. *An action a_i is weakly dominant if and only if it is the only cautiously rational action of player i.*

The proof of these propositions is short and simple, and, if you want, you can try to prove them by yourself. I omit the proofs here.

4. Rationalizability and Best Responses

Whereas up to this point we have assumed that all action sets are finite, we now also consider games with infinite action sets. We shall only use the concepts of strict dominance, and iterated deletion of strictly dominated actions, and we shall only consider the case that the dominating actions are pure actions. Thus, we can avoid all measure theoretic issues that arise when action sets are infinite.

Recall from Proposition 1 that a strategy that is a best response is not strictly dominated. Consider now games in which the action sets A_i are compact subsets of $\mathbb{R}$ and the utility functions u_i are continuous. In such games we can define for each player i the largest best response of player i, $\bar{a}_i$, and the smallest best response of player i, $\underline{a}_i$.

$$\bar{a}_i = \max \bigcup_{a \in A} BR_i(a) \quad \text{and} \quad \underline{a}_i = \min \bigcup_{a \in A} BR_i(a).$$

The following result uses the same assumptions as Proposition 2 in Topic 1.

Proposition 7. *If for every player $i \in I$:*

1. *the set A_i is a non-empty and compact subset of $\mathbb{R}$;*
2. *the function u_i is continuous in a;*
3. *the function u_i has increasing differences, i.e. if: $i \in I$, $a_i, a_i' \in A_i$, $a_i > a_i'$, for every $j \in I$ with $j \neq i$: $a_j, a_j' \in A_j$ and $a_j \geqslant a_j'$ then:*

$$u_i(a_i, a_{-i}) - u_i(a_i', a_{-i}) \geqslant u_i(a_i, a_{-i}') - u_i(a_i', a_{-i}'),$$

then for every player i every action $a_i > \bar{a}_i$ is strictly dominated by $\bar{a}_i$, and every action $a_i < \underline{a}_i$ is strictly dominated by $\underline{a}_i$.

Proof. First, we show that $\bar{a}_i$ is a best response to the vector of maximal actions of all other players: $(\max A_j)_{j \neq i}$. The proof is indirect: suppose there were some action $a_i < \bar{a}_i$ such that:

$$u_i\left(\bar{a}_i, (\max A_j)_{j \neq i}\right) - u_i\left(a_i, (\max A_j)_{j \neq i}\right) < 0.$$

The strictly increasing differences assumption means that then a_i yields higher utility not just when all other players choose their largest actions, but also when they choose any other actions. This is because the strictly increasing differences assumption implies for all $a_{-i} \in A_{-i}$:

$$u_i\left(\bar{a}_i, a_{-i}\right) - u_i\left(a_i, a_{-i}\right) \leqslant u_i\left(\bar{a}_i, (\max A_j)_{j \neq i}\right) - u_i\left(a_i, (\max A_j)_{j \neq i}\right),$$

and by assumption the right hand side of that inequality is negative. But then $\bar{a}_i$ is not a best response, which contradicts the assumption that it is the largest best response.

A similar argument shows that $\underline{a}_i$ is a best response to the vector of minimal actions of all other players: $(\min A_j)_{j \neq i}$.

Consider now any action $a_i > \bar{a}_i$. We now prove the assertion of the Proposition, that is, we prove that $\bar{a}_i$ strictly dominates a_i. Because a_i is not a best response, we have:

$$u_i\left(\bar{a}_i, (\max A_j)_{j \neq i}\right) - u_i\left(a_i, (\max A_j)_{j \neq i}\right) > 0.$$

The assumption of strictly increasing differences then implies that for all $a_{-i} \in A_{-i}$:

$$u_i\left(\bar{a}_i, a_{-i}\right) - u_i\left(a_i, a_{-i}\right) > 0.$$

But this says that $\bar{a}_i$ strictly dominates a_i.

A similar argument shows that $\underline{a}_i$ strictly dominates any $a_i < \underline{a}_i$. □

We are going to apply this result to study rationalizability in a Cournot duopoly. Let us assume that there are only two firms: $i = 1, 2$, that the firms' action sets are $A_i = [0, 10]$, and that inverse demand is given by: $p = 10 - (a_1 + a_2)$. We assume costs are zero, and therefore each firm's profit is:

$$u_i(a_1, a_2) = (10 - a_1 - a_2)a_i.$$

Recall that this game has strictly increasing differences if we revert the ordering of the real numbers for one of the two firms, say firm 2. This implies that we can apply our result to this game.

Each firm i's best response function can be found by maximizing the (quadratic) profit function. The result is:

$$BR_i(a) = \frac{10 - a_j}{2}.$$

Therefore, each firm's smallest best response is 0, and each firm's largest best response is 5. All actions larger than 5 are strictly dominated.

Now consider the smaller game in which each firm's action set is $[0, 5]$. In this game, each firm's smallest best response is 2.5, and each firm's largest best response is 5. Thus, all actions below 2.5 are strictly dominated.

Now consider the even smaller game in which each firm's action set is: $[2.5, 5]$. Each firm's smallest best response in this game is 2.5, and each firm's largest best response is: 3.75. Therefore, all actions above 3.75 are strictly dominated.

Note that, as we continue this process, we alternate between adjusting the upper bound, and adjusting the lower bound, of the remaining actions. Let's keep track of the upper bounds. The upper bound after the first round of elimination is 5. We then adjust the lower bound, to $\frac{10-5}{2}$, and then adjust the upper bound to $\frac{10-\frac{10-5}{2}}{2}$. In general, in each odd step n, the upper bound is given by:

$$\bar{a}_n = \frac{10 - \frac{10-\bar{a}_{n-2}}{2}}{2}.$$

This simplifies to:

$$\bar{a}_n = 2.5 + \frac{1}{4}\bar{a}_{n-2}.$$

This difference equation has the fixed point:

$$\bar{a}^* = 2.5 + \frac{1}{4}\bar{a}^* \Leftrightarrow \bar{a}^* = \frac{10}{3}.$$

Our initial value $a_1 = 5$ is larger than the fixed point. The coefficient in front of the lagged term in the difference equation is in absolute value less than 1. Therefore, the solution to the difference equation converges to the fixed point. Therefore, as the number of iterations tends to infinity, the upper bound of the surviving actions tends to $\frac{10}{3}$.

Let's consider the sequence of lower bounds. It starts with $\underline{a}_2 = 2.5$, and it continues according to the same difference equation as the upper bound:

$$\underline{a}_n = 2.5 + \frac{1}{4}\underline{a}_{n-2}.$$

By an argument analogous to the argument we used in the previous paragraph, the sequence of lower bounds also converges to $\frac{10}{3}$.

We have concluded that, if we apply iterated deletion of strictly dominated actions to this Cournot example, as the number of steps tends to infinity, for each firm the only surviving quantity is $\frac{10}{3}$. Note that this is also the Cournot equilibrium quantity.

What we have seen in this example generalizes to a large class of Cournot duopoly games.

Appendix for Topic 3

Proof of Proposition 1

STEP 1: We prove the "only if" part, that is, we assume that a_i^* is a best reply to a belief $\mu_i \in \Delta(A_{-i})$, and infer that a_i^* is not strictly dominated. The proof is indirect. Suppose a_i^* were strictly dominated by $\alpha_i \in \Delta(A_i)$. Then, obviously, α_i yields strictly higher expected utility given the belief μ_i than a_i^*:

$$\sum_{a_{-i} \in A_{-i}} U_i(\alpha_i, a_{-i})\mu_i(a_{-i}) > \sum_{a_{-i} \in A_{-i}} u_i(a_i^*, a_{-i})\mu_i(a_{-i}). \tag{1}$$

We thus have that α_i is a better response to μ_i than a_i^*, which is almost what we want to obtain, but not quite. To obtain the desired contradiction, we want to find a *pure* action that is a better response to μ_i than a_i^*. This can be done as follows. We re-write the left hand side of (1) as follows:

$$\begin{aligned}
\sum_{a_{-i} \in A_{-i}} U_i(\alpha_i, a_{-i})\mu_i(a_{-i}) &= \sum_{a_{-i} \in A_{-i}} \sum_{a_i \in A_i} \alpha_i(a_i) u_i(\alpha_i, a_{-i})\mu_i(a_{-i}) \\
&= \sum_{a_i \in A_i} \sum_{a_{-i} \in A_{-i}} \alpha_i(a_i) u_i(\alpha_i, a_{-i})\mu_i(a_{-i}) \\
&= \sum_{a_i \in A_i} \alpha_i(a_i) \left(\sum_{a_{-i} \in A_{-i}} u_i(a_i, a_{-i})\mu_i(a_{-i}) \right). \qquad (2)
\end{aligned}$$

Combining (1) and (2) we have:

$$\sum_{a_i \in A_i} \alpha_i(a_i) \left(\sum_{a_{-i} \in A_{-i}} u_i(a_i, a_{-i})\mu_i(a_{-i}) \right) > \sum_{a_{-i} \in A_{-i}} u_i(a_i^*, a_{-i})\mu_i(a_{-i}). \tag{3}$$

The left hand side of (3) is a convex combination of the expressions in large brackets in that term. This convex combination can be larger than the right hand side of (3) only if one of the expressions in large brackets is larger than the right hand side of (3), i.e., for some $a_i \in A_i$:

$$\sum_{a_{-i} \in A_{-i}} u_i(a_i, a_{-i})\mu_i(a_{-i}) > \sum_{a_{-i} \in A_{-i}} u_i(a_i^*, a_{-i})\mu_i(a_{-i}). \tag{4}$$

and thus a_i is a better response to μ_i than a_i^*, which contradicts the assumption a_i^* is a best response to μ_i among all pure actions.

Step 2:[1] We prove the "if" part, that is, we assume that a_i^* is not strictly dominated, and we show that there is a belief $\mu_i \in \Delta(A_{-i})$ to which a_i^* is a best response. The proof is constructive. We define two subsets, X and Y, of the set $\mathbb{R}^{|A_{-i}|}$, that is, the Euclidean space with dimension equal to the number of elements of A_{-i}. We shall think of the elements of these sets as payoff vectors. Each component indicates a payoff that player i receives when the other players choose some particular $a_{-i} \in A_{-i}$.

Now pick any one-to-one mapping $f : A_{-i} \longrightarrow \{1, 2, \ldots, |A_{-i}|\}$. For any action a_i of player i, we write $u_i(a_i, \langle a_{-i}\rangle) \in \mathbb{R}^{|A_{-i}|}$ for the vector of payoffs that player i receives when playing a_i, and when the other players play their various action combinations a_{-i}. Specifically, the k-th entry of $u_i(a_i, \langle a_{-i}\rangle)$ is the payoff $u_i(a_i, f^{-1}(k))$ where f^{-1} is the inverse of f. Intuitively, f defines an order in which we enumerate the elements of A_{-i}, and $u_i(a_i, \langle a_{-i}\rangle)$ lists the payoffs of player i when he plays a_i and the other players play a_{-i} in the order defined by f.

The set X is:

$$X = \{x \in \mathbb{R}^{|A_{-i}|} | x > u_i(a_i^*, \langle a_{-i}\rangle)\}. \tag{5}$$

Here, "$>$" is to be interpreted as: "strictly greater in every component." Therefore, the set X is the set of payoff vectors that are strictly greater in every component than $u_i(a_i^*, \langle a_{-i}\rangle)$, that is the payoff vector that corresponds to the undominated action a_i^*.

The set Y is:

$$Y = co\{y \in \mathbb{R}^{|A_{-i}|} | \exists a_i \in A_i : y = u_i(a_i, \langle a_{-i}\rangle)\}. \tag{6}$$

Here, "*co*" stands for "convex hull." The payoff vectors in Y are the payoff vectors that player i can achieve by choosing some mixed action. The weight that the convex combination that defines an element of y places on each element of $\{x \in \mathbb{R}^{|A_{-i}|} | \exists a_i \in A_i : x = u_i(a_i, \langle a_{-i}\rangle)\}$ corresponds to the probability which the mixed action assigns to each pure action $a_i \in A_i$.

It is obvious that both sets are nonempty and convex sets. Moreover, their intersection is empty. If X and Y overlapped, then every common element would correspond to the payoffs arising from a mixed action of player i that strictly dominates a_i^*. Because by assumption no such mixed action exists, X and Y cannot have any elements in common.

[1] An example and a graph that illustrate Step 2 follow after the end of the proof.

In the previous paragraph we have checked all the assumptions of the separating hyperplane theorem: we have two nonempty and convex sets that have no elements in common. The separating hyperplane theorem (Theorem 1.68 in Sundaram [4]) then says that there exist some row vector $i \in \mathbb{R}^{|A_{-i}|}$ which is not equal to zero in every component, and some $\delta \in \mathbb{R}$, such that:

$$i \cdot x \geqslant \delta \quad \forall x \in X \tag{7}$$

and

$$i \cdot y \leqslant \delta \quad \forall y \in Y. \tag{8}$$

Here "·" stands for the scalar product of two vectors. We treat all vectors in X and Y as column vectors. Therefore, the above scalar products are well-defined.

We now make two observations. The first is:

$$\delta = i \cdot u_i(a_i^*, \langle a_{-i} \rangle). \tag{9}$$

To show this note that by definition $u_i(a_i^*, \langle a_{-i} \rangle) \in Y$, and therefore, by (8), $i \cdot u_i(a_i^*, \langle a_{-i} \rangle) \leqslant \delta$. Next, for every $n \in \mathbb{N}$ define $x_n = u_i(a_i^*, \langle a_{-i} \rangle) + \varepsilon^n \cdot \iota$, where $\varepsilon \in (0,1)$ is some constant and ι is the column vector in $\mathbb{R}^{|A_{-i}|}$ in which all entries are "1". Observe that for every $n \in \mathbb{N}$ we have $x_n \in X$, so that by (7) we have: $i \cdot x_n \geqslant \delta$ for every n. On the other hand, we have: $\lim_{n \to \infty} x_n = u_i(a_i^*, \langle a_{-i} \rangle)$. By the continuity of the scalar product of vectors therefore: $i \cdot u_i(a_i^*, \langle a_{-i} \rangle) \geqslant \delta$. This, combined with our earlier observation $i \cdot u_i(a_i^*, \langle a_{-i} \rangle) \leqslant \delta$ implies (9).

Our second observation is:

$$i \geqslant 0 \tag{10}$$

where we interpret "$\geqslant$" to mean "greater or equal in every component," but not identical, and 0 stands for the vector consisting of zeros only. That i is not equal to 0 is part of the assertion of the separating hyperplane theorem. We shall prove indirectly that no component of i can be negative. Assume that i has a negative component. Without loss of generality, assume that it is the first. Now we consider the vector $x = u_i(a_i^*, \langle a_{-i} \rangle) + (1, \varepsilon, \varepsilon, \ldots, \varepsilon)$ where $\varepsilon > 0$ is some number. Clearly, the vector x that we define like this is contained in X. Moreover,

$$i \cdot x = i \cdot u_i(a_i^*, \langle a_{-i} \rangle) + i \cdot (1, \varepsilon, \varepsilon, \ldots, \varepsilon) = \delta + i \cdot (1, \varepsilon, \varepsilon, \ldots, \varepsilon) \tag{11}$$

where for the second equality we have used (9). We now want to evaluate $i \cdot (1, \varepsilon, \varepsilon, \ldots, \varepsilon)$. Unfortunately, we need some additional notation: we shall use for the k -th component of the vector i the symbol: i_k. Here, $k \in \{1, 2, \ldots, |A_{-i}|\}$. We then obtain:

$$i \cdot (1, \varepsilon, \varepsilon, \ldots, \varepsilon) = i_1 + \varepsilon \sum_{k=2}^{|A_{-i}|} i_k \tag{12}$$

Now observe that by the contrapositive assumption of our indirect proof $i_1 < 0$. Therefore, for small enough ε the right hand side of (12) is negative. Using this fact, we obtain from (11):

$$i \cdot x < \delta \tag{13}$$

which contradicts $x \in X$ and (7). This completes our indirect proof of (10).

Now we denote by $\|i\|$ the Euclidean norm of i. Because $i \neq 0$, if we define μ_i by:

$$\mu_i = \frac{1}{\|i\|} i, \tag{14}$$

then the Euclidean norm of μ_i is 1, and thus $\mu_i \in \Delta(A_{-i})$, that is, μ_i is a belief of player i. We complete the proof by showing that a_i^* is a best response to μ_i. (8) and (9) together imply:

$$i \cdot y \leqslant i \cdot u_i(a_i^*, \langle a_{-i} \rangle) \quad \forall y \in Y. \tag{15}$$

Dividing this inequality by $\|i\|$ we get:

$$\mu_i \cdot y \leqslant \mu_i \cdot u_i(a_i^*, \langle a_{-i} \rangle) \quad \forall y \in Y. \tag{16}$$

Now by definition of Y for every a_i: $u_i(a_i, \langle a_{-i} \rangle) \in Y$. Therefore:

$$\mu_i \cdot u_i(a_i, \langle a_{-i} \rangle) \leqslant \mu_i \cdot u_i(a_i^*, \langle a_{-i} \rangle) \quad \forall a_i \in A_i. \tag{17}$$

This is what we wanted to prove: a_i^* is a best response in A_i to μ_i.

We illustrate Step 2 of the proof of Proposition 1 with the following game in which player 1 chooses rows and player 2 chooses columns. Only the payoffs of player 1 are shown.

	L	R
T	2	2
MA	3	-2
MB	0	-1
MC	-1	0
B	-2	3

Example 1

Player 1's action T is not strictly dominated. We illustrate in Figure 1 the construction of beliefs to which action T is a best reply. The sets X and Y to which Step 2 of the proof of Proposition 1 refers are shown in the figure.

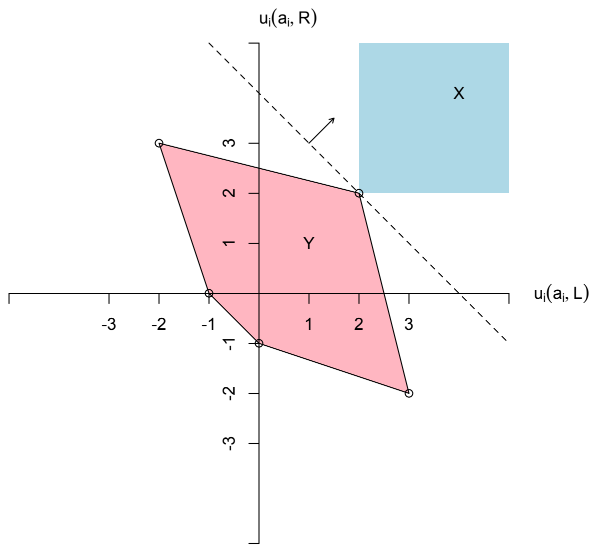

Figure 2

The hyperplane (straight line) separating X and Y is the dashed line in Figure 1. Figure 1 also shows the orthogonal vector for this hyperplane.

Topic 4: Multi-Stage Games

1. General Definition of Multi-Stage Games

Definition 8. *A multi-stage game is a list* $\left(I, (A_i)_{i\in I}, (\mathcal{A}_i^t)_{i\in I, t\in\{1,2,\ldots,T\}}, (u_i)_{i\in I}\right)$, *where:*

1. $I = \{1, 2, \ldots, n\}$ *is a finite set of players;*
2. $T \in \mathbb{N}$ *is the number of stages;*

and for every $i \in I$*:*

3. A_i *is the set of actions potentially available to player* i *at some stage of the game;*
4. $\mathcal{A}_i^1 \subseteq A_i$ *is the non-empty set of actions available to player* i *in stage* $t = 1$*;*
5. *for every* $t = 2, ..., T$*:*
$$\mathcal{A}_i^t : H^{t-1} \twoheadrightarrow A_i,$$
is a non-empty valued correspondence that describes the set of actions available to player i *in stage* t*;*
6. $u_i : H^T \to \mathbb{R}$ *is player* i*'s utility function.*

Here, for every $t = 1, 2, \ldots, T$*, we denote by* H^t *the set of histories that may occur up to stage* t*:*

$$H^1 = \bigtimes_{i=1}^{n} \mathcal{A}_i^1 \quad \text{and, for } t \geqslant 2: \quad H^t = \{(h^{t-1}, (a_1^t, a_2^t, \ldots, a_n^t)) | h^{t-1} \in H^{t-1} \text{ and } a_i^t \in \mathcal{A}_i^t(h^{t-1}) \text{ for all } i.\}$$

2. Examples of Multi-Stage Games

EXAMPLE 1: The players are two single-product firms: $i = 1, 2$. There are two stages. In the first stage each firm chooses whether to produce product X or product Y. In the second stage, each firm chooses a price p_i at which it offers its product. Firms have no cost, and therefore seek to maximize revenues.

There is a continuum of consumers. Consumers value one unit of product X at 1 Dollars. Consumers value one unit of product Y at v Dollars, where v is different for different consumers. Specifically, v is uniformly distributed on $[0, 2]$ with density $\frac{1}{2}$. Consumers buy from whatever firm offers the larger surplus: value - price. If consumers are indifferent between two firms, they pick one randomly. If consumers are indifferent between buying and not buying, they buy.

To avoid some trivial case distinctions, let us assume that a firm i that produces product X cannot charge more than 1 Dollar: $p_i \in [0, 1]$, and a firm i that produces product Y cannot charge more than 2 Dollars: $p_i \in [0, 2]$.

If both firms offer product X, then consumers will buy from the firm with the lower price only. The firms' profit functions are:

$$i_i(p_i, p_j) = \begin{cases} 0 & \text{if } p_i > p_j \\ \frac{1}{2}p_i & \text{if } p_i = p_j \\ p_i & \text{if } p_i < p_j \end{cases}$$

where $i \neq j$.

If firm i offers product X at price p_i whereas firm j offers product Y at price p_j, then a consumer will buy from firm i if: $1 - p_i > v - p_j \Leftrightarrow v < 1 + (p_j - p_i)$. If $p_i \leqslant p_j - 1$, then this means that all consumers buy from firm i. Otherwise, the mass of consumers buying from firm i is: $\frac{1}{2}(1 + (p_j - p_i))$. This means that firm i's profit function is:

$$i_i(p_i, p_j) = \begin{cases} \frac{1}{2}(1 + (p_j - p_i))p_i & \text{if } p_i > p_j - 1 \\ p_i & \text{if } p_i \leqslant p_j - 1 \end{cases}$$

Similarly, firm j's profit function is:

$$i_j(p_i, p_j) = \begin{cases} 0 & \text{if } p_j > p_i + 1 \\ \frac{1}{2}(1 + (p_i - p_j))p_j & \text{if } p_j \leqslant p_i + 1 \end{cases}$$

If both firms offer product Y, then consumers will buy from the firm with the lower price only. Firms' profit functions are:

$$i_i(p_i, p_j) = \begin{cases} 0 & \text{if } p_i > p_j \\ \frac{1}{4}(2 - p_i)p_i & \text{if } p_i = p_j \\ \frac{1}{2}(2 - p_i)p_i & \text{if } p_i < p_j \end{cases}$$

where $i \neq j$.

EXAMPLE 2: Three voters vote over three candidates: A, B, and C. There are three stages. In the first stage, voter 1 picks one of the three candidates. In the second stage, voter 2 picks one of the three candidates. In the third stage, voter 3 chooses among the two candidates picked by voters 1 and 2. Of course, if voters 1 and 2 picked the same candidate, voter 3 does not have any choice.

Voters' utilities are:

$$u_1(A) = 2 \text{ and } u_1(B) = 1 \text{ and } u_1(C) = 0.$$
$$u_2(C) = 2 \text{ and } u_2(A) = 1 \text{ and } u_2(B) = 0.$$
$$u_3(B) = 2 \text{ and } u_3(C) = 1 \text{ and } u_3(A) = 0.$$

Observe that in this example in each stage only one of the three players has a choice. One can fit this into the framework of general multi-stage games by assuming that all except one player have in each stage just one choice.

Games in which in each stage only one of the three players has a choice are called games of *perfect information.*

Games of perfect information can be represented by a *game tree.* In Figure 1 we show Example 2's game tree.

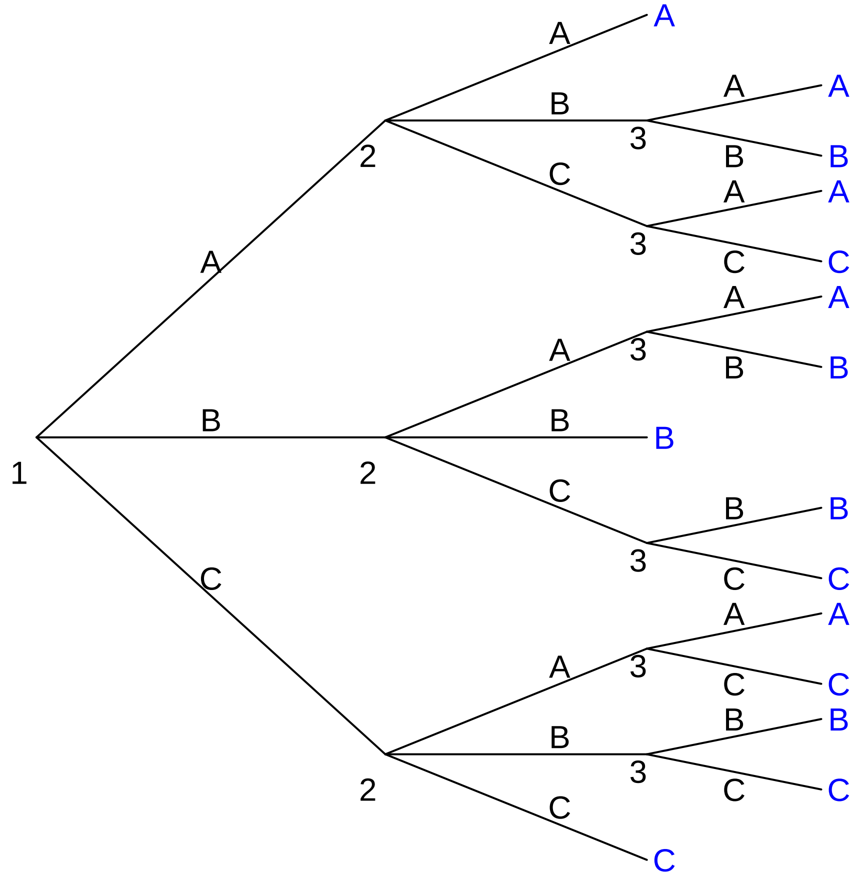
1
A
B
C
2
2
2
A
B
C
A
B
C
A
B
C
3
3
3
3
3
3
A
A
B
A
B
A
C
A
B
B
B
C
A
C
B
C
C
A
A
B
A
C
A
B
B
B
C
A
C
B
C
C

Figure 1

3. Backward Induction

We are interested in predicting how players will choose in multi-stage games. Building on the idea of Nash equilibrium, we shall develop an equilibrium concept for multi-stage games that we shall later call "subgame-perfect equilibrium." Rather than giving the full definition of this solution concept, we shall begin by describing the algorithm that is used in multi-stage games with a finite number of stages to find subgame-perfect equilibria. This algorithm is called "backward induction."

Backward induction

While $T > 1$:

Step 1: For every history $h^{T-1} \in H^{T-1}$ find a Nash equilibrium

$$a^*(h^{T-1}) = (a_1^*(h^{T-1}), a_2^*(h^{T-1}), \dots, a_n^*(h^{T-1}))$$

of the game with player set I, strategy sets $\mathcal{A}_i^T(h^{t-1})$ and utility functions: $v_i(a_1, a_2, \dots, a_n) = u_i(h^{T-1}, (a_1, a_2, \dots, a_n))$.

Step 2: Then consider the new multi-stage game with $T-1$ stages, where action correspondences are the same as in the first $T-1$ stages of the T-stage game, and where we define utility in this new game by:

$$u_i(h^{T-1}) = u_i(h^{T-1}, a^*(h^{T-1})).$$

Step 3: Set $T \leftarrow (T-1)$. Go back to Step 1.

When $T = 1$:

Step 4: Find a Nash equilibrium $a^* = (a_1^*, a_2^*, \dots, a_n^*)$ of the game with player set I, strategy sets $\mathcal{A}_i^1$, and utility functions: $u_i(a_1, a_2, \dots, a_n)$.

Backward Induction finds for each player, for every stage, for every possible history of the game up to that stage, an action.

Let us introduce some new terminology and notation:

Definition 9. *A (pure) strategy s_i of player i in a multi–stage game consists of:*

1. *an action $s_i^1 \in \mathcal{A}_i^1$ for the first stage;*
2. *a function: $s_i^t : H^{t-1} \to A_i$ for stages $t \geqslant 2$ such that:*

$$s_i^t(h^{t-1}) \in \mathcal{A}_i^t(h^{t-1}) \text{ for all } h^{t-1} \in H^{t-1}.$$

Let S_i denote the set of strategies of player i.

Backward Induction identifies for each player a strategy of that player

Definition 10. *Every strategy combination that Backward Induction finds is a subgame-perfect equilibrium.*

Note that we are considering an equilibrium notion: for every possible contingency, some stable convention on how to act in that contingency has developed that all players understand. In fact, Backward Induction is an algorithm for finding a subset of all Nash equilibria. Due to lack of time, we shall not discuss solution concepts related to rationalizability for multistage games. We shall discuss the interpretation of subgame-perfect equilibria in more detail in later sections.

More notes on the algorithm:

- Note that we focus on pure actions only. We shall disregard mixed actions for this entire topic.
- If at some stage, after some history, we don't find any Nash equilibrium, then there is no subgame-perfect equilibrium (in pure strategies).
- If at some stage, after some history, we find multiple Nash equilibria, then we choose one, and continue with the algorithm. We then have to repeat the algorithm for every Nash equilibrium that we could have picked.

4. Backward Induction in the Examples

Example 1: In the second stage, if both firms offer the same product, the only Nash equilibrium is: $p_1 = p_2 = 0$. This is the Bertrand equilibrium.

If firm i offers product X and firm j offers product Y, then firm i's best response will always satisfy:

$$p_i \geqslant p_j - 1$$

Therefore, we can maximize $\frac{1}{2}(1 + (p_j - p_i))p_i$ using the first order condition, and obtain:

$$p_i = \frac{1 + p_j}{2}.$$

Firm i cannot charge a price above 1. Its best response is therefore given by:

$$p_i = \begin{cases} \frac{1+p_j}{2} & \text{if } p_j \leqslant 1 \\ 1 & \text{if } p_j > 1. \end{cases}$$

Firm j's best response will always satisfy:

$$p_j \leqslant p_i + 1,$$

and therefore, we can maximize $\frac{1}{2}(1 + (p_i - p_j))p_j$ using the first order condition, and obtain:

$$p_j = \frac{1 + p_i}{2}.$$

This implies that there is a unique Nash equilibrium:

$$p_i = p_j = 1.$$

All consumers with $v < 1$ buy from firm 1. All consumers with $v > 1$ buy from firm 2. Both firms have profit 0.5.

We move now to the first stage. If firms anticipate the second stage equilibrium, the first stage has the following form:

	X	Y
X	0,0	0.5,0.5
Y	0.5,0.5	0,0

There are two Nash equilibria. One firm produces X, and the other firm produces Y. Firms avoid price competition through product differentiation.

Example 2: We can use the graphical representation of the game shown in Figure 1. In Figure 2 we have, starting at the end of the game tree, marked the optimal choices in red.

Note that voter 2 is indifferent at her second decision node between all three alternatives, because voter 1 offered B, and B is voter 3's most preferred alternative. Therefore, regardless of what voter 2 does, B will be the outcome.

If voter 1 chooses A, then voter 2 can put her most preferred outcome C into the mix, and make sure that C will be chosen.

If voter 1 chooses C, which is voter 2's preferred outcome, then voter 2 can either immediately choose C, or she can add a candidate to the mix of which she is sure that voter 3 will not choose it.

Thus voter 1 knows that, if she picks A or C, then voter 2 secures her most preferred outcome C. By contrast, if she picks B, then this will be the outcome, because it is voter 3's preferred outcome. Because voter 1 prefers B to C, she picks B.

The voter who moves third thus gets her most preferred outcome. One might ask whether this is true only for the particular preferences proposed here, or whether it is true more generally.

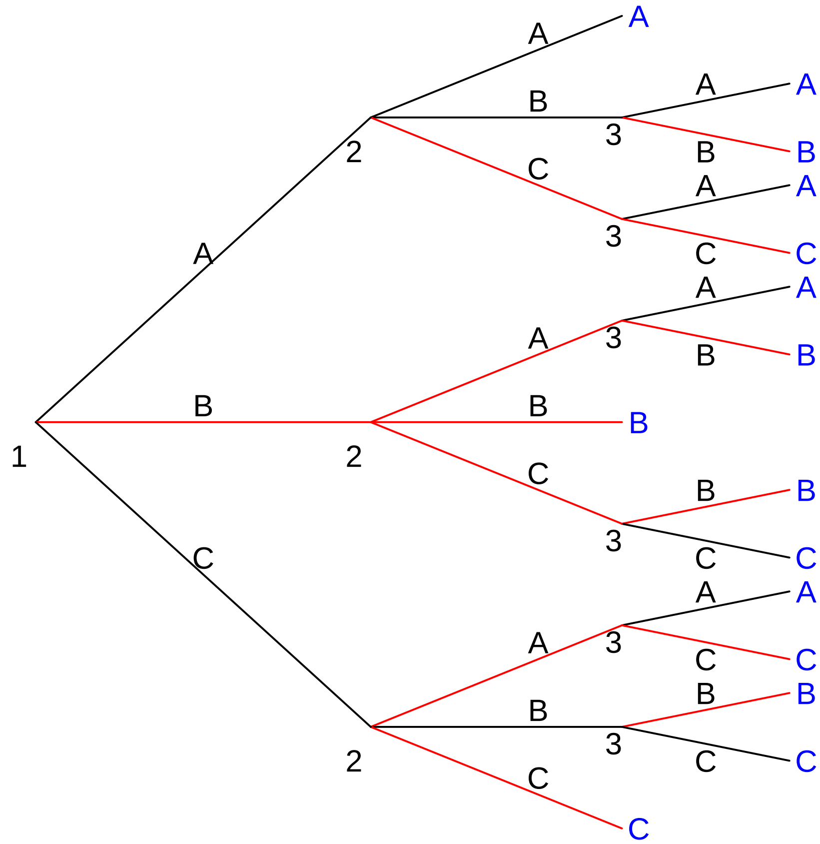

Figure 2

5. Subgame-Perfect Equilibria

In Definition 3 we defined "subgame-perfect equilibria" to be the strategy combinations that are identified by Backward Induction. That is, however, not the "real" definition of subgame-perfect equilibrium as the concept is defined in game theory. In this section we shall introduce the "real" definition of subgame-perfect equilibrium, and show that it is equivalent to Definition 3.

Recall the concept of a strategy from Definition 2. Suppose we have a list of strategies, one for each player: $(s_1, s_2, \ldots, s_n)$. Then we can determine for each stage what is going to happen if players follow these strategies:

Stage 1: $a^1 = (s_1^1, s_2^1, \ldots, s_n^1)$.

Stage 2: $a^2 = (s_1^2(a^1), s_2^2(a^1), \ldots, s_n^2(a^1))$.

Stage 3: $a^3 = (s_1^3(a^1, a^2), s_2^3(a^1, a^2), \ldots, s_n^3(a^1, a^2))$.

...

Stage T: $a^T = (s_1^T(a^1, a^2, \ldots, a^{T-1}), s_2^T(a^1, a^2, \ldots, a^{T-1}), \ldots, s_n^T(a^1, a^2, \ldots, a^{T-1}))$.

We call $(a^1, a^2, \ldots, a^T)$ the outcome path implied by $(s_1, s_2, \ldots, s_n)$. Note that for every strategy combination there is a unique implied outcome path, but the same outcome path may be implied by different strategies. For example, if a^1 is the same for two strategy combinations, then for a^2 it only matters which actions players' strategies prescribe following action profile a^1. All other prescriptions of the second period strategies don't matter for the outcome path.

We can assign utility directly to strategy combinations as follows:

$$(s_1, s_2, \ldots, s_n) \to (a^1, a^2, \ldots, a^T) \to \left(u_1(a^1, a^2, \ldots, a^T), u_2(a^1, a^2, \ldots, a^T), \ldots, u_n(a^1, a^2, \ldots, a^T)\right).$$

Abusing notation a little bit, we can write for player i's utility from the outcome path implied by $(s_1, s_2, \ldots, s_n)$: $u_i(s_1, s_2, \ldots, s_n)$ for $i = 1, 2, \ldots, n$.

But now we have a static game: the strategy sets are $S_1, S_2, \ldots, S_n$, and the utility functions are the ones we just constructed. We call this static game the "normal form" of the multi-stage game with which we started:

Definition 11. *The normal form of a multi-stage game is the game that has the strategy sets $S_1, S_2, \ldots, S_n$ as defined in Definition 9, and that has the utility functions which assign to every strategy combination the utility from the outcome path that the strategy combination implies.*

Without proof we state:

Proposition 8. *The strategy combinations found by Backward Induction are Nash equilibria of the normal form of the multi-stage game.*

We can strengthen this result a lot by considering the "subgames" of a multi-stage game. Consider any initial history h^{t-1} of a multi-stage game. The remaining game is essentially a game by itself. We can define it as follows:

Definition 12. *The subgame following history h^{t-1} of a given multi-stage game is itself a multi-stage game, and is defined as follows:*

1. $\hat{I} = I$;

2. $\hat{T} = T - (t-1)$;

and for every $i \in I$:

3. $\hat{A}_i = A_i$;

4. $\hat{\mathcal{A}}_i^1 = \mathcal{A}_i^t(h^{t-1})$;

5. *for every $\hat{t} = 2, ..., \hat{T}$: $\hat{\mathcal{A}}_i^{\hat{t}} : \hat{H}^{\hat{t}-1} \twoheadrightarrow A_i$ is given by:*

$$\hat{\mathcal{A}}_i^{\hat{t}}(\hat{h}^{\hat{t}-1}) = \mathcal{A}_i^{t-1+\hat{t}}(h^{t-1}, \hat{h}^{\hat{t}-1}) \text{ for all } \hat{h}^{\hat{t}-1} \in \hat{H}^{\hat{t}-1}$$

6. *$\hat{u}_i : \hat{H}^{\hat{T}} \to \mathbb{R}$ satisfies:*

$$\hat{u}_i(\hat{h}^{\hat{T}}) = u_i(h^{t-1}, \hat{h}^{\hat{T}}) \text{ for all } \hat{h}^{\hat{T}} \in \hat{H}^{\hat{T}}.$$

From any given strategy s_i for the original game we can derive an induced strategy $\hat{s}_i$ for the subgame by setting:

$$\hat{s}_i^1 = s_i^t(h^{t-1}),$$

and, for every $\hat{t} = 2, \ldots, \hat{T}$ and every $\hat{h}^{\hat{t}} \in \hat{H}^{\hat{t}}$:

$$\hat{s}_i^{\hat{t}}(\hat{h}^{\hat{t}-1}) = s_i^{t-1+\hat{t}}(h^{t-1}, \hat{h}^{\hat{t}-1}).$$

Proposition 8 implies:

Proposition 9. *The strategy combinations found by Backward Induction induce Nash equilibria of the normal form of every subgame of the multi-stage game.*

Now we can give the "true" definition of subgame-perfect equilibria:

Definition 13. *A strategy combination $(s_1, s_2, \ldots, s_n)$ for a multi-stage game is a subgame-perfect equilibrium of the multi-stage game if it is a Nash equilibrium of the normal-form of the multi-stage game, and if it induces for every subgame of the multi-stage game a Nash equilibrium of that subgame.*

Propositions 8 and 9 imply that the strategy combinations found by Backward Induction are a subset of the set of subgame-perfect equilibria. But, in fact, the following, stronger claim is true:

Proposition 10. *The set of strategy combinations found by Backward Induction equals the set of subgame-perfect equilibria of the multi-stage game.*

The formal proof of this result requires so much notation that it is not worthwhile to provide it here. I explain informally one interesting part of this proof. To show that every strategy combination found by Backward Induction is a Nash equilibrium of the game itself and all of its subgames, we have to show that in the game itself and in any subgame no player can gain by deviating from a strategy combination identified by Backward Induction. Let the first period of a subgame be period t. The definition of Backward Induction implies that no player has an incentive to deviate in period t, if the player does not alter their behavior in subsequent periods. We have to show that this implies that also more complicated deviations, in period t and in subsequent periods, are not beneficial. This is known as the "one-shot deviation principle." It says: "If any one-shot deviation is not worthwhile, then the same is true for all more complicated deviations." The reason this is true is as follows: Consider any complicated, multi-period deviation from some strategy s_i. Then, if no single-period deviation is beneficial, we can adjust the deviation so that behavior in the last period T is the same as in strategy s_i, and the deviation must still be profitable. But then we can also adjust deviation so that behavior in the period $T-1$ is the same as in strategy s_i, and the deviation must still be profitable. By induction, we infer that a one period deviation must be profitable. But this contradicts the assumption that no one-shot deviation is profitable.

6. Conceptual Aspects of the Subgame-Perfect Equilibrium Concept

Consider the following two player game, known as "ultimatum bargaining." Two players negotiate over the division of 10 Dollars. Player 1 can either propose that she herself gets 1 Dollars, or 5 Dollars, or 9 Dollars, in each case leaving the rest for player 2. Then player 2 can accept or reject player 1's proposal. If player 2 accepts the proposal, then players' utilities equals the money they get according to player 1's proposal. Otherwise, their utilities are zero. This is a game of perfect information. It is shown in Figure 3, with the two players' utilities listed at the end of the game tree.

I then show in Figure 4 the normal form of this game. If you investigate Figure 4, you will find that there are seven Nash equilibria: (9,1) and (YYY), (9,1) and (YNY), (9,1) and (NYY), (9,1) and (NNY), (5,5) and (YYN), (5,5) and (NYN), and (1,9) and (YNN). In each equilibrium, player 1's proposal is accepted. Moreover all proposals where player 1 demands for more money than he gets in the equilibrium are rejected. Proposals where player 1 demands less than he gets in the equilibrium may be accepted or rejected.

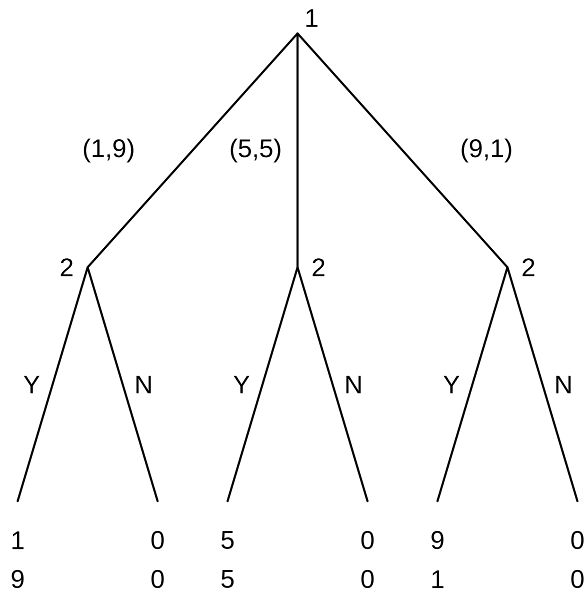

Figure 3

	YYY	YYN	YNY	YNN	NYY	NYN	NNY	NNN
(1,9)	1,9	1,9	1,9	1,9	0,0	0,0	0,0	0,0
(5,5)	5,5	5,5	0,0	0,0	5,5	5,5	0,0	0,0
(9,1)	9,1	0,0	9,1	0,0	9,1	0,0	9,1	0,0

Figure 4

The only subgame-perfect equilibrium is (9,1) and (YYY). The reason is that in the subgames in which player 2 moves, she chooses between accepting a proposal that gives her positive utility, and rejecting the proposal, which gives her zero utility. Player 2's only rational choice, and thus the only Nash equilibrium, is to say "yes." Therefore, in a subgame-perfect equilibrium player 2 has to choose (Y,Y,Y). But if player 2 accepts all offers, player 1 maximizes his own utility by proposing the offer that is most favorable to him, i.e. (9,1).

The other Nash equilibria involve "non-credible threats." In those equilibria, player 2 threatens to reject some offers. This is not a credible threat because, if player 2 has to carry out the threat, it will not be in his interest. Every offer gives him a positive utility. Rejecting the offer gives him zero utility. The reason the threat is part of a Nash equilibrium is that player 1 anticipates the threat, and therefore does not make the offer that player 2 threatens to reject. Therefore, player 2 never has to prove that he will actually carry out his threat. Subgame-perfect equilibrium rules out Nash equilibria that are sustained by threats (more generally: plans) that are not credible, in the sense that they are not optimal if the situation arises in which they need to be carried out.

That non-credible threats are ruled out is certainly desirable. But the concept of subgame-perfect equilibrium also has drawbacks. Consider the example in Figure 5. Players 1 and 2 take turns to "Stop" (S) or "Continue" (C) the game. The player who stops gets slightly more than the player who continues, but the longer they continue the larger the sum of players' utilities gets. The unique subgame-perfect equilibrium is that each player always stops.

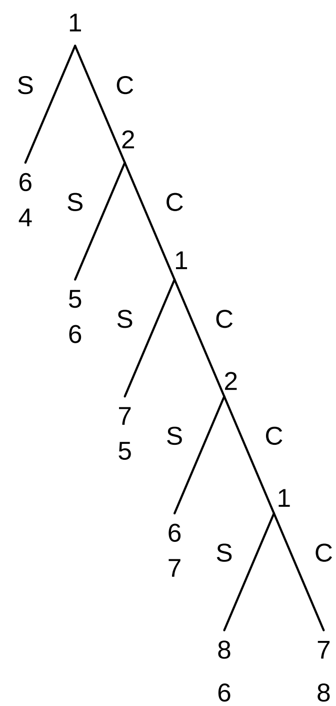

Figure 5

But suppose player 1 continued in his first move, contrary to the subgame-perfect equilibrium. Then player 2 might think that player 1 will perhaps also continue at his second node. Then player 2 might wish to continue at her first node, only to stop at her second node. This would give her utility 7, whereas stopping immediately would only give her utility 6. Why does subgame-perfect equilibrium predict that player 2 will stop? It is because, implicitly, subgame-perfect equilibrium assumes that any deviation from the equilibrium is a one-off error that will not be repeated. This seems a problematic assumption in practice. If a player deviates from the equilibrium convention once, it may seem likely that the same player will deviate from the equilibrium convention again later.

7. More Examples of Subgame-Perfect Equilibria

Suppose in stage 1 players play the game below on the left, and in stage 2 they play the game below on the right. The players' utilities are the sum of their utilities from stage 1 and stage 2.

	C	D
C	3,3	0,4
D	4,0	1,1

Stage 1

	A	B
A	3,3	0,0
B	0,0	1,1

Stage 2

Figure 6

Let us find the subgame-perfect equilibria in pure strategies of this game. Each of the four possible outcomes of stage 1 leads to a subgame in stage 2. Each of the four subgames in stage 2 are identical to the game on the right in Figure 6. The game on the right in Figure 6 has two Nash equilibria in pure strategies: (A, A) and (B, B). There are thus $2^4 = 16$ possible ways of assigning a Nash equilibrium to each of the stage 2 subgames. To find all subgame-perfect equilibria, we have to examine for each of these 16 possibilities the Nash equilibria of the first stage game.

This task looks too tedious, and therefore we shall focus on some examples. Suppose in stage 2 the same Nash equilibrium is played, regardless of the first stage outcome. Then the first stage has only one Nash equilibrium: (D, D). Note that the first stage game is the Prisoner's Dilemma. Now let's ask whether we can construct a subgame-perfect equilibrium in which players choose (C, C) in the first stage. Suppose that following (C, C) players play (A, A) in stage 2, and otherwise they play (B, B) in stage 2. Then a player may be tempted to deviate to D in stage 1, which increases their utility by 1. However, they anticipate that in the second stage they will lose a utility of 2, because instead of playing (A, A) in the second stage players will play (B, B). Thus, the deviation is not worthwhile, and we have found a subgame-perfect equilibrium where (C, C) is played in the first stage.

We won't study other subgame-perfect equilibria of the game in Figure 6.

Consider the game in Figure 7. Player 1 can give player 2 1, 2 or 3 Dollars. Player 2 can reward player 1 by giving him the same amount that player 2 gets. It is costless for player 2 to give player 1 money. Players' utility equals the amount of money that they receive.

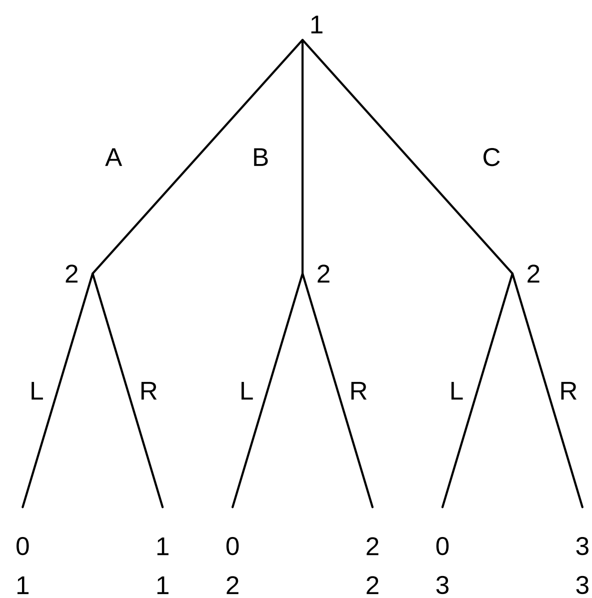

Figure 7

Player 2 is indifferent at every of his decision nodes. However, her equilibrium strategy has to pick one particular choice for each decision node. Therefore, there are $2^3 = 8$ possible solutions to the second stage. If we want to find the set of all subgame-perfect equilibria, we have to find for each of these player 1's best choice. I give just two examples of subgame-perfect equilibria of the game in Figure 7. Player 2 might never return any money to player 1. Then player 1' is indifferent between all choices, and can choose any of them in a subgame-perfect equilibrium. Another subgame-perfect equilibrium is that player 2 rewards player 1 only if player 1 chooses 3 Dollars. Then player 1 has a strict preference to give 3 Dollars.

The examples in this section illustrate how multiplicity of equilibria in the second stage may "create" multiplicity of equilibria in the first stage.

8. Games with Infinitely Many Stages

Sometimes we are interested in games in which the time horizon T is infinity ($T = \infty$) rather than finite. In this section we shall briefly consider two such examples. The first example is derived from the Prisoner's Dilemma which we encountered in Topic 1:

	C	D
C	3,3	0,4
D	4,0	1,1

Figure 8

Suppose the two players play this game in each of periods $t = 1, 2, \ldots$, where there is no final period. Suppose also that at the end of each period each player observes whether the other player cooperated or defected. Let each player's payoff be the sum of discounted payoffs in each period, where payoff in period t is discounted with discount factor δ_i^{t-1}, where $\delta_i \in (0, 1)$ is player i's discount factor. For simplicity, let's assume $\delta_1 = \delta_2 = \delta$.

The concepts of strategies, normal form, and subgames that we defined earlier can be adapted straightforwardly to games with infinitely many stages. I omit the details. We can also adapt the concept of subgame-perfect equilibrium, but obviously we cannot use Backward Induction to find subgame-perfect equilibria, because there is no final period with which we could start. The problem of finding subgame-perfect equilibria thus becomes more complicated. One helpful fact is that the one-shot deviation principle remains valid. I shall now give two examples of subgame-perfect equilibria of the infinitely repeated Prisoners' Dilemma, and use the one-shot deviation principle to verify that they are subgame-perfect equilibria.

One equilibrium is that both players defect in every period, regardless of what happened in previous periods. Thus, players repeat in every period the Nash equilibrium of the static game. Let's verify that no one-shot deviation is beneficial. If a player deviates in any given period, but sticks to the equilibrium in all future periods, the player gets payoff 0 rather than 1 in the given period, and the same payoff as if he did not deviate in all future periods. Therefore, the player loses utility. The deviation is not beneficial. Therefore, by the one-shot deviation principle, this is a subgame-perfect equilibrium.

There are, however, more subgame-perfect equilibria. Consider these strategies: every player cooperates as long as both players cooperated in all previous periods. Otherwise, they defect. The corresponding outcome path is that both players cooperate in all periods. Can a player gain by deviating after any initial history? If the player is supposed to cooperate, then she will obtain:

$$3 + \delta 3 + \delta^2 3 + \ldots = \frac{3}{1-\delta}.$$

If she deviates in the current period and defects, but sticks to the equilibrium strategy in all future periods, she gets:

$$4 + \delta 1 + \delta^2 1 + \ldots = 4 + \frac{\delta}{1-\delta}.$$

The deviation is not profitable if:

$$\begin{aligned} \frac{3}{1-\delta} &\geqslant 4 + \frac{\delta}{1-\delta} \Leftrightarrow \\ \delta &\geqslant \frac{1}{3}. \end{aligned}$$

Therefore, this is a subgame-perfect equilibrium as long as players are sufficiently patient.

The theory of repeated games is a large subfield of game theory. One set of results in this theory gives conditions under which repeated games in which players have sufficiently large discount factors have subgame-perfect equilibria in which payoffs are Pareto efficient. These results actually show that the efficient equilibria co-exist with a large number of equilibria.

The results (almost) say that every outcome is possible subgame-perfect equilibrium outcome. Results along these lines are known as "folk theorems," because they were folk-wisdom among game theory folk before they were proven.

Now we consider a second example of a multi-stage game with infinitely many stages. It is known as the "alternating offer" bargaining game due to Ariel Rubinstein. Two players negotiate over the division of a dollar. In the first stage, player 1 proposes a division $(x, 1-x)$, where $x \in [0,1]$ is what player 1 gets. Then player 2 says "yes" or "no." If player 2 says "yes" the game ends, and players get what was proposed. Otherwise, we enter the second period. In this period, player 2 proposes a division $(x, 1-x)$, and player 1 says "yes" or "no." If player 1 says "yes" the game ends, and players get what was proposed. Otherwise, we enter the third period. In this period, player 1 proposes. Etc. players alternate between being the proposer and the responder. An agreement achieved in period t gives player 1 payoff $\delta^{t-1}x$ and player 2 payoff $\delta^{t-1}(1-x)$. Here, $\delta \in (0,1)$. There is no final period.

Let's construct a subgame-perfect equilibrium in which player 1's proposal is the same, regardless of what has happened before, and is equal to $(x^*, 1-x^*)$, and player 2's proposal is $(y^*, 1-y^*)$ whenever it is her turn to propose, regardless of what happened before. Let's also assume that player 2 accepts any proposal that gives her at least utility $1-x^*$ and player 1 accepts any proposal that gives her at least utility y^*. For which values of x^* and y^* is this a subgame-perfect equilibrium? Two conditions must be satisfied: firstly, player 2 must be indifferent between accepting $1-x^*$ and rejecting $1-x^*$ so as to get $1-y^*$ in the next period:

$$1 - x^* = \delta_2(1 - y^*).$$

Secondly, player 1 must be indifferent between accepting y^* now, and rejecting y^* to get x^* one period later:

$$y^* = \delta_1 x^*.$$

The solution to these two equations is:

$$x^* = \frac{1-\delta_2}{1-\delta_1\delta_2} \quad \text{and} \quad y^* = \delta_1\frac{1-\delta_2}{1-\delta_1\delta_2}.$$

It is more insightful to write this as follows:

$$x^* = \frac{1-\delta_2}{1-\delta_1\delta_2} \quad \text{and} \quad 1 - y^* = \frac{1-\delta_1}{1-\delta_1\delta_2}.$$

This equation shows for each player what they get if they are the proposer.

For these values of x^* and y^* the strategies that we proposed form a subgame-perfect equilibrium. Players will agree immediately in period 1. Indeed, these strategies are the unique subgame-perfect equilibrium of the game. The proof that these strategies are the unique subgame-perfect equilibrium is a little more involved. The proof that they are subgame-perfect equilibria is simple. We only need to check one-shot deviations. By construction, no player, as a responder, will want to deviate. All that remains to be checked is whether a proposer can gain by making a proposal that is rejected. But this is not the case because in the next period this player will be granted a share in the Dollar that makes them indifferent between accepting and becoming proposer again in the next period and making the same proposal again. Thus, a player who does not make their equilibrium proposal will delay the agreement, but will not get a better outcome.

It is immediate from the formulas above that each player gets more as a proposer if their discount factor increases, and one can also calculate that they get more as a responder when their discount factor increases. This means that more patient players get better bargaining outcomes. There is also an intrinsic advantage to being the proposer. One way of seeing this is to focus on the case when $\delta_1 = \delta_2 = \delta$. Then the proposer gets:

$$\frac{1-\delta}{1-\delta^2} = \frac{1}{1+\delta},$$

which is obviously more than 0.5.

An interesting variation on this game is obtained when we assume that, whenever a player is the responder, the player also has the option to end the negotiations, and choose some "outside option." Suppose a player i who chooses their outside option obtains utility z_i, and suppose that z_i is less than what players receive as proposers:

$$z_1 < x^* \quad \text{and } z_2 < 1 - y^*.$$

Then the subgame-perfect equilibrium that we constructed remains a subgame-perfect equilibrium, with players never choosing their outside option. The intuitive meaning of this observation is that outside options in bargaining don't improve bargaining power unless the threat of taking them is credible.

Appendix for Topic 4

Games with Infinitely Many Stages And Their Subgame-Perfect Equilibria

In this part of the Appendix for Topic 4 we shall generalize the theory of multistage games from games with a finite number of stages to games with an infinite number of stages. We shall then formally prove the one shot deviation principle.

Definition 14. *A* multi-stage game *is a list* $\left(I, (A_i)_{i\in I}, (\mathcal{A}_i^t)_{i\in I, t\in\mathbb{N}}, (u_i)_{i\in I}\right)$*, where:*

1. $I = \{1, 2, \ldots, n\}$ *is a finite set of players;*

   and for every $i \in I$*:*

2. A_i *is the set of actions potentially available to player* i *at some stage of the game;*
3. $\mathcal{A}_i^1 \subseteq A_i$ *is the set of actions available to player* i *in stage* $t = 1$*;*
4. *for every* $t \in \mathbb{N}$ *with* $t \geqslant 2$*:*
$$\mathcal{A}_i^t : H^{t-1} \twoheadrightarrow A_i,$$
is a non-empty valued correspondence that describes the set of actions available to player i *in stage* t*;*
5. $u_i : H^\infty \to \mathbb{R}$ *is player* i*'s utility function.*

Here, for every $t \in \mathbb{N}$*, we denote by* H^t *the set of histories that may occur up to stage* t*:*

$$H^1 = \bigtimes_{i=1}^{n} \mathcal{A}_i^1 \quad \text{and, for } t \geqslant 2: \quad H^t = \{(h^{t-1}, (a_1^t, a_2^t, \ldots, a_n^t)) | h^{t-1} \in H^{t-1} \text{ and } a_i^t \in \mathcal{A}_i^t(h^{t-1}) \text{ for all } i\},$$

and we denote by H^∞ *the set of infinite histories:*

$$H^\infty = \{(a^1, a^2, a^3, \ldots) | (a^1, a^2, \ldots, a^t) \in H^t \text{ for all } t \in \mathbb{N}\}.$$

Definition 15. *A (pure) strategy s_i of player i in a multi-stage game consists of:*

1. *an action $s_i^1 \in \mathcal{A}_i^1$ for the first stage;*
2. *for every $t \in \mathbb{N}$ with $t \geqslant 2$ a function: $s_i^t : H^{t-1} \to A_i$ such that:*

$$s_i^t(h^{t-1}) \in \mathcal{A}_i^t(h^{t-1}) \text{ for all } h^{t-1} \in H^{t-1}.$$

Let S_i denote the set of strategies of player i.

Definition 16. *The history $h \in H^\infty$ implied by a list of strategies $(s_1, s_2, \ldots, s_n)$ is the infinite sequence $h = (a^1, a^2, \ldots)$ such that:*

1. *$a^1 = (s_1^1, s_2^1, \ldots, s_n^1)$,*

and, for every $t \in \mathbb{N}$ with $t \geqslant 2$:

2. *$a^t = (s_1^t(a^1, a^2, \ldots, a^{t-1}), s_2^t(a^1, a^2, \ldots, a^{t-1}), \ldots, s_n^t(a^1, a^2, \ldots, a^{t-1}))$.*

For every player $i \in I$ the utility implied by a list of strategies $(s_1, s_2, \ldots, s_n)$ is given by:

$$u_i(s_1, s_2, \ldots, s_n) = u_i(h) \text{ where } h \text{ is the history implied by strategies } (s_1, s_2, \ldots, s_n).$$

Definition 17. *The normal form of a multi-stage game is the game that has the strategy sets $S_1, S_2, \ldots, S_n$ from Definition 15 and the utility functions $u_1, u_2, \ldots, u_n$ from Definition 16.*

Definition 18. *For every $t \geqslant 2$ and for every history $h^{t-1} \in H^{t-1}$ the subgame following history h^{t-1} of a given multi-stage game is itself a multi-stage game, and is defined as follows:*

1. *$\hat{I} = I$;*

and for every $i \in I$:

2. *$\hat{A}_i = A_i$;*
3. *$\hat{\mathcal{A}}_i^1 = \mathcal{A}_i^t(h^{t-1})$;*

4. *for every* $\hat{t} \in \mathbb{N}$ *with* $\hat{t} \geqslant 2$*:* $\hat{\mathcal{A}}_i^{\hat{t}} : \hat{H}^{\hat{t}-1} \twoheadrightarrow A_i$ *is given by:*

$$\hat{\mathcal{A}}_i^{\hat{t}}(\hat{h}^{\hat{t}-1}) = \mathcal{A}_i^{t-1+\hat{t}}(h^{t-1}, \hat{h}^{\hat{t}-1}) \text{ for all } \hat{h}^{\hat{t}-1} \in \hat{H}^{\hat{t}-1};$$

5. $\hat{u}_i : \hat{H}^{\infty} \to \mathbb{R}$ *satisfies:*

$$\hat{u}_i(\hat{h}) = u_i(h^{t-1}, \hat{h}) \text{ for all } \hat{h} \in \hat{H}^{\infty}.$$

Here, for every $t \in \mathbb{N}$*, we denote by* $\hat{H}^t$ *the set of histories that may occur up to stage* t*:*

$$\hat{H}^1 = \bigtimes_{i=1}^{n} \hat{\mathcal{A}}_i^1 \quad \text{and, for } t \geqslant 2: \quad H^t = \{(h^{t-1}, (a_1^t, a_2^t, \ldots, a_n^t)) | h^{t-1} \in \hat{H}^{t-1} \text{ and } a_i^t \in \hat{\mathcal{A}}_i^t(h^{t-1}) \text{ for all } i\},$$

and we denote by $\hat{H}^{\infty}$ *the set of infinite histories:*

$$\hat{H}^{\infty} = \{(a^1, a^2, a^3, \ldots) | (a^1, a^2, \ldots, a^t) \in \hat{H}^t \text{ for all } t \in \mathbb{N}\}.$$

Definition 19. *For any* $t \in \mathbb{N}$, $h^{t-1} \in H^{t-1}$, $i \in I$, *and* $s_i \in S_i$, *the induced strategy of player* i, $\hat{s}_i$, *for the subgame following history* h^{t-1} *is given by:*

$$\hat{s}_i^1 = s_i^t(h^{t-1}),$$

and, for every $t \in \mathbb{N}$ *with* $t \geqslant 2$ *and every* $\hat{h}^{t-1} \in \hat{H}^{t-1}$*:*

$$\hat{s}_i^{\hat{t}}(\hat{h}^{\hat{t}-1}) = s_i^{t-1+\hat{t}}(h^{t-1}, \hat{h}^{\hat{t}-1}).$$

Definition 20. *A strategy combination* $(s_1, s_2, \ldots, s_n)$ *for a multi-stage game is a subgame-perfect equilibrium of the multi-stage game if it is a Nash equilibrium of the normal-form of the multi-stage game, and if for every subgame following some history* h^{t-1} *the induced strategies form a Nash equilibrium of that subgame.*

The One-Shot Deviation Principle

This result simplifies the task of finding subgame-perfect equilibria. It applies only to a special class of games, however. These games are called multi-stage games that are continuous at infinity. This class is defined as follows.

Definition 21. *A mutli-stage game is continuous at infinity if there is a sequence* $(\varepsilon_t)_{t \in \mathbb{N}}$ *of positive real numbers such that* $\lim_{t \to \infty} \varepsilon_t = 0$ *and for every* $t \in \mathbb{N}$ *with* $t \geqslant 2$ *and every history* $h^{t-1} \in H^{t-1}$ *and every player* $i \in I$ *the following is true: if* s_i *and* s_i' *are two different strategies of player* i *in the subgame following history* h^{t-1}*, and if* s_{-i} *is a list of strategies for all other players for the same subgame, then*

$$|\hat{u}_i(s_i, s_{-i}) - \hat{u}_i(s_i', s_{-i})| \leqslant \varepsilon_t.$$

Here, $\hat{u}_i$ *is player* i*'s utility function in the subgame following history* h^{t-1}.

Infinitely repeated games in which players have finite strategy sets and maximize the discounted present value of per period payoffs are examples of multi-stage games that are continuous at infinity. Here is an example of a game that is not continuous at infinity: Suppose there is just one player. In every period, the player can choose A or B. The player receives payoff 1 only if she chooses A infinitely many times. Otherwise she receives payoff 0. In this game, after every subgame, payoffs 0 and 1 still remain feasible, and therefore the difference to which the above definition refers, equals 1 in every period.

Proposition 11. *The One-Shot Deviation Principle If a multi-stage game is continuous at infinity, then strategies* $(s_1, s_2, \ldots, s_n)$ *constitute a subgame-perfect equilibrium of this game if and only if for every player* $i \in I$:

1. $u_i(s_i, s_{-i}) \geqslant u_i(\tilde{s}_i, s_{-i})$ *for every strategy* $\tilde{s}_i$ *such that* $\tilde{s}_i^t = s_i^t$ *for all* $t \geqslant 2$;

and for every $t \geqslant 2$ *and every history* $h^{t-1} \in H^{t-1}$,

2. $\hat{u}_i(\hat{s}_i, \hat{s}_{-i}) \geqslant \hat{u}_i(\tilde{s}_i, \hat{s}_{-i})$ *for every strategy* $\tilde{s}_i$ *of player* i *in the subgame that follows history* h^{t-1} *that satisfies* $\tilde{s}_i^t = \hat{s}_i^t$ *for all* $t \geqslant 2$.

The example of a game that is not continuous at infinity that we gave earlier illustrates that this result is not true if a game is not continuous at infinity. Consider the following strategy of the single player in that game: choose B in every period, regardless of history. With this strategy the player's utility equals 0. No single period deviation does not raise the payoff. Nonetheless, this strategy is not optimal: only strategies in which the player chooses A infinitely many times are optimal. Therefore, this strategy is not a subgame-perfect equilibrium of the single player game.

Proof. We first argue that, if no one shot deviation increases utility for any player, then no deviation over a finite number of stages increases utility for any player. This holds for the game itself, and also for all of its subgames. Because the proof for subgames is identical to the proof of the game itself, except that the notation needs to be adjusted, we only prove the claim for the game itself. Formally, we claim that $u_i(s_i, s_{-i}) \geqslant u_i(\tilde{s}_i, s_{-i})$ not only for all strategies $\tilde{s}_i$ such that $\tilde{s}_i^t = s_i^t$ for all $t \geqslant 2$, but also for all $\tilde{s}_i$ such that $\tilde{s}_i^t = s_i^t$ for all $t \geqslant \tau$ for some $\tau \in \mathbb{N}$. We prove this by induction over τ. Because the claim is true by assumption for $\tau = 1$, it remains to show that, if the claim is true for some $\tau \in \mathbb{N}$, it is also true for $\tau + 1$. Let $(a^1, a^2, a^3, \ldots)$ be the history implied by (s_i, s_{-i}), and let $(\tilde{a}^1, \tilde{a}^2, \tilde{a}^3, \ldots)$ be the history implied by $(\tilde{s}_i, s_{-i})$ where $\tilde{s}_i$ is allowed to differ from s_i in the first $\tau + 1$ periods. Because no deviation of length τ can increase utility, we have:

$$u_i(a^1, a^2, a^3, \ldots) \geqslant u_i(\tilde{a}^1, \tilde{a}^2, \tilde{a}^3, \ldots, \tilde{a}^\tau, \bar{a}^{\tau+1}, \bar{a}^{\tau+2}, \ldots)$$

where $\bar{a}^{\tau+1}, \bar{a}^{\tau+2}, \ldots$ stands for the history that follows $(\tilde{a}^1, \tilde{a}^2, \tilde{a}^3, \ldots, \tilde{a}^\tau)$ if from period $\tau + 1$ players play according to (s_i, s_{-i}). Because, in the subgame that follows history $(\tilde{a}^1, \tilde{a}^2, \tilde{a}^3, \ldots, \tilde{a}^\tau)$ no single period deviations can raise utility, we have:

$$u_i(\tilde{a}^1, \tilde{a}^2, \tilde{a}^3, \ldots, \tilde{a}^\tau, \bar{a}^{\tau+1}, \bar{a}^{\tau+2}, \ldots) \geqslant u_i(\tilde{a}^1, \tilde{a}^2, \tilde{a}^3, \ldots, \tilde{a}^\tau, \tilde{a}^{\tau+1}, \tilde{a}^{\tau+2}, \tilde{a}^{\tau+3}, \ldots)$$

Combining the two inequalities that we have shown we get:

$$u_i(a^1, a^2, a^3, \ldots) \geqslant u_i(\tilde{a}^1, \tilde{a}^2, \tilde{a}^3, \ldots, \tilde{a}^\tau, \tilde{a}^{\tau+1}, \bar{\bar{a}}^{\tau+2}, \bar{\bar{a}}^{\tau+3}, \ldots)$$

which is what we had to show.

We now extend the argument to deviations in infinitely many periods. Again we focus for simplicity on the game itself, but the argument also applies to all subgames. Let $\tilde{s}_i$ be any arbitrary deviation by player i. For any $\tau \in \mathbb{N}$ we have: a deviation in the first τ periods does not lead to any payoff increase. Moreover, the subsequent deviations raise payoff by at most ε_τ. Therefore, we have concluded that any deviation by player i cam raise payoff by at most ε_τ. This is true for all $\tau \in \mathbb{N}$. This implies that no deviation by player i can be profitable at all. This is because if it raised utility by $\delta > 0$ then we would have a contradiction: Because $\lim_{\tau \to \infty} \varepsilon_\tau = 0$, there is τ large enough such that $\varepsilon_\tau < \delta$, and the possibility of an increase in utility by δ would contradict what we have shown, namely that no deviation can increase utility by more than ε_τ. □

The one-shot deviation principle, when applied to games with a finite time horizon, is the justification for the argument that strategy combinations found by backward induction constitute a subgame-perfect Nash equilibrium. The backward induction algorithm ensures by definition that no one shot deviations are profitable. The one-shot deviation principle then implies that multi-period deviations are not profitable either. It is for this reason that the strategies found by backward induction form a Nash equilibrium in the game itself and in all of its subgames.

We used the one shot deviation principle for infinite horizon games in Section 8 of Topic 4, when considering examples of infinitely repeated games, and when considering the alternating offer bargaining game.

Existence of Subgame-Perfect Equilibria

We begin by considering games of perfect information, that is, games such that after any history h^t there is at most one player who has an action set $\mathcal{A}_i^t$ that has more than one element. The existence of subgame-perfect equilibria in pure strategies is obvious in games of perfect information in which the action sets are finite and the time horizon is finite. Backward induction will find such an equilibrium. To see how this result can be extended consider the following simple game of perfect information: There are just two stages. In the first stage, only player 1 moves. In the second stage, only player 2 moves. Each player's action set is the interval $A_i = [0, 1]$. Player 2's set of available actions equals $[0, 1]$ regardless of player 1's action. Each player has a utility function $u_i : [0, 1]^2 \to \mathbb{R}$ that is continuous. We shall prove that a subgame-perfect equilibrium in pure strategies exist.

Consider player 2's best responses. By the compactness of player 2's action set and the continuity of player 2's utility function for every action a^1 of player 1 there exists at least one optimal choice for player 2. There might be multiple best responses. For any $a^1 \in [0,1]$ denote by $BR^2(a^1)$ the set of player 2's best responses. A difficulty is that a continuous section from player 2's best response correspondence need not exist. Therefore, player 1's anticipated utility need not be continuous. This seems to pose a difficulty for the existence of an optimal choice for player 1.

However, this is only an apparent, not a real difficulty. Consider the graph of BR^2, that is the set:

$$\{(a^1, a^2) | a^2 \in BR^1(a^1)\}.$$

By arguments that we have encountered before this is a compact set. Therefore, a maximum of u^1 on this set exists. Denote this maximum by (a^{1*}, a^{2*}). Now consider the following strategies: Player 1 chooses a^{1*}. If player 1 chooses a^{1*} then player 2 chooses a^{2*}. Otherwise, player 2 chooses some arbitrary best response to a^1. These strategies form by by construction a subgame-perfect equilibrium in pure strategies.

This argument can be generalized significantly to obtain the following result:

Proposition 12. *Suppose that*

1. *for every player i the set A_i is a compact subset of some finite dimensional Euclidean space;*
2. *for every player i and every period t the correspondence $\mathcal{A}_i^t$ is continuous;*
3. *for every $t \geqslant 2$ and for every history $h^{t-1} \in H^{t-1}$ there is at most one player i for whom the set $\mathcal{A}_i^t$ has more than 1 element;*
4. *for every player i the utility function u_i is continuous on H^∞ in the product topology,*

then at least one subgame-perfect equilibrium exists.

The continuity assumption in the last bullet point implies continuity at infinity.

Now let's consider games in which some players move simultaneously. To obtain existence results we need to allow for the possibility that players randomize. To make this technically simple we shall assume that the sets A_i of potential actions of player i are all finite. Then we can define:

Definition 22. *A behavior strategy σ_i of player i in a multi–stage game consists of:*

1. *a probability distribution $\alpha_i^1 \in \Delta\left(\mathcal{A}_i^1\right)$ for the first stage;*
2. *a function: $\sigma_i^t : H^{t-1} \to \Delta\left(A_i\right)$ for stages $t \geqslant 2$ such that:*

$$supp(\sigma_i^t(h^{t-1})) \subseteq \mathcal{A}_i^t(h^{t-1}) \text{ for all } h^{t-1} \in H^{t-1}.$$

Let $\mathcal{S}_i$ denote the set of behavior strategies of player i.

We use here the term "behavior strategy" instead of "mixed strategy" because for dynamic games there are actually two concepts of randomized strategies. One is the one described in the definition above. The other one will be discussed later inTopic 6, and it is referred to as "mixed strategies."

In Definition 22 the symbol "$\Delta(.)$" stands for the set of all probability distributions over the set in brackets. The symbol "*supp*(.)" stands for the support of the probability distribution in brackets. Formally, the support is the smallest closed set that has probability measure 1. When the probability distribution is defined over a finite set, the support is simply the set of all elements of the set that have strictly positive probability.

Proposition 13. *If all action sets A_i are finite, and if the game is continuous at infinity, then at least one subgame-perfect equilibrium in behavior strategies exists.*

Proof. Suppose first that the time horizon is finite. Then backwards induction finds all subgame-perfect equilibria. Every subgame that begins in the last stage of the game has at least one Nash equilibrium in mixed strategies, by Nash's existence theorem. If we pick one such equilibrium for every last stage subgame, then we can move on to the second to the last stage, and assume that players correctly anticipate the outcome of the last stage. Then Nash's existence theorem shows that an equilibrium of the second to last stage exists for every history. Therefore, by the one-shot deviation principle, we have obtained a subgame-perfect equilibrium for each subgame that begins in the second to last period. Iterating this argument, we obtain a subgame-perfect equilibrium for the whole game.

To extend the proof to the case that the time horizon is infinite, one constructs an appropriate limit of subgame-perfect equilibria of truncated versions of the infinite horizon game. I omit the details. □

Markov Perfect Equilibria

We introduce in this section a further refinement of the concept of subgame-perfect equilibrium: Markov perfect equilibrium. Markov-perfect equilibria are subgame-perfect equilibria in which in each period t each player's strategy depends only on some state variable, but not on the entire history of the game. The state variable is meant to summarize which aspects of the history are relevant for the future of the game.

The set of Markov perfect equilibria may be a small subset only of the set of all subgame perfect equilibria. Consider the infinitely repeated Prisoners' Dilemma. The sets of available strategies as well as the players' utilities are in all subgames essentially the same. The history of the game is irrelevant for which strategies are available and what players' utilities are. Markov perfect equilibria are therefore those that involve history-independent strategies. But this means that the infinitely repeated Prisoners' Dilemma has a unique Markov-Perfect equilibrium: all players will always defect.

A state variable on which a player at stage t conditions their decision can be viewed as a partition of the histories up to time t, i.e. as a partition of the set H^{t-1}.

Definition 23. *A partition $\mathcal{H}^{t-1}$ of H^{t-1} is a set of subsets of H^{t-1} such that:*

1. *if A and B are two elements of $\mathcal{H}^{t-1}$ then either $A = B$ or $A \cap B = \varnothing$;*

2. *the union of all elements of $\mathcal{H}^{t-1}$ is H^{t-1}.*

For any $h^{t-1} \in H^{t-1}$ we denote by $\mathcal{H}^{t-1}(h^{t-1})$ the element of $\mathcal{H}^{t-1}$ that contains h^{t-1}.

Definition 24. *If $\mathcal{H}^{t-1}$ and $\hat{\mathcal{H}}^{t-1}$ are both partitions of H^{t-1}, then we say that* <u>$\mathcal{H}^{t-1}$ *is finer than* $\hat{\mathcal{H}}^{t-1}$ *(or, equivalently,* $\hat{\mathcal{H}}^{t-1}$ *is coarser than* $\mathcal{H}^{t-1}$*)*</u> *if $\mathcal{H}^{t-1} \neq \hat{\mathcal{H}}^{t-1}$ and if for any history $h^{t-1} \in H^{t-1}$ we have:*

$$\mathcal{H}^{t-1}(h^{t-1}) \subseteq \hat{\mathcal{H}}^{t-1}(h^{t-1}).$$

The requirement of Markov perfection will be formalized as the requirement that a player's strategy assigns the same action to two histories whenever the two histories belong to the same element of some particular partition $\mathcal{H}^{t-1}$.

Definition 25. *A strategy s_i of player i is* <u>*measurable with respect to a partition* $\mathcal{H}^{t-1}$</u> *if:*

$$\mathcal{H}^{t-1}(h^{t-1}) = \mathcal{H}^{t-1}(\tilde{h}^{t-1}) \Rightarrow s_i^t(h^{t-1}) = s_i^t(\tilde{h}^{t-1}).$$

The measurability requirement is more restrictive for coarser partitions than for finer partitions. Maskin and Tirole (Markov Perfect Equilibrium, I. Observable Actions, *Journal of Economic Theory* 100 (2001), 191-219) defined a Markov Perfect equilibrium to be a subgame-perfect equilibrium in which each all players' strategies for any period t are measurable with respect to the coarsest partition for which the requirement "makes sense." I shall give a slightly simplified version of their definition.

Definition 26. *A partition $\mathcal{H}^{t-1}$ is* <u>*consistent*</u> *if whenever $\mathcal{H}^{t-1}(h^{t-1}) = \mathcal{H}^{t-1}(\tilde{h}^{t-1})$ for every player i:*

1. *the set of strategies of player i in the subgame following h^{t-1} is the same as the set of strategies of player i in the subgame following $\tilde{h}^{t-1}$;*

2. *there exist a number α and a positive number β such that:*

$$u_i(s) = \alpha + \beta\tilde{u}_i(s) \text{ for all strategy combinations in the subgame that follows } h^{t-1}.$$

Here, u_i is the utility function of player i in the subgame that follows h^{t-1} and $\tilde{u}_i$ is the utility function of player i in the subgame that follows $\tilde{h}^{t-1}$.

The second condition in this definition refers to increasing *affine* transformations rather than arbitrary increasing transformations because I want to present a definition that also makes sense when randomization is allowed. In this case the utility functions are interpreted as Bernoulli utility functions, and preferences over lotteries don't change if the utility function is subjected to an increasing affine transformation, but other increasing transformations will change the preferences over lotteries.

Definition 27. *A subgame-perfect equilibrium of a multi-stage game is called a* <u>*Markov Perfect Equilibrium*</u> *if for every $t \in \mathbb{N}$ and every player i player i's strategy for period t is measurable with respect to the coarsest consistent partition of H^{t-1}.*

Note that coarsest consistent partitions always exist.

There are several ways in which this definition could be modified so that the concept of Markov Perfect Equilibrium becomes even more restrictive:

- one might allow the transformation in item 2. of the definition of consistency to depend on s_{-i};
- one might allow in the definition of consistency that strategy sets are identical only up to relabeling;
- one might allow that histories of different length belong to the same element of a partition.

Some of these issues are pursued in Maskin and Tirole (2001), but we shall not discuss them here.

In some games, such as infinitely repeated games, the concept of Markov perfect equilibrium can serve as a refinement of the concept of subgame perfect equilibrium. In the infinitely repeated Prisoners Dilemma, for example, there are often many subgame-perfect equilibria, but only one Markov perfect equilibrium: defect always.

In some other games all subgame-perfect equilibria are also Markov-perfect equilibria. An example is Rubinstein's alternating offer bargaining game.

Why might players play Markov-perfect equilibria? It has been suggested that these are appealing because they involve simple strategies. But note that simplicity may come at a cost. If the players other than i condition on aspects of history that should not matter in a Markov-perfect equilibrium, then for player i the payoff cost of not conditioning on those aspects of history may be very high. The literature has considered other ways of defining simplicity of strategies. Perhaps the most prominent alternative is to consider only strategies that can be implemented by finite automata, and to measure the complexity of the strategy by some measure of the complexity of the finite automaton, such as the number of states of the automaton, or also the complexity of the transition function of the finite automaton.

Topic 5: Static Games of Incomplete Information

We now return to the static setting, and we introduce a new class of games, games of *incomplete* information. We mean by this games in which at least some players don't know all components of the game, that is, the player set, the strategy sets, or the utility functions. In fact, we shall focus on incomplete information about the utility functions. The other forms of incomplete information can be dealt with in a very similar way, and are less frequently considered.

I emphasized earlier that the concept of Nash equilibrium does not really rely on the assumption that players know the game. However, when interpreted as a stable convention, it does assume that all players learn to play the game in a context that is stable, that is, the same game is played repeatedly, and that all players recognize each instance when they play this same game. Incomplete information can be interpreted as an environment in which the game changes, and players cannot completely distinguish one game from another.[2]

A special case of *incomplete* information is the case of *asymmetric* information, that is, the case, in which players are incompletely informed, and different players have different knowledge. This is the case of greatest interest. Our examples will not just be games of incomplete information, but they will be games of asymmetric information.

The framework for analyzing games of incomplete information that we shall introduce is due to John Harsanyi. John Harsanyi essentially invented the theory of games with incomplete information. This invention made it possible to create game theoretic models that formalize the classic themes of the economics of information, such as adverse selection, or signaling. The economics of information predates game theory. This branch of economics emphasizes that many market phenomena and many institutions can be interpreted as adaptations of the economy to asymmetric information. The development of the theory of games with incomplete information made it possible to examine this argument, which remains one of the central themes of economics, in a unified framework.

[2]This interpretation of incomplete information is a little vague. But these notes are not the right place to discuss this interpretation in more detail.

1. Games and Type Spaces

Let $I = \{1, 2, \ldots, n\}$ be the set of players, and for every $i \in I$ let A_i be the set of actions of players i. We want to introduce incomplete information about the utility functions $u_i : A \to \mathbb{R}$ where $A = \bigtimes_{i \in I} A_i$. We want to allow for the possibility that players don't know other players' utility functions, but also that players don't know their own utility function. A simple way of modeling this mathematically is to assume that there is some parameter θ that enters the utility functions, and that players are incompletely, and potentially asymmetrically, informed about θ. Let us denote by Θ the set of possible values of θ. We assume that Θ is a subset of some finite dimensional Euclidean space $\mathbb{R}^m$. Players' utility functions are then:

$$\begin{aligned} u_i : A \times \Theta &\to \mathbb{R} \\ (a, \theta) &\mapsto u_i(a, \theta) \end{aligned}$$

Players will have beliefs about θ which take the form of probability distributions over Θ. But they will also have beliefs about other players' beliefs, and beliefs about other players' beliefs about other players' beliefs, etc. In order to allow uncertainty not only about Θ, but also uncertainty about other players' beliefs, etc., we have to introduce a model in which has different potential *types* with different beliefs. We therefore introduce the notion of a *type space.*

Definition 28. *A* type space $\mathcal{T}$ *consists of:*

- *for every player i a set T_i of possible types of player i;*
- *for every player i a mapping $\beta_i : T_i \to \Delta(\Theta \times T_{-i})$, where $T_{-i} = \bigtimes_{j \neq i} T_j$, and where $\Delta(\Theta \times T_{-i})$ is the set of all probability measures on $\Theta \times T_{-i}$.*

The types of player i correspond to the different beliefs that player i may have. The mapping β_i describes player i's beliefs. These are i's beliefs about θ and i's beliefs other players' types. If we also assume informally that player i "knows" the mapping β_j, then i's beliefs about other players' types in particular describe what i believes about other players' beliefs about θ. Indeed, we can construct for each type of each player an infinite hierarchy of beliefs, beliefs about beliefs, beliefs about beliefs about beliefs, etc. We could have tried to explicitly model such infinite hierarchies of beliefs. But it is much easier to work with type spaces than to describe infinite hierarchies of beliefs explicitly. In the Appendix for this topic we describe in more detail how infinite hierarchies of beliefs can be derived from a type space.

2. Simplifying the Framework

For analyzing incomplete information games the set Θ that we introduced in the previous section is, in fact, redundant. Once we know u_i, we can determine each player's utility directly as a function of the types of the other players. We therefore define games of incomplete information dropping the set Θ from the definition:

Definition 29. *A* game of incomplete information *is a list* $(I, (A_i)_{i\in I}, (T_i)_{i\in I}, (u_i)_{i\in I}, (\beta_i)_{i\in I})$, *where:*

- $I = \{1, 2, \ldots, n\}$ *is a finite set of players;*

and for every $i \in I$*:*

- *the set* A_i *is the set of actions of players* i*;*
- *the set* T_i *is the set of types of players* i*;*
- *the function* $u_i : A \times T \to \mathbb{R}$ *is player* i*'s utility function;*
- *the function* $\beta_i : T_i \to \Delta(T_{-i})$ *describes player* i*'s beliefs.*

3. Special Cases

We say that a game of incomplete information has *private values* if each player's utility only depends on that player's type:

$$u_i(a, (t_i, t_{-i})) = u_i(a, (t_i, t'_{-i})),$$

for all $i \in I$, $a \in A$, $t_i \in T_i$ and $t_{-i}, t'_{-i} \in T_{-i}$.

We say that a game of incomplete information has a *common prior* if there is a probability measure μ on T such that for every player $i \in I$ and every type $t_i \in T_i$ the belief $\beta_i(t_i)$ is the conditional distribution of t_{-i} under μ where we condition on player i's type being t_i.

We say that a game of incomplete information has *independent types* if there is a common prior μ, and moreover, for every player i there is a probability measure μ_i on T_i such that μ is the product of the measures μ_i. Note that if types are independent, then all types of player i have the same beliefs about other players' types.

4. Bayesian Nash Equilibria

A strategy of player i is a complete contingent plan, where player i pretends not to know which type she is, and therefore which beliefs she will hold. Even if she knows those beliefs, she has to figure out what she would do had she other beliefs because the other players don't know her beliefs.

Definition 30. *A* strategy of player *i in a game of incomplete information is a mapping*

$$\sigma_i : T_i \to A_i.$$

Note that we again focus on pure strategies and ignore, for simplicity, mixed strategies.

A Bayesian Nash equilibrium is now a list of strategies such that each strategy is a best response to all the other strategies.

Definition 31. *A* Bayesian Nash equilibrium *of a game of incomplete information is a list of strategies $(\sigma_1^*, \sigma_2^*, \ldots, \sigma_n^*)$ such that for each player i and for each type $t_i \in T_i$:*

$$\mathbb{E}_{\beta_i(t_i)}\left[u_i\left(\sigma_i^*(t_i), \left(\sigma_j^*(t_j)\right)_{j \neq i}, t_{-i}\right)\right] \geqslant \mathbb{E}_{\beta_i(t_i)}\left[u_i\left(a_i, \left(\sigma_j^*(t_j)\right)_{j \neq i}, t_{-i}\right)\right] \quad \textit{for all } a_i \in A_i.$$

Note that we require that for each player i, for each type t_i of that player, the action that player i plans to take when being that player is optimal. We could instead have written a model in which there is an ex ante probability distribution over player i's types, and we could have phrased the optimality condition ex ante, thus requiring that player i's strategy yields maximal expected utility when player i does not yet know her type. But the ex ante perspective, and the perspective of the definition provided above, are equivalent, up to some mathematical details which are unimportant. For example, they are exactly equivalent if each player has only finitely many types, and if each type has positive ex ante probability.

5. Examples

Example 1: Cournot Competition With Incomplete Information

We return to the linear Cournot game. The set of players consists of two firms: $I = \{1, 2\}$. Each firm chooses the quantity that it wishes to produce: $A_i = \mathbb{R}_+$. Inverse demand is given by: $p = 10 - a_1 - a_2$. We allow price to be negative. Let us assume that $c_1 = 0$, but c_2 may be 0 or it may be 2. We introduce the type space: $T_1 = \{0\}$, and $T_2 = \{0, 2\}$. The common prior is that each value of t_2 has probability $\frac{1}{2}$. Firm i's profits are: $u_i(a_1, a_2) = a_i(10 - a_1 - a_2) - c_i a_i$.

A strategy for firm 1 is a quantity a_1, but a strategy for firm 2 assigns a quantity a_2^0 to type $t_2 = 0$, and quantity a_2^2 to type $t_2 = 2$. Each of these quantities has to maximize expected profits for the type who chooses the quantity.

$$a_1 \in \arg\max \left[\frac{1}{2}(10 - a_1 - a_2^0)a_1 + \frac{1}{2}(10 - a_1 - a_2^2)a_1\right]$$

$$a_2^0 \in \arg\max \left[(10 - a_1 - a_2^0)a_2^0\right]$$

$$a_2^2 \in \arg\max \left[(10 - a_1 - a_2^2)a_2^2 - 2a_2^2\right]$$

You can check that first order conditions are necessary and sufficient for solutions to these maximization problems. Therefore, we obtain:

$$\begin{aligned} a_1 &= \frac{10 - \frac{a_2^0 + a_2^2}{2}}{2} \\ a_2^0 &= \frac{10 - a_1}{2} \\ a_2^2 &= \frac{8 - a_1}{2} \end{aligned}$$

The unique solution to these equations is:

$$a_1 = \frac{22}{6}, \quad a_2^0 = \frac{19}{6}, \quad \text{and} \quad a_2^2 = \frac{13}{6}.$$

This is the unique Bayesian Nash equilibrium of this game.

Note that firm 1 with marginal cost 0 produces more than firm 2 with marginal cost 0. Intuitively, firm 2 knows that firm 1 thinks that firm 2 may have marginal cost 2. Therefore, firm 1 behaves more aggressively, that is, produces more. This has the consequence that firm 2 has to produce less, even if its true marginal cost are the same as those of firm 1.

Example 2: A First Price Auction With Private Values and Independent Types:

n bidders bid in an auction for a painting. Each bidder i submits a bid $b_i \geqslant 0$. Bids are submitted simultaneously. The bidder with the highest bid wins. If several bidders have made the highest bid, the winner is randomly chosen from the set of bidders who made the highest bid. The winner's utility is $t_i - b_i$, where t_i is the type of bidder i. It represents the value of the painting to bidder i. The losers' utility is 0. The types t_i are independent and identically distributed. They follow the uniform distribution on the interval $[0, 1]$.

We want to find a Bayesian Nash equilibrium of this game. For simplicity, we focus on strategies where all bidders have the same strategy, and each bidder i's bid b_i is a linear and increasing function of t_i: $b_i = ct_i$ with $c > 0$. For which value(s) of c is this a Bayesian Nash equilibrium?

Consider bidder i with value t_i and suppose bidder i believes all other players follow a strategy $b_j = ct_j$. Bidder i will not want to bid above $b_i = c$, because no other bidder places a higher bid. Any bid equal to c or higher wins with probability 1, and therefore a bid higher than c just incurs higher cost, without increasing the probability of winning.

When bidder i places bid $b_i \leqslant c$, bidder i wins if all other bidders' place a bid below b_i, which happens when their values satisfy: $t_j < \frac{b_i}{c}$. All other bidders' types are below $\frac{b_i}{c}$ with probability

$$\left(\frac{b_i}{c}\right)^{n-1}.$$

This follows from the assumption that values are i.i.d and uniformly distributed.

When bidder i wins, her utility is $t_i - b_i$. When she loses, her utility is 0. Therefore, if she places bid b_i her expected utility is:

$$\left(\frac{b_i}{c}\right)^{n-1}(t_i - b_i).$$

Let us find the derivative of this expression with respect to b_i. It is:

$$(n-1)\left(\frac{b_i}{c}\right)^{n-2}\frac{1}{c}(t_i - b_i) - \left(\frac{b_i}{c}\right)^{n-1}.$$

This simplifies to:

$$\frac{1}{c}\left(\frac{b_i}{c}\right)^{n-2}\left[(n-1)t_i - nb_i\right].$$

Let us set the term in square brackets equal to zero:

$$\begin{aligned}(n-1)t_i - nb_i &= 0 \Leftrightarrow \\ b_i &= \frac{n-1}{n}t_i.\end{aligned}$$

Bidder i's utility increases as long as $b_i < \frac{n-1}{n}t_i$, and it decreases when $b_i > \frac{n-1}{n}t_i$. Taking into account that bidder i does not want to bid above c, we obtain as bidder i's best response:

$$b_i = \min\left\{\frac{n-1}{n}t_i, c\right\}.$$

We have a Bayesian Nash equilibrium in which all bidders choose $b_i = ct_i$ if this best response is the same as $b_i = ct_i$. This is the case if:

$$c = \frac{n-1}{n}.$$

Thus, within the class of Bayesian Nash equilibria that we have considered, there is a unique such equilibrium. All bidders choose:

$$b_i = \frac{n-1}{n} t_i.$$

Note that as n increases, the bids increase, and as $n \to \infty$ the bids approach the type of each bidder, t_i.

Example 3: A First Price Auction With Common Values and Independent Types

n bidders bid in an auction for a painting. Each bidder i submits a bid $b_i \geqslant 0$. Bids are submitted simultaneously. The bidder with the highest bid wins. If several bidders have made the highest bid, the winner is randomly chosen from the set of bidders who made the highest bid. The winner's utility is $v - b_i$, where v is the "true" value of the painting. It is the same for all bidders, unlike in the previous example. The losers' utility is 0.

Each bidder i observes a signal t_i. These signals are independently distributed and follow the uniform distribution on the interval $[0, 1]$. The true value of the painting is equal to the average of the private signals:

$$v_i = \frac{t_1 + t_2 + ... + t_n}{n}.$$

To find a Bayesian Nash equilibrium of this game we focus again on equilibria where all bidders have the same strategy, and each bidder i's bid b_i is a linear and increasing function of their signal t_i: $b_i = ct_i$ with $c > 0$. We shall investigate for which value(s) of c this is a Bayesian Nash equilibrium. Consider bidder i who has observed signal t_i and who believes all other players follow the strategy $b_j = ct_j$. As in the previous example, bidder i will only want to consider bids that satisfy: $b_i \leqslant c$.

When bidder i places bid b_i, bidder i wins if all other bidders place a bid below b_i, which happens when their signals satisfy: $t_j < \frac{b_i}{c}$. This happens with the probability that we calculated before:

$$\left(\frac{b_i}{c}\right)^{n-1}.$$

What is bidder i's expected utility, if she wins? Here, it is crucial that we calculate the expected value of the painting correctly. We have to calculate this expected value *conditional on bidder i winning*. But winning means that all other bidders' signals are below $\frac{b_i}{c}$. Therefore, each other bidder's conditional expected type is: $\frac{b_i}{2c}$. Therefore, the conditional expected value of the painting is:

$$\frac{1}{n}t_i + \frac{n-1}{n}\frac{b_i}{2c}.$$

Notice that $\frac{b_i}{c}$ is less than 1, so that the second term on the right hand side is less than $\frac{n-1}{n}\frac{1}{2}$, which is the unconditional expected value of the contribution of other bidders' signals to the true value. In a sense, winning is "bad news." It means the other bidders have received worse signals than the winner has. This is called the "winner's curse." When calculating conditional expected utility we have to anticipate the winner's curse.

Now we can write down the expected utility of bidder i if she makes bid b_i. It is:

$$\left(\frac{b_i}{c}\right)^{n-1}\left(\frac{1}{n}t_i + \frac{n-1}{n}\frac{b_i}{2c} - b_i\right).$$

When reading this example, it is crucial that you make sure that you understand this formula. From here onwards, the analysis of the example is algebra that is similar to, but more complex than, the algebra in the previous example.

Let us differentiate the expected utility of bidder i with respect to b_i. We get:

$$(n-1)\left(\frac{b_i}{c}\right)^{n-2}\frac{1}{c}\left(\frac{1}{n}t_i + \frac{n-1}{n}\frac{b_i}{2c} - b_i\right) + \left(\frac{b_i}{c}\right)^{n-1}\left(\frac{n-1}{n}\frac{1}{2c} - 1\right).$$

After a few steps of algebra which I leave out this simplifies to:

$$\frac{1}{c}\left(\frac{b_i}{c}\right)^{n-2}\left[\frac{n-1}{n}t_i + (n-1)\frac{b_i}{2c} - nb_i\right].$$

Let us inspect this expression. If $b_i > 0$, but very close to zero, this expression is positive. What happens as we increase b_i? If the derivative increases, then clearly the highest possible bid is optimal. But this means that the optimal bid is independent of t_i, and therefore, we certainly don't have an equilibrium in which the equilibrium bid is a linear function of t_i. Therefore, we shall focus on the case that the derivative decreases as b_i increases, so that the expected utility is concave. This is the case if:

$$\begin{aligned} \frac{n-1}{2c} - n &< 0 \Leftrightarrow \\ c &> \frac{1}{2}\frac{n-1}{n} \end{aligned}$$

Let us focus on the case that this inequality is satisfied.

Now we can find the optimal bid by setting the first derivative of expected utility with respect to b_i equal to zero. This yields:

$$b_i = \frac{n-1}{n}\frac{1}{n - \frac{n-1}{2c}}t_i.$$

Remembering that a bid above c is not sensible, we can write the optimal bid then as:

$$b_i = \min\left\{\frac{n-1}{n}\frac{1}{n - \frac{n-1}{2c}}t_i, c\right\}.$$

This is firm i's best response if its type is t_i and the other bidders have linear bidding strategies $b_j = ct_j$ where c is larger than $\frac{1}{2}\frac{n-1}{n}$.

As in the previous example we obtain an equilibrium if we solve the following equation for c:

$$c = \frac{n-1}{n}\frac{1}{n - \frac{n-1}{2c}}$$

A few steps of algebra give as the solution:

$$c = \frac{1}{2}\frac{n+2}{n}\frac{n-1}{n}.$$

We have thus found our equilibrium. Notice that c is indeed larger than $\frac{1}{2}\frac{n-1}{n}$.

We want to investigate what happens as n becomes larger. Here are some examples:

$$\begin{aligned} n = 2 \quad &\Rightarrow \quad c = 0.5 \\ n = 3 \quad &\Rightarrow \quad c \approx 0.5556 \\ n = 4 \quad &\Rightarrow \quad c = 0.5625 \\ n = 5 \quad &\Rightarrow \quad c = 0.56 \\ n = 6 \quad &\Rightarrow \quad c \approx 0.5556 \end{aligned}$$

One can check that, as n increases further, c falls, and converges back to 0.5. This reflects that an increase in n has two effects: it increases competition, thus making higher bids more attractive, but it also makes the winner's curse worse, thereby making lower bids more attractive.

Appendix for Topic 5

Type Spaces and Hierarchies of Beliefs

We begin by recalling the first definition of a type space that we offered in a main text. In fact, this definition is independent of the notion of a game. We can consider a group of agents $i \in I$ all of whom form beliefs, beliefs about beliefs, etc. about some uncertain $\theta \in \Theta$. In he first section of this appendix we will therefore develop the language of type spaces and infinite hierarchies of beliefs without any explicit reference to games.

Definition 32. *A* type space $\mathcal{T}$ *consists of:*

- *for every player i a set T_i of possible types of player i;*
- *for every player i a mapping $\beta_i : T_i \rightarrow \Delta(\Theta \times T_{-i})$, where $T_{-i} = \bigtimes_{j \neq i} T_j$, and where $\Delta(\Theta \times T_{-i})$ is the set of all probability measures on $\Theta \times T_{-i}$.*

Let us spell out how we can determine for each type in a type space what the type believes about θ, but also what the type believes about others' beliefs, and what the type believes about others' beliefs about others' beliefs, etc.

We begin observing that for every type t_i we can determine type t_i's *first order belief*, that is, type t_i's belief about θ. We denote the first order belief by $b_{i,1}$. It is an element of $B_{i,1} = \Delta(\Theta)$. It is given by:[3]

$$b_{i,1}(\theta) = \beta_i(t_i)\left(\{(\theta, t_{-i}) | t_{-i} \in T_{-i}\}\right) \text{ for all } \theta \in \Theta.$$

Denote by $\hat{b}_{i,1}$ the function that assigns to each type t_i that type's first order belief. The range of $\hat{b}_{i,1}$ is thus contained in $B_{i,1}$. Define $\hat{b}_{-i,1}$ to be the function that assigns to every vector of

[3]Formally, the following equation specifies for every singleton set $\{\theta\}$ the probability that $b_{i,1}$ assigns to this set, i.e. $b_{i,1}(\{\theta\})$. To simplify the notation, I write $b_{i,1}(\theta)$ instead of $b_{i,1}(\{\theta\})$. Also note that, to fully specify the probability measure $b_{i,1}$, we would have to define $b_{i,1}$ not just for singletons, but for all measurable sets. For simplicity, I skip this step here. It should be not too difficult to extend the equation for singletons to an analogous equation for all measurable sets.

types of agents other than i, t_{-i}, the vector of these agents' first order beliefs. The range of $\hat{b}_{-i,1}$ is thus a subset of $B_{-i,1} = \times_{j \neq i} B_{j,1}$.

Type t_i's *second order belief* is type t_i's belief about θ and about the other agents' first order beliefs. We denote the second order belief by $b_{i,2}$. It is an element of $\Delta(\Theta \times B_{-i,1})$. It is given by:

$$b_{i,2}(\theta, b_{-i,1}) = \beta_i(t_i)\left(\left\{(\theta, t_{-i}) | \hat{b}_{-i,1}(t_{-i}) = b_{-i,1}\right\}\right) \text{ for all } (\theta, b_{-i,1}) \in \Theta \times B_{-i,1}.$$

Denote by $\mathcal{B}_{i,12}$ the set of all sequences $(b_{i,1}, b_{i,2})$ of first and second order beliefs such that the marginal distribution of $b_{i,2}$ on Θ is $b_{i,1}$. Denote by $\hat{b}_{i,12}$ the function that assigns to teach type t_i the sequence of that type's first and second order belief. Note that the range of $\hat{b}_{i,12}$ is contained in $\mathcal{B}_{i,12}$. Denote by $\hat{b}_{-i,2}$ the function that assigns to every vector of types of agents other than i, t_{-i}, the vector of these agents' first and second order beliefs. The range of $\hat{b}_{-i,2}$ is thus contained in $\mathcal{B}_{-i,12} = \times_{j \neq i} \mathcal{B}_{j,12}$.

Type t_i's *third order belief* is type t_i's belief about θ and about the other agents' first and second order beliefs. We denote the third order belief by $b_{i,3}$. It is an element of $\Delta(\Theta \times \mathcal{B}_{-i,12})$. It is given by:

$$b_{i,3}(\theta, b_{-i,1}, b_{-i,2}) = \beta_i(t_i)\left(\left\{(\theta, t_{-i}) | \hat{b}_{-i,12}(t_{-i}) = (b_{-i,1}, b_{-i,2})\right\}\right)$$
$$\text{for all } (\theta, b_{-i,1}, b_{-i,2}) \in \Theta \times \mathcal{B}_{-i,12}.$$

We denote by $\mathcal{B}_{i,123}$ the set of all sequences $(b_{i,1}, b_{i,2}, b_{i,3})$ such that the marginal distribution of $b_{i,2}$ on Θ is $b_{i,1}$, and the marginal distribution of $b_{i,3}$ on $\Theta \times \mathcal{B}_{i,1}$ is $b_{i,2}$. We denote by $\hat{b}_{i,123}$ the function that assigns to teach type t_i the sequence of that type's first, second, and third order belief. Note that the range of $\hat{b}_{i,123}$ is $\mathcal{B}_{i,123}$. Denote by $\hat{b}_{-i,123}$ the function that assigns to every vector of types of agents other than i, t_{-i}, the vector of these agents' first, second, and third order beliefs. The range of $\hat{b}_{-i,123}$ is thus contained in $\mathcal{B}_{-i,123} = \times_{j \neq i} \mathcal{B}_{j,123}$.

We can continue this construction ad infinitum. Suppose we had defined first, second, ..., n-th order beliefs. Type t_i's $(n+1)$*-th order belief* is type t_i's belief about θ, and about the other agents' first, second, ..., n-th order beliefs. It is denoted by $b_{i,n+1}$. It is an element of

$\Delta(\Theta \times \mathcal{B}_{-i,12\ldots n})$. It is given by:

$$b_{i,n+1}(\theta; b_{-i,1}, b_{-i,2}, \ldots, b_{-i,n}) = \beta_i(t_i)\left(\left\{(\theta, t_{-i}) | \hat{b}_{-i,12\cdots n}(t_{-i}) = (b_{-i,1}, b_{-i,2}, \ldots, b_{-i,n})\right\}\right)$$

$$\text{for all } (\theta; b_{-i,1}, b_{-i,2}, \ldots, b_{-i,n}) \in \Theta \times \mathcal{B}_{-i,12\ldots n}.$$

We obtain for each type t_i of a player i an infinite sequence ("hierarchy") of beliefs. Denote by $\mathcal{B}_i$ the set of all sequences $(b_{i,1}, b_{i,2}, b_{i,3} \ldots)$ such that for every $n \in \mathbb{N}$ the marginal distribution of $b_{i,n+1}$ on $\Theta \times \mathcal{B}_{12\cdots(n-1)}$ is $b_{i,n}$. The infinite hierarchy of beliefs is an element of $\mathcal{B}_i$. We denote the function that assigns to every type of player i the infinite hierarchy of player i's beliefs by $\hat{b}_i$:

$$\hat{b}_i : T_i \to \mathcal{B}_i.$$

Type spaces thus give us a simple way of specifying infinite hierarchies of beliefs. Instead of going through the cumbersome process of writing down an infinite hierarchy of beliefs, we instead write down a type space.

One might ask whether every infinite hierarchy of beliefs that is potentially of interest can be reproduced in some type space. The answer to this question is "yes." Indeed, one can prove something a little stronger: There is one type space in which for every agent i the range of the function $\hat{b}_i$ is exactly the set of all infinite hierarchies of beliefs $\mathcal{B}_i$. Thus, for every infinite hierarchy of beliefs taken from $\mathcal{B}_i$ we can find some type t_i with that infinite hierarchy of beliefs. This type space is called the "universal type space." To define the universal type space we can in fact label each type t_i with the infinite hierarchy of beliefs that correspond to t_i. Because the type space is *universal* this means that for every $i \in I$:

$$T_i = \mathcal{B}_i.$$

The existence of a universal type space is then proven by showing that for all $i \in I$ there exists a mapping:

$$\beta_i : T_i \to \Delta(\Theta \times T_{-i})$$

such that for every type $t_i \in T_i$ we have:

$$\hat{b}_i(t_i) = t_i.$$

This proof is due to Jean-François Mertens and Shmuel Zamir, Formulation of Bayesian Analysis for Games with Incomplete Information, *International Journal of Game Theory* 14 (1985), 1-29. A simpler version of the argument is contained in Adam Brandenburger and Eddie Dekel, Hierarchies of Beliefs and Common Knowledge, *Journal of Economic Theory* 59 (1993), 189-198.

Partitional Models of Beliefs

Let Ω be the set of all states of the world. There is a mapping $\hat{\theta} : \Omega \to \Theta$ that determines for each state of the world ω the corresponding true value of the parameter θ. Each player i has a subjective probability measure on Ω: $\pi_i \in \Delta(\Omega)$. For each agent i there is also a partition P_i of Ω. The interpretation is this: if the state $\omega \in \Omega$ occurs, then agent i is informed of the element of P_i that contains ω. Let's denote that element by $P_i(\omega)$. Agent i's beliefs about the state of the world are then updated so that they equal the prior π_i conditioned on $P_i(\omega)$.

This model is in fact equivalent to the type space model introduced in the previous section. The set of types T_i of agent i is the same as the set of elements of type i's partition P_i. Every type t_i 's beliefs about θ and the other agents' types t_{-i} is given by:

$$\beta_i(t_i)(\theta, t_{-i}) = \frac{\pi_i \left(\left\{ \omega \in t_i | \hat{\theta}(\omega) = \theta \text{ and } \omega \in t_j \text{ for all } j \neq i \right\} \right)}{\pi_i(t_i)}$$

Thus we can construct from every partitional model of beliefs a type space model of beliefs. We can also transform every type space model of beliefs into an equivalent partitional model of beliefs. This is easy to show, but we skip the proof here.

The Impossibility of Agreeing to Disagree

A special class of partitional model of beliefs are those that are derived from a common prior. In this case the prior π_i is the same for all i. We shall write π for this prior. We shall now demonstrate that the assumption of a common prior has strong implications for infinite hierarchies of beliefs. Assuming a common prior is *not at all* without loss of generality.

For simplicity let's assume that Ω is finite, that $n = 2$, and that $\pi(\omega) > 0$ for all $\omega \in \Omega$. The two agents' partitions are P_1 and P_2. Let us write $P_1 \wedge P_2$ for the finest common coarsening of P_1 and P_2 (also called the "meet" of P_1 and P_2). For every $\omega \in \Omega$ let's denote by $(P_1 \wedge P_2)(\omega)$ the element of $P_1 \wedge P_2$ that includes ω.

Now consider any "event" E where $E \subseteq \Omega$. This event can for example include all states in which the parameter θ takes on some particular value. Or it might include all states in which some agent's first order belief is given by some particular probability measure, or, more generally, has some given features. We can determine for every state $\omega \in \Omega$ and every $i \in I$ which probability agent i attaches to E in ω. This probability is given by:

$$q_i(E, \omega) = \frac{\pi(E \cap P_i(\omega))}{\pi(P_i(\omega))}.$$

Now let us define the concept of "common knowledge" of an event E. By this we mean informally that every agent knows that E has occurred, and every agent knows that every agent knows that E has occurred, etc. For simplicity, we shall identify "knowledge that E has occurred" with "attaching probability 1 to E." The formal definition of common knowledge is as follows:

Definition 33. *Event E is* common knowledge *at ω if:*

$$(P_1 \wedge P_2)(\omega) \subseteq E.$$

To see that this formal definition reflects the informal definition, let us construct agents' infinite hierarchy of beliefs regarding E in a state ω at which E is common knowledge. Let us denote the element of $P_1 \wedge P_2$ that includes ω by $\mathbb{P}$. Consider any agent i, and any state $\omega \in \mathbb{P}$. Because $P_1 \wedge P_2$ is a coarsening of P_i we have $P_i(\omega) \subseteq \mathbb{P}$. By assumption we have: $\mathbb{P} \subseteq E$. Therefore, $q_i(E, \omega) = 1$, that is, agent i believes with probability 1 that E has occurred.

Let K_iE denote the set of all $\omega \in \Omega$ such that agent i believes with probability 1 that E has occurred. We have shown in the previous paragraph that $\mathbb{P} \subseteq K_iE$. Therefore, we can also conclude that $\mathbb{P} \subseteq K_1E \cap K_2E \cap E$. Repeating the argument int the previous paragraph, with the event $K_1E \cap K_2E \cap E$ now taking the role of the event E in the previous paragraph, we can now conclude that every agent i believes with probability 1 that $K_1E \cap K_2E \cap E$ has occurred, that is, every agent i believes with probability 1 that "E has occurred and both agents believe with probability 1 that E has occurred."

Iterating the argument, we find that at every state $\omega \in \mathbb{P}$ the event E is common knowledge in the informal sense of the word that we described earlier. Thus, the formal definition implies the informal definition.

Now let $E \subseteq \Omega$ be some event, and let $q_1, q_2 \in [0,1]$ be two probabilities. Define:

$$\mathcal{E}(E, q_1, q_2) = \{\omega \in \Omega | q_1(E, \omega) = q_1 \text{ and } q_2(E, \omega) = q_2\}.$$

Thus, $\mathcal{E}(E, q_1, q_2)$ is the event that agent 1's belief attaches probability q_1 to E and agent 2's belief attaches probability q_2 to E.

Proposition 14. *Suppose at some state $\omega \in \Omega$ we have that $\mathcal{E}(E, q_1, q_2)$ is common knowledge. Then $q_1 = q_2$.*

This is Aumann's famous theorem on the impossibility of agreeing to disagree (Robert Aumann, Agreeing to Disagree, *Annals of Statistics* 4 (1976), 1236-1239). It says that if the probabilities two agents attach to some event are common knowledge among the two agents then these probabilities have to be the same. It cannot be that any disagreement about the probability of some event is common knowledge among us. This is a consequence of the common prior assumption. It is trivial that the theorem would be false if agents had subjective priors.

Proof. Define $\mathbb{P} = (P_1 \wedge P_2)(\omega)$. Note that $\mathbb{P}$ must be the union of some finite number of elements of P_1. Let us enumerate these elements of P_1 as $P_1^1, P_1^2, \ldots, P_1^k$. For every $\omega \in \mathbb{P}$ we have:

$$q_1(E, \omega) = q_1.$$

Therefore, for every $j = 1, 2, \ldots, k$ we have:

$$\frac{\pi(E \cap P_1^j)}{\pi(P_1^j)} = q_1.$$

We can re-write this as:

$$\pi(E \cap P_1^j) = q_1 \pi(P_1^j).$$

Let us sum over all j:

$$\sum_{j=1}^{k} \pi(E \cap P_1^j) = \sum_{j=1}^{k} \left(q_1 \pi(P_1^j) \right).$$

This is the same as:

$$\pi(E \cap \mathbb{P}) = q_1 \pi(\mathbb{P}).$$

Using a similar argument we get:

$$\pi(E \cap \mathbb{P}) = q_2 \pi(\mathbb{P}).$$

But this implies:

$$q_1 = q_2.$$

□

The Impossibility of Speculative Trade

We will consider two risk-neutral individuals who are betting on the outcome of a coin toss. The coin toss can come up either Head or Tails. If Head comes up, then agent 2 has to pay a Dollar to agent 1. If Tails comes up, then agent 1 has to pay a Dollar to agent 2. The coin may or may not be fair, that is, the probability that Head comes up may be more than 0.5 or less than 0.5. The reason why the agents might engage in a bet is that each of them has private information about the bias of the coin. We shall describe this using a partitional model of beliefs.

We assume again that Ω is finite, that ther e is a common prior i, and that $\pi(\omega > 0$ for all $\omega \in \Omega$. A function $\hat{\theta}$ assigns to every state $\omega \in \Omega$ the probability $\hat{\theta}(\omega) \in [0,1]$ with which the coin lands on Head. Thus, $\hat{\theta}(\omega)$ describes the true bias of the coin. The two agents' information partitions are P_1 and P_2. These information partitions describe private knowledge about the true bias of the coin. The agents' bets may be based on this private knowledge.

Now consider the following event:

$$E_1 = \left\{ \omega \in \Omega | \frac{\sum_{\tilde{\omega} \in P_1(\omega)} (\pi(\omega)\hat{\theta}(\omega))}{\sum_{\tilde{\omega} \in P_1(\omega)} \pi(\omega)} \geqslant 0.5 \right\}.$$

This is the event that agent 1 is willing to accept the bet. Similarly, define the event that agent 2 is willing to bet as:

$$E_2 = \left\{ \omega \in \Omega | \frac{\sum_{\tilde{\omega} \in P_2(\omega)} (\pi(\omega)\hat{\theta}(\omega))}{\sum_{\tilde{\omega} \in P_2(\omega)} \pi(\omega)} \leqslant 0.5 \right\}.$$

With these definitions, $E_1 \cap E_2$ is the event that both agents are willing to accept the bet.

We might anticipate that, if the two agents bet with each other, the fact that they bet with each other becomes common knowledge. If each agent is only willing to bet if the bet is advantageous to the agent, then this means that when agents bet the event $E_1 \cap E_2$ is common knowledge.

Proposition 15. *Suppose that at some state $\omega \in \Omega$ the event $E_1 \cap E_2$ is common knowledge. Then:*

$$\frac{\sum_{\tilde{\omega} \in P_1(\omega)} (\pi(\omega)\hat{\theta}(\omega))}{\sum_{\tilde{\omega} \in P_1(\omega)} \pi(\omega)} = \frac{\sum_{\tilde{\omega} \in P_2(\omega)} (\pi(\omega)\hat{\theta}(\omega))}{\sum_{\tilde{\omega} \in P_2(\omega)} \pi(\omega)} = 0.5.$$

This proposition says that if both agents are willing to bet, then neither agent actually expects to gain a strict advantage from the bet. If, for example, betting were a little costly, then neither agent would bet.

Proof. Define $\mathbb{P} = (P_1 \wedge P_2)(\omega)$. Note that $\mathbb{P}$ must be the union of some finite number of elements of P_1. Let us enumerate these elements of P_1 as $P_1^1, P_1^2, \ldots, P_1^k$. For every P_1^j we have to have:

$$\frac{\sum_{\tilde{\omega} \in P_1^j} (\pi(\omega)\hat{\theta}(\omega))}{\pi(P_1^j)} \geqslant 0.5.$$

That is, conditional on each set P_1^j the expected value of θ has to be at least 0.5. The conditional expected value of θ conditional on $\mathbb{P}$ is the weighted average of the conditional expected value of θ conditional on each of the sets P_1^j, as the following equation makes transparent.

$$\frac{\sum_{\tilde{\omega}\in\mathbb{P}}(\pi(\omega)\hat{\theta}(\omega))}{\pi(\mathbb{P})} = \frac{\sum_{\tilde{\omega}\in P_1^1}(\pi(\omega)\hat{\theta}(\omega))}{\pi(\mathbb{P})} + \frac{\sum_{\tilde{\omega}\in P_1^2}(\pi(\omega)\hat{\theta}(\omega))}{\pi(\mathbb{P})} + \ldots + \frac{\sum_{\tilde{\omega}\in P_1^k}(\pi(\omega)\hat{\theta}(\omega))}{\pi(\mathbb{P})}$$
$$= \frac{\pi(P_1^1)}{\pi(\mathbb{P})}\frac{\sum_{\tilde{\omega}\in P_1^1}(\pi(\omega)\hat{\theta}(\omega))}{\pi(P_1^1)} + \frac{\pi(P_1^2)}{\pi(\mathbb{P})}\frac{\sum_{\tilde{\omega}\in P_1^2}(\pi(\omega)\hat{\theta}(\omega))}{\pi(P_1^2)} + \ldots + \frac{\pi(P_1^k)}{\pi(\mathbb{P})}\frac{\sum_{\tilde{\omega}\in P_1^k}(\pi(\omega)\hat{\theta}(\omega))}{\pi(P_1^k)}$$

Because each of the conditional expected values in this weighted sum is at least 0.5, we can conclude that also the weighted sum is at least 0.5:

$$\frac{\sum_{\tilde{\omega}\in\mathbb{P}}(\pi(\omega)\hat{\theta}(\omega))}{\pi(\mathbb{P})} \geqslant 0.5,$$

By an analogous argument, referring to agent 2's information partition, we can conclude:

$$\frac{\sum_{\tilde{\omega}\in\mathbb{P}}(\pi(\omega)\hat{\theta}(\omega))}{\pi(\mathbb{P})} \leqslant 0.5,$$

and hence:

$$\frac{\sum_{\tilde{\omega}\in\mathbb{P}}(\pi(\omega)\hat{\theta}(\omega))}{\pi(\mathbb{P})} = 0.5.$$

But because the conditional expected value on the left hand side of this equation equals the weighted sum of conditional expected values, conditioning on information sets of player 1, we obtain:

$$\frac{\sum_{\tilde{\omega}\in P_1^j}(\pi(\omega)\hat{\theta}(\omega))}{\pi(P_1^j)} = 0.5.$$

for every j. Similarly, if we use P_2^j as a generic symbol for the information sets of player 2 that make up $\mathbb{P}$, we get:

$$\frac{\sum_{\tilde{\omega}\in P_2^j}(\pi(\omega)\hat{\theta}(\omega))}{\pi(P_2^j)} = 0.5.$$

Thus, we have found in particular, for the information sets to which ω belongs:

$$\frac{\sum_{\tilde{\omega}\in P_1(\omega)}(\pi(\omega)\hat{\theta}(\omega))}{\sum_{\tilde{\omega}\in P_1(\omega)}\pi(\omega)} = \frac{\sum_{\tilde{\omega}\in P_2(\omega)}(\pi(\omega)\hat{\theta}(\omega))}{\sum_{\tilde{\omega}\in P_2(\omega)}\pi(\omega)} = 0.5.$$

$\square$

This result is a special version of a more general result by Paul Milgrom and Nancy Stokey ("Information, Trade and Common Knowledge", *Journal of Economic Theory* 26 (1982), 17-27.) The conclusion of this result is often referred to as the "impossibility of speculative trade." Here, the potential trade is "speculative" because if the true θ were commonly known, then it could not be the case that both agents strictly benefit from trading. All possible trade is based on "insider knowledge" about the bias of the coin. In this sense agents "speculate." But if there is a common prior, both agents are risk neutral, and the willingness to trade becomes common knowledge, then neither party can expect to strictly gain from trade.

The Electronic Mail Game

Two players, $i = 1, 2$ play with probability $1 - \pi$ game G_a:

	A	B
A	M,M	1,-L
B	-L,1	0,0

and with probability π game G_b:

	A	B
A	0,0	1,-L
B	-L,1	M,M

where $L > M > 1$ and $\pi < 0.5$.

In the literature this game is called the "Coordinated Attack Problem." Initially only player 1 knows which game is the true game. Then the following sequence of events unfolds. If the game is G_b player 1's computer automatically sends an email to player 2. With probability $\varepsilon > 0$ player 1's message gets lost. With probability $1 - \varepsilon > 0$ player 1's message reaches player 2. If player 2 receives an email from player 1, her computer automatically sends a confirmation to player 1. With probability $\varepsilon > 0$ player 2's message gets lost. With probability $1 - \varepsilon > 0$ player 2's message reaches player 1. If player 1 receives an email from player 2, his computer automatically sends a confirmation to player 2. With probability $\varepsilon > 0$ player 1's message gets lost. With probability $1-\varepsilon > 0$ player 1's message reaches player 2. ... Etc. If the game is G_a no messages are sent by anyone.

We can model this using the partitional model of beliefs. The state space is:

$$\begin{aligned}\Omega &= \{(0,0),(1,0),(1,1),(2,1),(2,2),(3,2),\ldots\} \\ &= \{(q_1,q_2) \in \mathbb{N}_0^2 | 0 \leqslant q_1 - q_2 \leqslant 1\}.\end{aligned}$$

The information partitions are:

$$P_1((0,0)) = \{(0,0)\}$$

$$\text{and, if } q_1 \geqslant 1 : P_1((q_1,q_2)) = \{(q_1,q_1-1),(q_1,q_1)\}$$

$$P_2((q_1,q_2)) = \{(q_2,q_2),(q_2+1,q_2)\}$$

The prior distribution is given by:

$$\pi((0,0)) = 1-\pi.$$

$$\text{and, if } q_1 \geqslant 1 : \pi((q_1,q_2)) = \pi(1-\varepsilon)^{q_1+q_2-1}\varepsilon.$$

Now consider a state (q_1,q_2) where q_1,q_2 are both "very large:" In such a state player 1 knows that the game is G_b. Player 1 knows that player 2 knows that the game is G_b. Player 1 knows that player 2 knows that player 1 knows that player 2 knows that the game is G_b. ... Player 1 knows that player 2 knows that player 1 knows that player 2 knows .... that player 2 knows that the game is G_b. And also player 2 knows that the game is G_b. Player 2 knows that player 1 knows that player 2 knows that the game is G_b. ... Player 2 knows that player 1 knows that player 2 knows that player 1 knows .... that player 2 knows that the game is G_b. Thus the game G_b is "almost common knowledge."

Now notice that when G_b is common knowledge, there are two pure strategy Nash equilibria: (A,A) and (B,B). The following result shows that even in states where G_b is "almost common knowledge" the game has only one Nash equilibrium: both players choose A.

Proposition 16. *There is a unique Bayesian Nash equilibrium. In this Bayesian Nash equilibrium both players choose A at all information sets.*

Proof. When $q_1 = 0$, player 1 chooses A, because he knows A to be strictly dominant.

When $q_2 = 0$, player 2 attaches probability $\frac{1-\pi}{1-\pi+\pi\varepsilon}$ to the event that the game is G_a and player 1 chooses A. The belief that is most favorable to choosing B is that player 2 believes that player 1 chooses B when the game is G_b. The probability that the game is G_b is $\frac{\pi\varepsilon}{1-\pi+\pi\varepsilon}$. Notice that this latter probability is strictly smaller than 0.5. If player 2 believed with probability 0.5 the game is G_a and player 1 chooses A and with probability 0.5 the game is G_b and player 2 chooses B, then player 2's expected payoff from choosing A is: $(M + 1)/2$, and player 2's expected payoff from choosing B is $(-L + M)/2$. Obviously the latter is less than the former. Because the actual probability of G_b is less than 0.5, player 2's best response is unambiguously to choose A.

We continue with an inductive argument. Suppose we had proved the claim for all $(q_1, q_2) \leqslant (q, q)$, where $q \geqslant 0$. In all remaining states both players are sure that the game is G_b. When $q_1 = q + 1$ player 1 attaches at least probability $\frac{\varepsilon}{\varepsilon+(1-\varepsilon)\varepsilon}$ to the event that player 2 has not received her message, and is in her information set q. But note that player 2 chooses A in her information set q, by the inductive assumption. Because $\frac{\varepsilon}{\varepsilon+(1-\varepsilon)\varepsilon} > \frac{1}{2}$, player 1 thus attaches at least probability 0.5 to the event that player 2 chooses A. But because

$$\frac{1}{2} \cdot 0 + \frac{1}{2} \cdot 1 > \frac{1}{2} \cdot (-L) + \frac{1}{2} \cdot M,$$

it follows that player 1's best response is A.

When $q_2 = q + 1$ player 2 attaches at least probability $\frac{\varepsilon}{\varepsilon+(1-\varepsilon)\varepsilon}$ to the event that player 1 has not received her message, and is in her information set $q + 1$. But we just concluded that player 1's choice in her information set $q + 1$ is A. Therefore, by the same argument that we used in the previous paragraph player 2's unique best choice is A.

This argument can be continued ad infinitum, and the result is proven. □

Global Games

There are two players, $i = 1, 2$. Each player has two strategies: "I(invest)" and "N(ot invest)". They payoffs are given by the following matrix:

	I	N
I	θ, θ	$\theta - 1, 0$
N	$0, \theta - 1$	$0, 0$

Here, $\theta \in \mathbb{R}$ is the “state of the economy.” Note that, if $\theta > 1$, I is a dominant strategy. If $\theta \in [0, 1]$, (I, I) and (N, N) are pure strategy Nash equilibria. If $\theta < 0$, N is a dominant strategy.

Now assume that players don't know the value of θ. Suppose that θ is uniformly distributed on $\mathbb{R}$. Observe that this is not really a well-defined probability measure. Nonetheless, for a slightly informal presentation of the argument, let's work with this ”improper prior” as it is sometimes called in the literature.

Neither player observes θ. But each player i gets a private signal: $x_i = \theta + \varepsilon_i$ where ε_i is normally distributed with mean 0 and standard deviation σ. Assume that the ε_i are independent.

The “improper prior” matters when we now calculate each player i's posteriors when observing $x_i = x$. Two posteriors matter: what does player i believe about θ, and what does player i believe about the signal x_j observed by the other player. I now claim that under the improper prior the posterior about θ is the normal distribution with mean x and standard deviation σ, and that player i's posterior about x_j is the normal distribution with mean x and standard deviation $\sqrt{2}\sigma$.

Where does the improper prior enter? We have removed from the calculation here the prior mean and variance of θ. If the variance of the prior tends to infinity, then neither prior mean nor prior variance matter for the posterior.

These calculations simplify the proof of the following proposition.

Proposition 17. *This game has an (essentially) unique Bayesian Nash equilibrium: each player invests if $x_i > 1/2$ and does not invest if $x_i < 1/2$.*

To understand why this result is interesting suppose $x_i \in (0, 1)$ and $\sigma \to 0$. Then it is “almost common knowledge” that $\theta = x_i$. Yet only one of the two pure strategy Nash equilibria is played in equilibrium.

Proof. We prove the result by iterated elimination of strictly dominated strategies (IESDS). We begin by showing that for $x_i > 0.5$ the only strategy surviving IESDS is I.

In Step 1 we argue that I is strictly dominant when $x_i > 1$. When $x_i > 1$, the expected value of θ is greater than 1. Therefore it is strictly dominant to Invest.

For Step n suppose we had concluded after n steps, for both players i, that I is the only surviving strategy when $x_i > \bar{x}_n$, where $\bar{x}_n > 0.5$. Then the probability that player i attaches to the event that player $j \neq i$ does not invest is at most:

$$\phi\left(\frac{\bar{x}_n - x_i}{\sqrt{2}\sigma}\right).$$

where ϕ is the cumulative distribution function of the standard normal distribution. Player i's expected utility from investing when observing x_i is therefore at least:

$$x_i - \phi\left(\frac{\bar{x}_n - x_i}{\sqrt{2}\sigma}\right).$$

This expression is increasing in x_i. There is a unique $\hat{x}$ such that this expression is 0 if $x_i = \hat{x}$. Observe that: $0.5 < \hat{x} < \bar{x}_n$. We set: $\bar{x}_{n+1} = \hat{x}$.

Now consider the sequence: $(\bar{x}_n)_{n \in \mathbb{N}}$. This sequence is decreasing. Also, the sequence is bounded below by 0.5. Therefore the sequence is converging. Denote the limit by $\bar{x}_\infty$.

For given $\bar{x}$ define $b(\bar{x})$ to be the unique solution to:

$$x_i - \phi\left(\frac{\bar{x} - x_i}{\sqrt{2}\sigma}\right) = 0.$$

The sequence $(\bar{x}_n)_{n \in \mathbb{N}}$ satisfies for all $n \in \mathbb{N}$:

$$\bar{x}_{n+1} = b(\bar{x}_n).$$

Because b is continuous, we must have:

$$b(\bar{x}_\infty) = \bar{x}_\infty.$$

Therefore:

$$\begin{aligned} \bar{x}_\infty - \phi\left(\frac{\bar{x}_\infty - \bar{x}_\infty}{\sqrt{2}\sigma}\right) &= 0 \Leftrightarrow \\ \bar{x}_\infty &= \phi(0) = 0.5. \end{aligned}$$

We have shown that for $x_i > 0.5$ the only strategy surviving IESDS is I. A symmetric argument shows that for $x_i < 0.5$ the only strategy surviving IESDS is N. $\square$

A very large literature has developed surrounding the theory of global games. Global game models have been used to model phenomena such as bank runs, or currency crises, where games with multiple equilibria are involved, and the theory of global games might help to predict when one equilibrium rather than another is played. The impact of providing additional information about θ has also been investigated, and the results of such investigations have been interpreted as providing information about things such as monetary policy announcements.

A paper by Weinstein and Yildiz (*Econometrica* 2007) has suggested a reason for being cautious in interpreting the theory of global games. They showed that for every Nash equilibrium of a complete information game there is some incomplete information perturbation in which the original game is "almost common knowledge" and which has a *unique* Bayesian Nash equilibrium that is close to the original Nash equilibrium.

Topic 6: Extensive Games

1. An Example

Let us consider a dynamic game of incomplete information. Players 1 and 2 play one of the following two games, where player 1 chooses rows and player 2 chooses columns.

	H	N
T	-1,-1	-3,0
N	2,-1	0,0

x

	H	N
T	1,1	-1,0
N	2,1	0,0

y

Figure 9

The true game is with probability 0.6 the game on the left, labeled x, and with probability 0.4 the game on the right, labeled y. Player 1 knows which is the true game, but player 2 does not. Player 1 moves first. Player 2 observes player 1's choice, and then moves second. We can graphically represent this situation as follows:

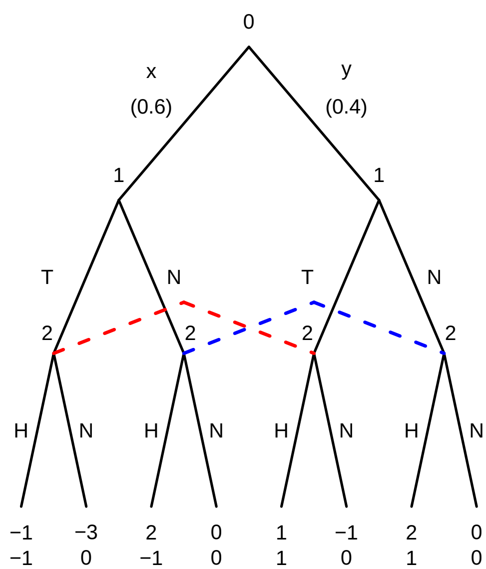

Figure 10

Reading Figure 10 from top to bottom, the game begins with a random move by player 0, which selects either the game labeled x, or the game labeled y, where the former is chosen with probability 0.6 and the latter is chosen with probability 0.4. We interpret player 0 as "Nature." Player 0 has no utilities. All moves by player 0 are random. After player 0, player 1 moves, and chooses either T or N. Then player 2 moves, choosing either H or N. The utilities are listed at the end of the tree, with player 1's utilities coming first, and player 2's utilities coming second.

The red and blue dotted lines in Figure 10 are crucial. They connect nodes of the game tree at which player 2 moves, but player 2 does not know which of the connected nodes she is at. Consider, for example, the two nodes connected by a red dotted line. These are the two nodes where player 2 observes that player 1 has chosen T, but player 2 does *not* know whether "Nature" has chosen x or y. Thus, player 2, when it is her turn, does not know which of these two nodes she is at. The two nodes connected by a blue dotted line correspond to player 1's choice of N. Player 2 does observe player 1's choice, but she does not observe "Nature's" choice. This is why we have connected these two nodes with a blue dotted line.

Nodes that are connected by a dotted line belong to the same "information set" of player 2. In fact, each node at which a player has to make a decision belongs to an information set. However, we don't show information sets for player 1 because they each have only one element. When player 1 chooses, he knows which of the two nodes he is at, because he knows whether the game is x or y.

The example in Figure 10 is a classic signaling game. Player 1 is a potential employee of player 2. If Nature chooses x, then this means that player 1 is not a good match for player 2. If Nature chooses y, then this means that player 1 is a good match. Player 1 chooses to "train" or "not to train." Training is costly for player 1. In state x the cost is 3. In state y the cost is 1. Player 2 observes whether player 1 trains, but does not know whether player 1 is a good match as an employee. Player 2 has to choose whether to "hire" or "not to hire" player 1. If player 2 hires player 1, the benefit to player 1 is equal to 2. Player 1's utility is given by this benefit minus the cost of training. Player 2 has utility -1 if she hires an employee who is not a good match, and it is 1 if she hires an employee who is a good match. Note that training does not have any impact on the benefit that the employer derives from an employee, and for the employee it is only a cost, not a benefit.

2. Informal Definition of Extensive Games

Extensive games are games that are represented by figures such as Figure 10. The basic structure is a directed graph that begins with an initial node. From each node some edges originate that connect the given node to subsequent nodes. We only consider finite extensive games here, so that the game ends with terminal nodes. Utilities are associated with each terminal node.

Next to each node we indicate the player who chooses at that node. This player may be player 0, in which case it is a chance move. The edges that connect nodes are labeled with the names of the corresponding actions. If the player who chooses at a node is player 0, then the edges are also labeled with the probability with which each of the nodes is chosen by "Nature."

The set of nodes at which some player i moves is partitioned into "information sets." The interpretation is that, if an information set is reached, the player knows that this information set has been reached, but the player does not observe which of the nodes in the information set has been reached. Therefore, all nodes that are in the same information set must be associated with the same player, and they must also all offer this player the same set of possible actions. In the graphs that I draw here, I omit information sets that contain just one element. That is, if for some node no information set is indicated then this means that the node is in an information set by itself, i.e. the player who chooses at that node knows that she is at that node.

Defining extensive games formally requires some cumbersome notation. Therefore, here, I adopt instead the approach of pointing at Figure 10, and defining extensive games as any figure that looks like Figure 10.

Another example is on the next page. This is in fact the same multi-stage game that we already discussed in Section 7 of Topic 4. Notice that to represent simultaneous moves we have introduced an artificial order of moves, specifically player 1 moves first and player 2 moves second, but have then indicated using information sets that player 2, when moving after player 1, does not observe player 1's previous move.

Of course, the games of perfect information that we showed in Topic 4 as trees are also extensive games. In those examples, we did not have to indicate any information sets because at every decision node that player who moved at that decision node knew at which decision node she was.

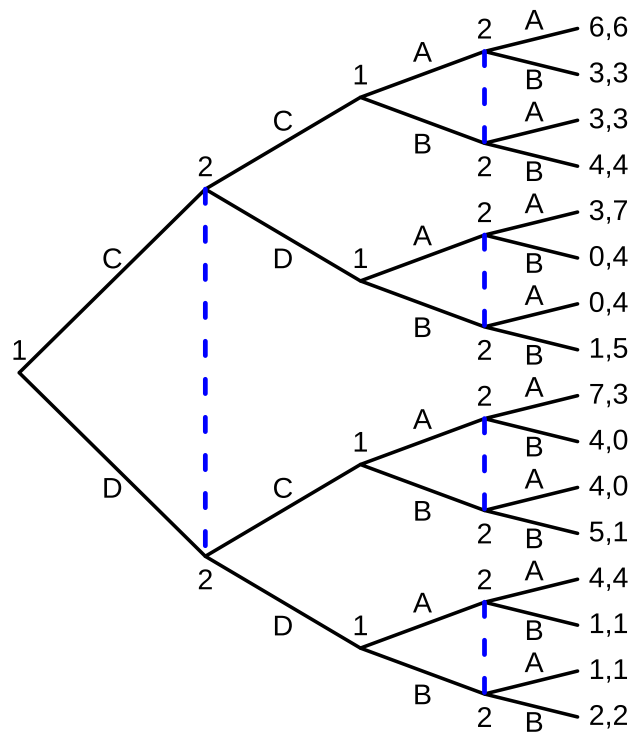

Figure 11

3. Strategies, Nash Equilibria, and Subgame-Perfect Equilibria

A strategy of a player in an extensive game assigns to each information set of that player an action. This is the natural generalization of the concept of a strategy that we introduced in Topics 4 and 5. As before, the best way to think of strategies is to think of them as "complete contingent plans," where the contingencies that players plan for include the possibility that they themselves at some point do not what they intended to do. We focus again on pure strategies, and do not discuss in these notes randomization in extensive games.

If we are given a list of strategies, one for each player, then we can determine which of the terminal nodes of the game we will reach if players follow those strategies. Therefore, we can also determine players' utilities. If there are random moves in the game, such as in Figure 10, we will obtain a probability distribution over terminal nodes, and then we can determine each players' expected utility.

Here is an example. Suppose in the game in Figure 10 player 1's strategy is to choose T at her left decision node, but N at her right decision node, and player 2's strategy is to choose H at her red information set, but to choose N at her blue information set, then the result will be that with probability 0.6 players get utilities $(-1,-1)$, and with probability 0.4 players' utilities will be $(0,0)$, so that the expected utilities will be $(-0.6,-0,6)$.

Once we have each player's strategy set and for each player, for each profile of strategies, the associated expected utility of that player, then we have again obtained the normal form of the given extensive game, and we can determine Nash equilibria of the normal form. The strategy sets can easily be quite large, and explicitly writing down the normal form of an extensive game and then finding the Nash equilibria can be quite tedious.

As we explained in the context of Topic 4, we might not be interested in all Nash equilibria, but only in the Nash equilibria that are also subgame-perfect, that is, that form Nash equilibria of the game itself, and of all subgames. In Figure 12 I have indicated one of the subgames of the extensive game shown in Figure 11. The subgame is surrounded by a red line. The important features of the subgame are these: it consists of one initial node and all subsequent nodes and edges. Moreover, there is no information set that is partially in the subgame, and partially outside of the subgame. This means that, whenever they have reached the subgame, players know that they are playing the subgame.

In Topic 4 we have focused on Nash equilibria that are also Nash equilibria of all subgames, and we can do the same in extensive games. As in multi-stage games, also in general extensive games, such Nash equilibria are called subgame-perfect equilibria. But note that some extensive games have no subgames. The game in Figure 10, for example, has no subgame.

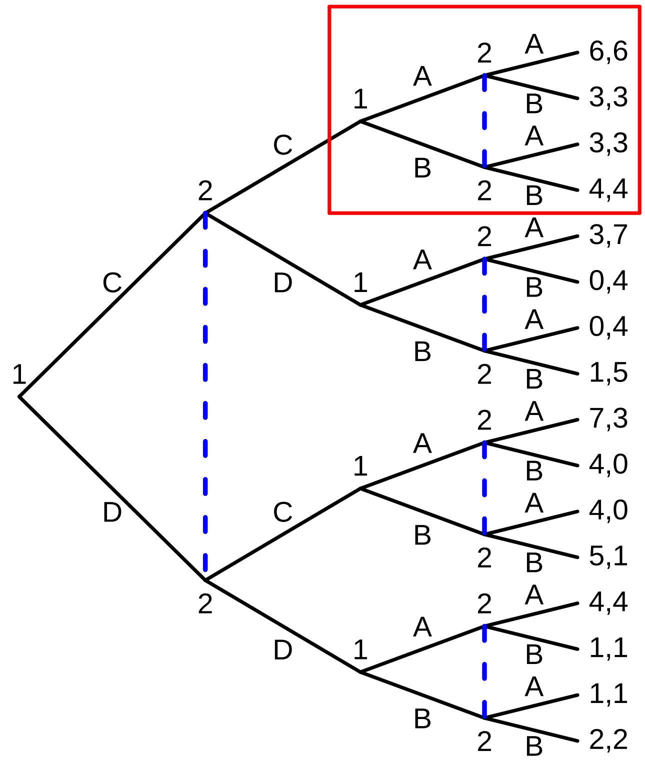

Figure 12

4. Sequential Equilibria

We now take a closer look at the game in Figure 10. Figure 13 shows the normal form of that game. Strategies of player 1 indicate first what player 1 chooses to do when she knows that she is not a good match, and second what she chooses to do when she knows she is a good match. Strategies of player 2 indicate first whether player 2 hires a potential employee who is trained, and second whether player 2 hires a potential employee who is not trained.

	HH	HN	NH	NN
TT	-0.2,-0.2	-0.2,-0.2	-2.2,0	-2.2,0
TN	0.2,-0.2	-0.6,-0,6	-1,0.4	-1.8,0
NT	1.6,-0.2	0.4,0.4	0.8,-0.6	-0.4,0
NN	2,-0.2	0,0	2,-0.2	0,0

Figure 13

The game in Figure 13 has two Nash equilibria: (NT, HN). and (NN, NN). In the first Nash equilibrium, a player 1 who is a good match signals this by training, whereas a player 1 who is not a good match does not train. Player 2 hires only those player 1s who are trained. In the second Nash equilibrium, nobody ever trains, and nobody ever gets hired. In the language that economists use to describe equilibria of signaling games, the first equilibrium is a "separating equilibrium," and the second equilibrium is a "pooling equilibrium."

A lot of thought has gone into the question whether the pooling equilibrium is plausible. Intuitively, when meeting an untrained player 1, player 2 thinks that with probability 0.6 this player is a bad match, and with probability 0.4 this player is a good match. Therefore, in expected terms, it is not worthwhile to hire the player 1. But what if player 2 meets a trained player 1? In equilibrium, this is not supposed to happen. Is player 2's plan not to hire a trained player 1 rational? Using a metaphor that we employed earlier: Is the "threat" not to hire a trained player 1 credible?

Whether it is rational for player 2 not to hire a trained player 1 depends on what player 2 believes about the quality of the match with such a player 1. If, for example, player 2 believes that a trained player 1 is with probability 0.6 a bad match, and with probability 0.4 a good match, then it is rational not to hire player 1. But if the probabilities were the reverse probabilities, then it would be rational to hire player 1.

The concept of "sequential equilibrium," which is a new equilibrium concept that we introduce in this section, requires us to describe for each player not only the player's strategy but also the player's beliefs. For every information set of every player, we have to indicate with which probability the player believes she is at any of the nodes that are included in this information set. This is obviously not relevant when an information set has only one element.

But for each of the two information sets of player 2 in Figure 10, we have to indicate which probability player 2 assigns to the events that player 1 is a good match or a bad match. Players' choices have to be expected utility maximizing given their beliefs, and given the other players' strategies. Every sequential equilibrium is a Nash equilibrium, and indeed it is a subgame-perfect equilibrium, but not every Nash equilibrium, nor every subgame-perfect equilibrium is a sequential equilibrium.

Let us examine whether the two Nash equilibria of the game in Figure 10 are sequential equilibria. For the first equilibrium, we could specify the beliefs of player 2 as follows: Player 2 believes with probability 1 that player 1 is a good match if player 1 acquires training, and otherwise player 2 believes with probability 1 that player 1 is a bad match. These are indeed the only beliefs that are consistent with the equilibrium. Sequential equilibrium requires us to specify beliefs that are consistent with the equilibrium.

For the second equilibrium, we specify that player 2 believes that a player 1 who is not trained is with probability 0.6 a bad match, and with probability 0.4 a good match. This is again the only specification of beliefs that are consistent with the equilibrium. But which beliefs should we specify for the information set that is reached when player 1 does train? The equilibrium does not pin down those beliefs. Sequential equilibrium allows us to specify any beliefs that could be calculated as conditional probabilities if for each choice of each player there were some tiny but positive probability with which the player does not make the equilibrium choice but some other choice. These tiny probabilities may be different for different players, and, for given player, for different choices of that player.

We could associate with the second Nash equilibrium the beliefs of player 2 that player 1, when trained, is with probability 0.6 a bad match and with probability 0.4 a good match. This would be the correct conditional probabilities if bad and good matches had the same tiny probability of getting trained by chance. Then the belief of player 2 at her information set should just be equal to her prior. Therefore, the second Nash equilibrium is also a sequential equilibrium. Indeed, there are many beliefs that support this sequential equilibrium. For example, player 2 could believe that bad matches are twice as likely to get trained by chance as good matches. Then, when encountering a player 1 who is trained, her posterior belief that this player is a bad match would be:

$$\frac{0.6 \cdot 2 \cdot \varepsilon}{0.6 \cdot 2 \cdot \varepsilon + 0.4 \cdot \varepsilon} = \frac{3}{4}.$$

This would also be a belief that would form, together with the equilibrium strategies, and the implied belief when player 1 is trained, a sequential equilibrium. Thus, we can see that many beliefs are allowed by sequential equilibrium. Indeed, one can easily show that *any* belief about a trained player that makes it expected utility maximizing not to hire her can be used to support the given equilibrium as a sequential equilibrium.

Which beliefs about untrained players are sensible, though? Given that it is more expensive for player 1s who are a bad match to get trained than for player 1s who are a good match, shouldn't it be more likely that a trained player is a good match? This simple question has been the focus of a lot of research in game theory, in particular in the 1980s and 1990s. The focus of this literature are belief-based refinements of sequential equilibrium. We shall not take up this question here, however, and instead accept that the signaling game of Figure 10 has sequential equilibria in which the different types separate, but also sequential equilibria in which the different types pool. No unambiguous prediction follows from the concept of sequential equilibrium for this game.

Appendix for Topic 6

Formal Definition of Extensive Games

I begin by providing the formal definition of an extensive games. A useful auxiliary term is that of a "tree," and to define that term we first need to define "directed graphs."

Definition 34. *A "directed graph" is a pair $G = (V, E)$ where*

- *V is a finite set ("vertices", "nodes");*
- *$E \subseteq V \times V$ ("edges")*

Definition 35. *Let $G = (V, E)$ be a directed graph and suppose $x, \hat{x} \in V$ where $x \neq \hat{x}$. If $(x, \hat{x}) \in E$ we say that $\hat{x}$ is a "direct successor" of x and that x is a "direct predecessor" of $\hat{x}$.*

Definition 36. *Let $G = (V, E)$ be a directed graph and suppose $x \in V$. We say that x is a "terminal node" of x if there does not exist a $\hat{x} \in V$ that is a direct successor of x.*

Definition 37. *Let $G = (V, E)$ be a directed graph, and suppose $x, \hat{x} \in V$ where $x \neq \hat{x}$. A "path" from x to $\hat{x}$ is a sequence of vertices and edges:*

$$(x^1, e^1, x^2, e^2, \ldots, x^K, e^K, x^{K+1})$$

such that for all $k = 1, 2, \ldots, K$:

- *$e^k = (x^k, x^{k+1})$*
- *$e^k \in E$*

and:

- *$x^1 = x$; $x^{K+1} = \hat{x}$.*

If there is a path from x to $\hat{x}$, we say that $\hat{x}$ is a "successor" of x and that x is a "predecessor" of $\hat{x}$.

Definition 38. *A "tree" is a triple $G = (V, E, x^0)$ where*

- *(V, E) is a directed graph;*
- *$x^0 \in V$ (the "root");*
- *for every $x \in V$ with $x \neq x^0$ there is a unique path from x^0 to x.*

Now we are ready to provide the central definition of this section.

Definition 39. *A game in extensive form is a list of the following form:*

- *a finite, non-empty set $N = \{1, 2, \ldots, n\}$ (the players);*
- *a tree $G = (V, E, x^0)$ (the game tree);*
- *a partition of the set of elements of V that are not terminal nodes into subsets V_i, for $i \in N \cup \{0\}$ (the set of nodes where player i chooses; player 0 is Nature);*
- *for every $\bar{x} \in V_0$ a probability distribution over $\{(x, \hat{x}) \in E | x = \bar{x}\}$ (probabilities for each of Nature's moves);*
- *for every player $i \in N$ a partition of V_i into k_i subsets V_i^j $(j = 1, 2, \ldots, k_i)$ such that:*

  $$\bar{x}, \bar{\bar{x}} \in V_i^j \Rightarrow |\{(x, \hat{x}) \in E | x = \bar{x}\}| = |\{(x, \hat{x}) \in E | x = \bar{\bar{x}}\}|$$

  (player i's information sets)
- *for every player $i \in N$, for every information set V_i^j of player i, a set $A_i(V_i^j)$ (the actions available to player i at information set V_i^j);*
- *for every $\bar{x} \in V_i^j$, a bijection $f_{i,\bar{x}}$ of the form:*

  $$f_{i,\bar{x}} : \{(x, \hat{x}) \in E | x = \bar{x}\} \to A_i(V_i^j),$$

  (labeling player i's choices at each information set);

- *for every $i \in N$ a function $u_i : \bar{V} \to \mathbb{R}$ where $\bar{V}$ is the set of terminal nodes (player i's utility function).*

The extensive games in Figures 10 and 11 illustrate this definition.

Strategies and Memory

Definition 40. *A "(pure) strategy of player i" is a function:*

$$s_i : \{V_i^j | j \in \{1, 2, \ldots, k_i\}\} \to \bigcup_{j \in \{1,2,\ldots,k_j\}} A_i(V_i^j),$$

such that $s_i(V_i^j) \in A_i(V_i^j)$ for every $j = 1, 2, \ldots, k_i$.

We denote the set of all pure strategies of player i by S_i.

We next introduce two different definitions of randomized strategies. Both definitions generalize the definition of a mixed strategy in static games of complete information. The first definition is:

Definition 41. *A "mixed strategy of player i" is a probability distribution $\sigma_i \in \Delta(S_i)$.*

Notice that, although we use here the term "mixed strategies," this is only one of two different generalizations of the notion of mixed strategies to extensive games. We shall use a different term in the second definition. But also the second definition generalizes what is referred to as "mixed strategies" in static games.

The above definition of mixed strategies formalizes an "ex ante" notion of randomization. Playing a mixed strategy in the sense of this definition means that, before the game begins, the player picks randomly one of the many possible contingent plans for playing the game, and then sticks to this plan for the remainder of the game.

We denote the set of all mixed strategies of player i, as defined above, by Σ_i.

The next generalization of the concept of a "mixed strategy" is an ex post notion of randomization. Instead of randomly picking a complete contingent plan for the remainder of the game, where this contingent plan is deterministic, we now formalize the behavior of a player, whenever asked to make a choice, picks randomly their next choice, and who only plans ex ante the probabilities with which she will randomize, but not the actual choice. Thus, randomization is planned initially, but is carried out later. We refer to such a complete contingent plan of random choices as a "behavior strategy."

Definition 42. *A "behavior strategy of player i" is a function:*

$$b_i : \{V_i^j | j = 1, 2, \ldots, k_i\} \rightarrow \Delta \left(\bigcup_{j \in \{1,2,\ldots,k_j\}} A_i(V_i^j) \right),$$

such that for every $j = 1, 2, \ldots, k_j$ the support of $b_i(V_i^j)$ is a subset of $A_i(V_i^j)$.

We denote the set of all behavior strategies of player i by $\mathcal{B}_i$.

An interesting question is whether it is really necessary to distinguish these two forms of randomization, or whether it is sufficient to consider just one of them. Mathematically, each of the two concepts would seem redundant if for every mixed strategy there is an "equivalent" behavior strategy, and vice versa. To study this question formally, we need to define the notion of "equivalent."

Define:

$$\mathfrak{S}_i = \Sigma_i \cup \mathcal{B}_i.$$

For every

$$\mathfrak{s} = (\mathfrak{s}_1 \ldots \mathfrak{s}_n) \in \bigtimes_{i \in N} \mathfrak{S}_i$$

and every

$$x \in V$$

we denote by

$$\rho(x; \mathfrak{s})$$

the probability that x will be visited when players choose according to $\mathfrak{s}$.

Definition 43. *Two strategies $\mathfrak{s}_i, \hat{\mathfrak{s}}_i$ are equivalent if for all $\mathfrak{s}_{-i} \in \mathfrak{S}_{-i}$ and all $x \in V$:*

$$\rho(x; (\mathfrak{s}_i, \mathfrak{s}_{-i})) = \rho(x; (\hat{\mathfrak{s}}_i, \mathfrak{s}_{-i}))$$

We now give an example of a behavior strategy for which no equivalent mixed strategy exists. Consider the following game:

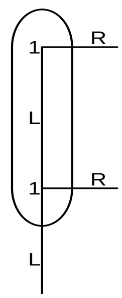

The “Absent Minded Driver” Example

This example is a one player “game.” A driver drives down the highway and needs to find the right exit. There are two exists. At each exit the driver can either turn off the motorway (choice R) or continue to drive on the motorway (choice L).[4] The information set indicates the key feature of this example: The driver is forgetful. When he sees an exit, he does not remember whether this is the first or the second exit.

In this game there is only one information set, and there are only two choices at that information set. Therefore, the driver in this game has only two pure strategies: R and L. The first pure strategy implies that the driver exits at the first exit. The second pure strategy implies that the driver never takes any exit. There is no pure strategy that implies that the driver takes the second exit.

Next, consider mixed strategies. A mixed strategy assigns some probability q to the pure strategy R, and the complementary probability $1-q$ to the pure strategy L. Note that therefore, when the driver adopts a mixed strategy, she will never exit at the second exit. With probability q the driver always exits, and thus exits at exit 1, and with the complementary probability $1-q$ the driver never exits.

[4]For our discussion here it does not matter which payoffs the driver gets if he takes either of the two exits, or continues at both exits.

Finally, consider behavior strategies. A behavior strategy assigns a probability p to exiting (R) whenever the driver encounters an exit. Thus, with probability p she will exit at the first exit. With probability $(1-p)p$ (probability of continuing at the first exit times probability of exiting at the second exit) she will exit at the second exit. With the remaining probability of $(1-p)^2$ she will continue at both exits.

Consider the mixed strategy that corresponds to $q = 0.5$. There is no corresponding behavior strategy. This is because any behavior strategy would have to correspond to a probability p that is strictly between zero and 1, and therefore, the driver would exit with positive probability at the second exit. But under the mixed strategy the driver never exits at the second exit.

Now consider the behavior strategy that corresponds to $p = 0.5$. There is no corresponding mixed strategy. This is because no mixed strategy leads the driver to exit at the second exit with positive probability, but this is what the behavior strategy accomplishes.

In this example, the single player in the game suffers from a particular form of imperfect memory. She does not know whether she has visited a given information set before. If we rule out this form of imperfect memory, then the two concepts of randomization that we are considering here are ordered in the sense that the set of outcome distributions that can be implemented by one form of randomization is contained in the set of outcome distributions that can be implemented by the other form of randomization. Specifically, we have:

Proposition 18. *Consider a game in extensive form. Let $i \in N$ be a player such that for all $j \in \{1, 2, \ldots, k_i\}$: $|A_i(v_i^j)| \geqslant 2$. Then there is for every behavior strategy of player i an equivalent mixed strategy of player i if and only if on every path that connects the root with a terminal node there is for every information set V_i^j of player i at most one node that belongs to the path.*

We omit the proof of this result. The condition on which this proposition is based is referred to as "no absent-mindedness." A player is absent-minded if she does not know remember whether she has visited an information set before, that is, if there is a path from the root of the extensive game tree to a terminal node that intersects an information set twice.

The converse of the claim in the previous proposition need not hold, i.e. there might be outcome distributions that can be implemented by a mixed strategy but not by any behavior strategy. Here is an example.

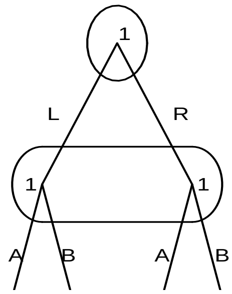

Forgetting Your own Previous Choice

This is again a single player game. The player first chooses between L and R, and then between A and B, but between the first and the second choice the player forgets what his first choice was. The player has four pure strategies. They are represented by the ordered pairs in $\{L, R\} \times \{A, B\}$. Now consider the mixed strategy that picks (L, A) with probability 0.5 and (R, B) with probability 0.5. The implied distribution over terminal nodes is that the left most and the right most terminal nodes are reached with probability 0.5 each. No behavioral strategy can achieve this. The reason is that a behavioral strategy cannot involve correlations between the choices at the first and at the second information set. Therefore, if a mixed strategy reaches the left most and the right most terminal node each with positive probability it must also reach the two other terminal nodes with positive probability.

To obtain full equivalence between behavioral and mixed strategies we need to assume that there is no imperfect recall at all. We define this as follows.

Definition 44. *Consider a game in extensive form. Let $i \in N$. We say that i has perfect recall if $x, \hat{x} \in V_i^j$ for some j implies the following.*

If $x_i^1, x_i^2, \ldots, x_i^L = x$ are the decision nodes of player i on the path from the root to x, and if $\hat{x}_i^1, \hat{x}_i^2, \ldots, \hat{x}_i^{\hat{L}} = \hat{x}$ are the decision nodes of player i on the path from the root to $\hat{x}$, then:

- $L = \hat{L}$ *(no absent-mindedness);*

- *for every* $\ell \in \{1, 2, \ldots, L-1\}$ *there is an information set* V_i^k *of player* i *such that* $x_i^\ell \in V_i^k$ *and* $\hat{x}_i^\ell \in V_i^k$ *(no forgetting of what* i *knew before)*;
- *for every* $\ell \in \{1, 2, \ldots, L-1\}$, *if* $(x_i^\ell, \bar{x})$ *is on the path from the root to* x, *and if* $(x_i^\ell, \bar{\bar{x}})$ *is on the path from the root to* $\hat{x}$, *then:*

$$f_{i,x_i^\ell}(x_i^\ell, \bar{x}) = f_{i,\hat{x}_i^\ell}(\hat{x}_i^\ell, \bar{\bar{x}}).$$

(no forgetting of what i *did before)*.

We state the main result of this section without proof:

Proposition 19 (Kuhn's Theorem)**.** *Consider a game in extensive form. If a player* $i \in N$ *has perfect recall then there is for every mixed strategy of* i *an equivalent behavior strategy of player* i.

Subgame-Perfect Equilibria

In the main text we introduced the concepts of Nash equilibrium and subgame-perfect equilibrium of a game in extensive form informally. In this section, we give a formal definition of subgame-perfect equilibrium.

Definition 45. *For a given game in extensive form let* x *be a vertex that is not the root:* $x \neq x_0$, *and that is also non-terminal. Let* $i \in N \cup \{0\}$ *be such that* $x \in V_i$. *Denote by* $V(x)$ *the set that consists of* x *and all successors of* x. *Suppose that for every* $x' \in V(x)$, *if* $x' \in V_i^j$ *for some* i *and* j, *then* $V_i^j \subseteq V(x)$. *Then the subgame starting at* x *of the given extensive game is the extensive game consisting of the following:*

- *The same set* N *of players as in the original extensive game.*
- *The set* $V(x)$ *of vertices.*
- *The set* $E \cap V(x) \times V(x)$ *of edges.*
- *The root* $x^0 = x$.

- *The player partition* $V_i(x) = V_i \cap V(x)$ *for all* $i \in N$.
- *For every* $x' \in V_0(x)$ *the same probability distribution over Nature's moves as in the original game.*
- *The information sets* $V_i^j(x) = \{V_i^j | V_i^j \subseteq V_i(x)\}$ *for all* $i \in N$.
- *For every information set* $V_i^j(x)$ *the same set of actions, and the same labeling of actions, as in the original game.*
- *The utility function* u_i *that is the restriction of the utility function in the original game to the terminal nodes that are in* $V(x)$, *for every* $i \in N$.

Definition 46. *Consider an extensive game. Let* b_i *be a behavior strategy of player* i. *Let* x *be a non-terminal vertex. Suppose that there is a subgame starting in* x. *Then we denote by* $b_i[x]$ *the restriction of* b_i *to the information sets of player* i *in the subgame that starts in* x.

Definition 47. *A list* $(b_1, b_2, \ldots, b_n)$ *of behavior strategies forms a subgame-perfect equilibrium of an extensive game the strategies form a Nash equilibrium, and if, for every non-terminal* $x \in V$ *in which a subgame starts, the restrictions of the strategies to the subgame starting in* x, $(b_1[x], b_2[x], \ldots, b_n[x])$, *form a Nash equilibrium of the subgame.*

We now introduce a version of the algorithm of "one-shot deviation principle" for arbitrary finite games in extensive form.

Definition 48. *Consider an extensive game. Let* $b = (b_1, b_2, \ldots, b_n)$ *be a list of behavior strategies. Let* x *be a vertex in which a subgame begins. The truncation of the game at* x *given* b *is the extensive game that one obtains by:*

- *removing all successors of* x, *so that* x *becomes a terminal node;*
- *associating with* x *the payoffs that players would have obtained by playing the strategies* $(b_1[x], b_2[x], \ldots, b_n[x])$ *in the subgame starting at* x.

Definition 49. *Consider an extensive game.*

- *A subgame beginning in x is called a shortest subgames if there are no subgames that begin in a successor of x.*
- *A subgame beginning in x is called a longest subgames if there are no subgames that begin in a predecessor of x.*

Definition 50. *Consider an extensive game. Let $(b_1, b_2, \ldots, b_n)$ be a list of behavior strategies. The shortest truncation of the game given b is the extensive game that one obtains by successively truncating the game in each of its longest subgames.*

The following result is a "one-shot deviation principle" for arbitrary finite extensive games.

Proposition 20. *Behavior strategies $(b_1, b_2, \ldots, b_n)$ form a subgame-perfect equilibrium if and only if they form a Nash equilibrium of the shortest truncation of every subgame.*

The one shot deviation principle is the result that justifies the backward induction algorithm:

1. Find the shortest subgames.
2. Determine for each shortest subgame a Nash equilibrium.
3. Truncate the game by removing all shortest subgames.
4. Start again at 1.

The backwards induction algorithm finds all subgame-perfect equilibria. Because the algorithm can always be implemented (by Nash's existence theorem for mixed strategy Nash equilibria), we have the following:

Proposition 21. *Every finite extensive game has at least one subgame-perfect equilibrium in pure or behavior strategies.*

Trembling Hand Perfect Equilibrium

We begin with the following example, which is due to Reinhard Selten:

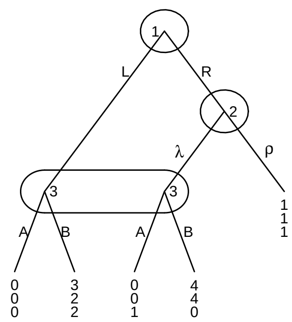

Selten's Example

This game has two Nash equilibria in pure strategies: (R, ρ, A) and (L, ρ, B). This game has no subgames. Therefore, both of these Nash equilibria are also subgame-perfect. But one might argue that the second equilibrium relies on a move that violates the "spirit of subgame-perfection:" If player 2's decision node were reached, then player 2 should not choose ρ as the equilibrium prescribes, but he should choose λ. If he chooses ρ he obtains a utility of 1, but if chooses λ anticipating player 3's equilibrium move B, he will get a utility of 4. Thus, player 2's equilibrium move is not utility maximizing if player 2's decision node is reached. The equilibrium is subgame-perfect because no subgame starts in player 2's decision node. This is because player 3's information set is such that player 3 does not know whether a move by player 2 or a move by player 1 preceded his choice. Thus, the sequence of moves that begins with player 2's choice cannot be separated from the rest of the game.

Reinhard Selten presented this example in his paper "Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games," *International Journal of Game Theory*, 1975. He commented: "There cannot be any mistakes if the players are absolutely rational. Nevertheless, a satisfactory interpretation of equilibrium points in extensive games seems to require that the possibility of mistakes is not completely excluded. This can be achieved by a point of view which looks at complete rationality as a limiting case of incomplete rationality."

Trembling-hand perfect equilibrium is one formalization of the idea that Selten expressed in the above quote. We now introduce its formal definition.

Definition 51. *A perturbation for player i is a function δ_i that assigns to every information set V_i^j of player i a strictly positive measure on $A_i(V_i^j)$ such that:*

$$\sum_{a_i \in A_i(V_i^j)} \delta_i(V_i^j)[a_i] < 1.$$

To simplify notation we shall write below $\delta_i(a_i)$ instead of $\delta_i(V_i^j)[a_i]$.

Definition 52. *Let Γ be an extensive game. Let $\delta = (\delta_1, \delta_2, \ldots, \delta_n)$ be a list of perturbations, one for each player. We denote by $\Gamma(\delta)$ the δ-perturbed game, i.e. the game in which each player i at every information set V_i^j has to choose each action $a_i \in A_i(V_i^j)$ at least with probability $\delta_i(a_i)$.*

We denote by $B_i(\delta_i)$ the set of behavior strategies of player i in the δ-perturbed game. In the perturbed game, $(b_1, b_2, \ldots, b_n)$ form a Nash equilibrium of the δ-perturbed game $\Gamma(\delta)$ if for each player i behavior strategy $b_i$maximizes $u_i(b_1, b_2 \ldots, b_n)$ among all strategies in $B_i(\delta_i)$.

We can now provide the central definition of this section:

Definition 53. *A list of behavior strategies $b^* = (b_1^*, b_2^*, \ldots, b_n^*)$ is a (trembling hand) perfect equilibrium if there exists a sequence of lists of perturbations:*

$$\left(\delta^k\right)_{k \in \mathbb{N}} = \left(\delta_1^k, \delta_2^k, \ldots, \delta_n^k\right)_{k \in \mathbb{N}}$$

and a sequence of lists of behavior strategies:

$$\left(b^k\right)_{k \in \mathbb{N}} = \left(b_1^k, b_2^k, \ldots, b_n^k\right)_{k \in \mathbb{N}}$$

such that:

- *for every player $i \in I$, every information set V_i^j of player i, and every action $a_i \in A_i(V_i^j)$: $\lim_{k \to \infty} \delta_i^k(a_i) = 0$;*

- *for every $k \in \mathbb{N}$ b^k is a Nash equilibrium of* $\Gamma(\delta^k)$*;*
- $\lim\limits_{k \to \infty} b^i = b^*$.

Existence of trembling-hand perfect equilibrium in finite games in extensive form is relatively easy to show.

Proposition 22. *Every finite game in extensive form has at least one trembling-hand perfect equilibrium.*

To see this pick any sequence of perturbations. By Nash's existence theorem, each perturbed game has at least one Nash equilibrium in behavior strategies. The sequence of Nash equilibria of the perturbed games has moreover a convergent subsequence because it is contained in a compact set. The limit of such a subsequence is a trembling-hand perfect equilibrium.

Proposition 23. *Every perfect equilibrium is a subgame-perfect equilibrium, but not every subgame-perfect equilibrium is a perfect equilibrium.*

We omit the relatively straightforward proof of the first part of this proposition. To illustrate the second part of this proposition consider this game:

	L	R
T	1,1	0,0
B	0,0	0,0

In this game, (B, R) is a subgame-perfect equilibrium (because it is a Nash equilibrium, and there are no subgames), but it is not a perfect equilibrium (because T and L are the only best responses to fully mixed strategies).

We now provide a proof that in Selten's example (R, ρ, A) is a perfect equilibrium. Choose $\varepsilon \in (0, 0.25)$. Define for all $k \in \mathbb{N}$:

$$\begin{aligned} \delta^k(a_i) &= \varepsilon^k \quad \text{for all } a_i \in \{L, \rho, A, B\} \\ \delta^k(a_i) &= \sqrt{\varepsilon^k} \quad \text{for all } a_i \in \{R, \lambda\} \end{aligned}$$

We will show that there exists a $\bar{k} \in \mathbb{N}$ such that for all $k \in \mathbb{N}$ with $k \geqslant \bar{k}$ the game $\Gamma(\delta^k)$ has a Nash equilibrium in which players assign the maximum feasible probability to R, ρ, and A. This obviously implies that (R, ρ, A) is a perfect equilibrium. We begin by considering player 1's choice in the proposed Nash equilibrium. R is preferred to L if:

$$3\varepsilon^k \leqslant 4\varepsilon^k\sqrt{\varepsilon^k} + (1 - \sqrt{\varepsilon^k})$$

This obviously holds for sufficiently large k. Next, we consider player 2's choice. ρ is preferred to λ if:

$$4\varepsilon^k \leqslant 1$$

Again, this holds for sufficiently large k. Finally, we turn to player 3's choice. A is preferred to B if:

$$\begin{aligned} 2\varepsilon^k &\leqslant (1 - \varepsilon^k)\sqrt{\varepsilon^k} \Leftrightarrow \\ 2\sqrt{\varepsilon^k} &\leqslant 1 - \varepsilon^k \end{aligned}$$

This also holds for sufficiently large k. This completes our proof.

We can also also easily see that in Selten's example (L, ρ, B) is not a perfect equilibrium. For player 2's choice of ρ to be optimal, we would have to have for sufficiently large k:

$$4(1 - \varepsilon^k) \leqslant 1$$

This does not hold for sufficiently large k.

Sequential Equilibrium

The concept of sequential equilibrium was proposed by David Kreps and Robert Wilson in their article "Sequential Equilibria," *Econometrica*, 1982, pp. 863-894. One of their purposes was too simplify Selten's definition of trembling-hand perfect equilibrium. They also wanted to introduce a formal language that would make it easier to talk about possible refinements of trembling-hand perfect equilibrium based on ideas about which “beliefs” that players may hold during the game are “plausible” and which are not.

Consider again Selten’s example. The reason why (L, ρ, B) is not a plausible Nash equilibrium of this game is that player 2’s choice of ρ is not utility maximizing, provided that player 2’s information is reached, and provided that player 2 believes that in the future players will follow the equilibrium strategies. Kreps and Wilson introduced the requirement of “sequential rationality,” which requires that at every information set players act to maximize expected utility, conditional on the information set being reached, and believing that in the future all players will stick to the equilibrium strategies. This requirement generalizes the idea of “subgame-perfection” to arbitrary extensive games.

While it is straightforward to formalize this requirement for information sets that have just a single element, it is not immediate what the requirement of “sequential rationality” would mean at information sets that have more than one element. At such information sets, a player’s expected utility maximizing strategies are well-defined only if we first specify the player’s beliefs about which of the different decision nodes in the same information set has been reached. Therefore, Kreps and Wilson introduced this definition:

Definition 54. *A* belief system of player *i is a function μ_i that assigns to every information set V_i^j of player i a probability distribution over the elements of V_i^j.*

A “sequential equilibrium” in Kreps and Wilson’s strategy consists not only of a list of strategies, one for each player, but also a list of a belief systems, one for each player. This is because only if the belief systems are specified we can check whether players are sequentially rational. This motivates the following definition.

Definition 55. *An* assessment *is a pair (b, μ) such that*

- *$b = (b_1, b_2, \ldots, b_n)$ is a list of behavior strategies, one for each player;*
- *$\mu = (\mu_1, \mu_2, \ldots, \mu_n)$ is a list of belief systems, one for each player.*

The important point is that a “sequential equilibrium” in Kreps and Wilson’s definition is not just a list of strategies, but an assessment, that is, a list of strategies, and a list of belief systems. This is an important shift in perspective in comparison to the concepts of Nash equilibrium and trembling-hand perfect equilibrium.

We can now define the concept of sequential rationality.

Definition 56. *An assessment (b, μ) is* sequentially rational *if for every player $i \in N$ and for every information set V_i^j of player i the strategy b_i maximizes player i's expected utility conditional on reaching V_i^j and conditional on player i's belief $\mu_i(V_i^j)$ at information set V_i^j.*

The requirement of sequential rationality is not enough for an assessment to be a plausible equilibrium concept. The belief system must to a certain degree be compatible with the behavior strategies. For example, when information sets are reached with positive probability when the equilibrium strategies are played, the belief system should reflect that players continue to believe in the equilibrium strategies, and have updated their beliefs according to Bayes rule. But we will also want to impose requirements for beliefs that are formed at information sets that are not reached with positive probability when the equilibrium strategies are played. Motivated by Selten's trembling-hand perfect equilibrium concept, Kreps and Wilson introduced the following requirement for assessments.

Definition 57. *An assessment (b, μ) is* consistent *if there are a sequence $(b^k)_{k \in \mathbb{N}} = (b_1^k, b_2^k, \ldots, b_n^k)_{k \in \mathbb{N}}$ of behavior strategies, and a sequence $(\mu^k)_{k \in \mathbb{N}} = (\mu_1^k, \mu_2^k, \ldots, \mu_n^k)_{k \in \mathbb{N}}$ of belief systems, such that:*

- *for every $k \in \mathbb{N}$, every $i \in N$, information set V_i^j of player i, and action $a_i \in A_i(V_i^j)$ we have:*

$$b_i^k(V_i^j)[a_i] > 0;$$

- *for every $i \in N$, information set V_i^j of player i, and action $a_i \in A_i(V_i^k)$ we have:*

$$\lim_{k \to \infty} b_i^k(V_i^j)[a_i] = b_i(V_i^j)[a_i];$$

- *for every $k \in \mathbb{N}$, $i \in N$, and every information set V_i^j of player i, $\mu_i^k(V_i^j)$ is the conditional probability implied by b^k;*

- *for every $i \in N$ and every information set V_i^j of player i:*

$$\lim_{k \to \infty} \mu_i^k(V_i^j) = \mu_i(V_i^j).$$

Intuitively, we can think of the behavior strategies b_i^k in this definition as players' behavior strategies in a perturbed version of the original game. These strategies have at every information set full support, i.e. all actions receive positive probability. Therefore, at every information set beliefs conditionally on reaching that information set about how the information set was reached are well defined. No information set is reached with zero probability. Kreps and Wilson's definition of consistency of an assessment requires that there is a sequence of full support strategies, and a corresponding sequence of conditional beliefs, such that the full support strategies converge to the strategy component of the assessment, and the conditional beliefs converge to the belief system component of the assessment.

We are now ready to define "sequential equilibrium."

Definition 58. *An assessment (b, μ) is a* sequential equilibrium *if it is consistent and sequentially rational.*

The following proposition shows formally that sequential equilibrium is a less restrictive refinement of Nash equilibrium than trembling-hand perfect equilibrium.

Proposition 24. *If b is a trembling-hand perfect equilibrium, then there is a belief system μ such that (b, μ) is a sequential equilibrium.*

We don't prove this result formally. But we explain the intuition behind the result. Recall that Selten's definition of trembling hand perfect equilibrium requires the full support strategies to be Nash equilibria of the perturbed game. This requirement is removed in Kreps and Wilson's definition of consistency. This requirement implies that all full support strategies are expected utility maximizing, subject to the minimum probability bounds in the perturbed games, conditional on every information set, and conditional on the beliefs that are derived for this information set by Bayesian updating. By the continuity of expected utility, it follows that strategies are also maximizing conditionally on the limit beliefs, provided that such limit beliefs exist. The existence of limit beliefs can be ensured by considering an appropriate subsequence of the Nash equilibria of perturbed games that converge to a trembling-hand perfect equilibrium.

Kreps and Wilson also showed that in almost all games there is no difference between trembling-hand perfect equilibrium and sequential equilibrium. Specifically, they proved:

Proposition 25. *In almost all games, if (b, μ) is a sequential equilibrium, then there exists a trembling-hand perfect equilibrium b' that implies the same probability distribution over terminal nodes as b does.*

I have not defined what Kreps and Wilson meant by "almost all games." I shall not give the precise definition. But roughly speaking it says that games in which the proposition is not true are very exceptional, and are "knife edge."

To see that not every sequential equilibrium corresponds to a trembling hand perfect equilibrium it suffices to consider static games. In static games, every Nash equilibrium is also a sequential equilibrium. Therefore, the example that we used in the discussion following Proposition 23 to show that not every subgame-perfect equilibrium is trembling-hand perfect also demonstrates that not every sequential equilibrium is trembling-hand perfect.

The following result, together with the earlier results, situates sequential equilibrium precisely between subgame-perfect equilibrium and trembling hand perfect equilibrium:

Proposition 26. *If (b, μ) is a sequential equilibrium, then b is a subgame-perfect equilibrium.*

We return to Selten's example from the previous subsection. Consider the first Nash equilibrium, (R, ρ, A). We saw earlier that this equilibrium is trembling-hand perfect. Therefore, we know from Proposition 24 that there is a belief system that forms together with these strategies a sequential equilibrium. Only the beliefs of player 3 matter. This is because the other two players' information sets are singletons. To obtain beliefs for player 3 we can consider the conditional probabilities implied by the full support strategies that we constructed as Nash equilibrium strategies for perturbed games when proving that (R, ρ, A) is trembling-hand perfect. Focusing on the left decision node in player 3's information set, the conditional probability of reaching this node is:

$$\frac{\varepsilon^k}{\varepsilon^k + (1-\varepsilon^k)\sqrt{\varepsilon^k}} = \frac{1}{1 + \frac{1-\varepsilon^k}{\sqrt{\varepsilon^k}}}.$$

As $\varepsilon^k \to 0$, this converges to 0. Thus, the Nash equilibria of the perturbed game help us to obtain a suitable limit belief for player 3: player 3 places probability 0 on the node on the left in his information set and probability 1 on the node on the right in his information set. This belief obviously renders the choice of A utility maximizing. Thus, sequential rationality is satisfied.

Moreover, consistency is satisfied by construction, given the way in which we derived player 3's beliefs from the Nash equilibria of perturbed versions of the game in Selten's example.

In Selten's example it is also clear that the strategy combination (L, ρ, B) is not part of a sequential equilibrium. This is because regardless of the belief system ρ is not a sequentially rational choice for player 2. We conclude that in this example trembling-hand perfect equilibria and sequential equilibria are the same.

I conclude this section by discussing two examples in which sequential equilibria have some surprising properties. Both examples are taken from an article by David Kreps and Gary Ramey ("Structural Consistency, Consistency, and Sequential Rationality," *Econometrica* 1987, 1331-1348). Consider first Example 1 from Kreps and Ramey (1987). The set of Nash equilibria of this game is as follows: player 1 plays R_1. Player 2 chooses R_2. Finally, player 3 chooses L and R with positive probability each, where each is chosen with probability of at least 1/3. Are there belief systems that make these choices a sequential equilibrium? Consistency of beliefs will imply that player 2's belief at her information set places probability 1 on the right hand side node in her information set. To make player 3's randomization sequentially rational, a belief system must put equal probabilities on the center and on the right decision node in player 3's information set. Moreover, consistency implies that the center decision node is infinitely more likely than the left decision node in player 3's information set. Thus, there is only one candidate for a consistent belief system that makes these choices sequentially rational: player 3's beliefs must put probability 1/2 on the center and on the right decision node in his information set.

What is remarkable about the unique sequential equilibrium of Example 1 in Kreps and Ramey (1987) is that the beliefs of player 3 imply that player 3 believes that player 1 and 2's choices are correlated. The beliefs are thus not based on any hypothesis about player 1 and 2's behavior if we require that such a hypothesis includes the assumption that players 1 and 2 behave independently.

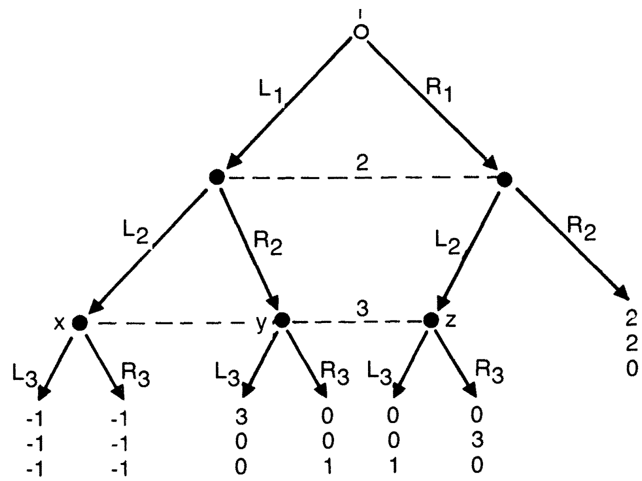

Example 1 from Kreps and Ramey (1987)

Consider next Example 2 from Kreps and Ramey (1987). In this example, there are three players: player 1a, player 1b, and player 2. Players 1a and 1b have identical utilities, and their utilities are the top of the two numbers indicated at each terminal node. The bottom number is player 2's utility. The game begins in one of the two decision nodes that are marked by an open circle. Each of these nodes has probability 0.5 of being the initial decision node. These two nodes belong to information sets of player 1a and player 1b respectively. Thus, each of the two players 1 is with probability 0.5 the first player to move. The first player to move can either "quit" or "pass." Then player 2 can choose either L or R, or player 2 can choose to pass. If player 2 chose to pass, then the other player 1 can either quit or pass. The information sets indicate that none of the three players knows whether player 1a or player 1b is the first to move.

In any Nash equilibrium of this game players 1a and 1b choose to quit. Player 2 must choose L or R with probability no more than 1/3. This implies that player 2's belief system must put strictly positive probability on both decision nodes in player 2's information set. But this means that player 2 believes that both player 1a and player 1b deviate from their equilibrium strategies. Nonetheless, when calculating expected payoff from the various choices at his information set, player 2 bases this calculation on the assumption that the other players will make equilibrium choices in the future. This seems contradictory.

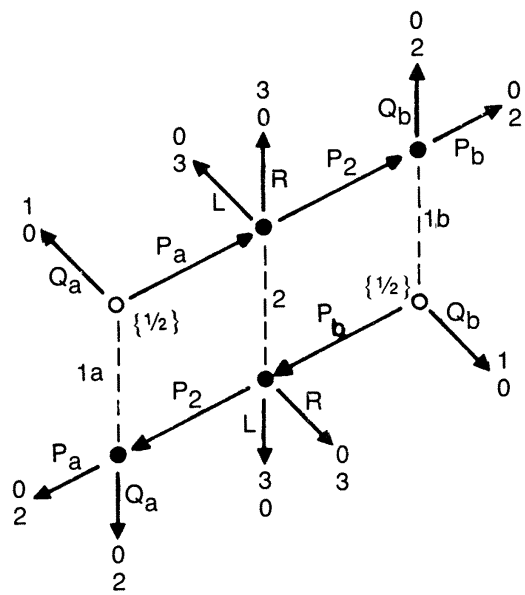

Example 2 from Kreps and Ramey (1987)

Perfect Bayesian Equilibrium

It may be difficult to verify that the belief system in an assessment satisfies the consistency condition of sequential equilibrium. This is because the construction of appropriate perturbed strategies, and the verification that the conditional probabilities implied by the perturbed strategies converge to he given belief system, may be a demanding task. Therefore, researchers have sought to replace the consistency condition by conditions that can be verified without having to construct perturbed strategies. The equilibrium notion of "perfect Bayesian equilibrium" is based on this idea. A variety of definitions of this concept have been proposed, and when the term is used in the literature, one can not always be sure about the precise meaning, although there is consensus about the broad outlines of the definition of perfect Bayesian equilibrium. Here, I shall review the definition of perfect Bayesian equilibrium that was suggested by Fudenberg and Tirole in "Perfect Bayesian Equilibrium and Sequential Equilibrium," *Journal of Economic Theory* 53 (1991), 236-260, in Section 2 of that article.

Fudenberg and Tirole do not define perfect Bayesian equilibrium for all extensive games, but only for a subclass of such games, namely multi-stage games with observed actions and with incomplete information. This is the subclass of games to which the notion of perfect Bayesian equilibrium is most commonly applied. We begin by defining this subclass of extensive games. The set of players is: $N = \{1, 2, \ldots, n\}$. Each player i has a finite set of possible types, which we denote by Θ_i. There is a common prior ρ on $\Theta = \times_{i \in N} \Theta_i$. We assume that types are independent, so that ρ is the product of its marginals ρ_i on Θ_i. Types are drawn before players make any choices, and players observe their own type but not other players' types. After each player has chosen their own type the game begins.

The game has stages $t = 1, 2, \ldots, T$. In each period players simultaneously choose actions. At the end of the period each player observes the outcome of that period. The set of actions available to a player is allowed to depend on the history of the game. We introduce the artificial initial history $h_0 \in H_0 = \{o\}$. The set of histories up to period $t-1$ is some set H^{t-1}. The set of actions available to player i in period t is a finite set $A_i^t(h^{t-1})$ where $h^{t-1} \in H^{t-1}$, and the set of histories up to period t is: $H^t = \{(h^{t-1}, (a_1^t, a_2^t, \ldots, a_n^t)) \,|\, h^{t-1} \in H^{t-1}$ and $a_i^t \in A_i^t(h^{t-1})$ for all $i \in N\}$. Player i's utility is a function $u_i : \Theta \times H^T \to \mathbb{R}$.

A behavior strategy b_i for player i is a sequence of functions b_i^t for $t = 1, 2, \ldots, T$ such that for each t the function b_i^t maps every pair $(\theta_i, h^{t-1}) \in \Theta_i \times H^{t-1}$ into a probability distribution $b_i^t(\theta_i, h^{t-1})$ over $A_i(h^{t-1})$. A belief system μ_i for player i is a sequence of functions μ_i^t for $t = 1, 2, \ldots, T$ such that for each t the function μ_i^t maps every $(\theta_i, h^{t-1}) \in \Theta_i \times H^{t-1}$ into a probability distribution $\mu_i^t(\theta_i, h^{t-1})$ over Θ_{-i}. Define b to be a vector of behavior strategies, one for each player, and define μ to be a list of belief systems, one for each player. An assessment is a pair (b, μ). We will define a perfect Bayesian equilibrium to be an assessment that satisfies a certain set of conditions. We introduce these conditions next.

The first condition for beliefs that Fudenberg and Tirole propose is that for every observed history there is a joint distribution of types, conditional on this history, such that each player's beliefs about the other players' types is derived from this joint distribution by conditioning on the player's own type.

Conditional Common Prior: For every $t = 1, 2, ..., T$ and for every $h^{t-1} \in H^{t-1}$ there is a probability distribution $\rho^t(h^t)$ on Θ such that for every player i the belief $\mu_i^t(\theta_i, h^{t-1})$ assigns to every $\theta_{-i} \in \Theta_{-i}$ the conditional probability $\rho^t(h^{t-1})[\theta_i, \theta_{-i}] / \sum_{\tilde{\theta}_{-i} \in \Theta_{-i}} \rho^t(h^{t-1})[\theta_i, \tilde{\theta}_{-i}]$.

This assumption says that the commonly observed history of action leads all players to update the prior distribution of types ρ in the same way, and that each individual player's beliefs are then derived by conditioning the shared new distribution of types on the knowledge of their own type.

Sequential rationality now requires that for every initial history the strategies for the remainder of the game form a Bayesian Nash equilibrium given the conditional common prior at that history. Formally, we define sequential rationality as follows:

Sequential Rationality: For every $t = 1, 2, ..., T$, every $h^{t-1} \in H^{t-1}$, and every $i \in N$:

$$u_i(b|\theta_i, h^{t-1}, \rho^t(h^{t-1})) \geqslant u_i(b'_i, b_{-i}|\theta_i, h^{t-1}, \rho^t(h^{t-1})) \text{ for all behavior strategies } b'_i \text{ of player } i.$$

Here, notation such as $u_i(b|\theta_i, h^{t-1}, \rho^t(h^{t-1}))$ stands for the expected utility of player i if players choose the behavior strategy profile b, if player i's type is θ_i, conditional on history h^{t-1} being reached, and conditional on the probability distribution of types being $\rho^t(h^{t-1})$.

We are now going to add three further conditions for beliefs. The first condition is:

Conditional Independence: For every $t = 1, 2, \ldots, T$ and every $h^{t-1} \in H^{t-1}$ there are probability measures $\rho_i^t(h^{t-1})$ on Θ_i such that for every $\theta \in \Theta$:

$$\rho^t(h^{t-1})[\theta] = \prod_{i \in N} \rho_i^t(h^{t-1})[\theta_i].$$

The second condition is that beliefs should be updated using Bayes' rule whenever possible. It is important that this condition applies not only on the equilibrium path, but also off the equilibrium path. That is, if a history is realized that has zero probability according to the equilibrium strategies, then players form new beliefs, but they update these beliefs using Bayes rule if, from now on, they observe actions that receive positive probability under the equilibrium strategies.

Bayesian Updating Whenever Possible: For every $t = 1, 2, \ldots, T-1$, and every $h^{t-1} \in H^{t-1}$, consider an action profile a^t such that $(h^{t-1}, a^t) \in H^t$, and consider a player $i \in N$ such that there exists a $\tilde{\theta}_i \in \Theta_i$ with $\rho^t(h^{t-1})[\tilde{\theta}_i] \cdot b_i^t(\tilde{\theta}_i, h^{t-1})[a_i^t] > 0$. Then for every $\theta_i \in \Theta_i$:

$$\rho_i^{t+1}(h^{t-1}, a^t)[\theta_i] = \frac{\rho_i^t(h^{t-1})[\theta_i] \cdot b_i^t(\theta_i, h^{t-1})[a_i^t]}{\sum_{\tilde{\theta}_i \in \Theta_i} \left(\rho_i^{t-1}(h^{t-1})[\tilde{\theta}_i] \cdot b_i^t(\tilde{\theta}_i, h^{t-1})[a_i^t]\right)}.$$

The final condition is that beliefs about player i should only depend on player i's actions, and not on other player's actions. This is based on the idea that only player i's actions can signal anything that is relevant to player i's type. Other players cannot signal any information about player i's type because they do not know the type.

No Signaling What You Don't Know: For every $t = 1, 2, \ldots, T-1$, and all $(h^{t-1}, a^t), (h^{t-1}, \tilde{a}^t) \in H^t$, if $a_i^t = \tilde{a}_i^t$, then:

$$\rho_i(h^{t-1}, a^t)[\theta_i] = \rho_i(h^{t-1}, \tilde{a}^t)[\theta_i] \text{ for all } \theta_i \in \Theta_i.$$

We can now define:

Definition 59. *An assessment (b, μ) is a perfect Bayesian equilibrium if it satisfies sequential rationality, and satisfies the conditions (i) Conditional Common Prior, (ii) Conditional Independence, (iii) Bayesian Updating Whenever Possible, and (iv) No Signaling What You Don't know.*

Sometimes, the concept of "weak perfect Bayesian equilibrium" is used. This concept only requires sequential rationality, and Bayesian updating on the equilibrium path. It omits the other conditions in the definition of a perfect Bayesian equilibrium.

Fudenberg and Tirole (1991) proved the following result:

Proposition 27. *Suppose that each player has only two types: $|\Theta_i = 2|$ for all $i \in N$. Then the set of perfect Bayesian equilibria and the set of sequential equilibria are the same. The same conclusion is true if there are only two stages: $T = 2$.*

They also demonstrated by an example that, if the assumptions of this proposition are not satisfied, not every perfect Bayesian equilibrium needs to be sequential. Here is their example: Suppose some player i has three possible types: $\Theta_i = \{\theta_i^1, \theta_i^2, \theta_i^3\}$ and that, after some history, players' beliefs about this player's type put probability 1 on type θ_i^1. Suppose after this history, player i with type θ_i^2 is supposed to choose a_i whereas player i with type θ_i^3 is supposed to choose $\hat{a}_i$. The above definition of perfect Bayesian equilibrium does not impose any constraints on how players' beliefs are updated when player i is observed to choose a_i or $\hat{a}_i$.

Consider the following beliefs: the other players believe with probability 1 that player i is of type θ_i^3 if she chooses a_i, and they believe with probability 1 that player i is of type θ_i^2 if she chooses $\hat{a}_i$. Although these beliefs are not ruled out by perfect Bayesian equilibrium, they are not consistent. To see this let k be the index of perturbations that justify consistent beliefs. Let $\mu_i^{1,k}, \mu_i^{2,k}$, and $\mu_i^{3,k}$ be the conditional probabilities with which player i is of types $\theta_i^1, \theta_i^2, \theta_i^3$. Let $\varepsilon^{1,k}$ and $\hat{\varepsilon}^{1,k}$ be the probabilities with which player i choose a_i and $\hat{a}_i$ if she is of type θ^1. Define $\varepsilon^{2,k}, \varepsilon^{3,k}, \hat{\varepsilon}^{2,k}$, and $\hat{\varepsilon}^{3,k}$ analogously. Note that $\lim_{k\to\infty} \varepsilon^{2,k} = \lim_{k\to\infty} \hat{\varepsilon}^{3,k} = 1$ and $\lim_{k\to\infty} \hat{\varepsilon}^{2,k} = \lim_{k\to\infty} \varepsilon^{3,k} = 0$. Now suppose that action $\hat{a}_i$ is observed. The conditional probability of player i being of type θ_i^3 is:

$$\frac{\mu_i^{3,k}\varepsilon^{3,k}}{\mu_i^{1,k}\varepsilon^{1,k} + \mu_i^{2,k}\varepsilon^{2,k} + \mu_i^{3,k}\varepsilon^{3,k}}.$$

This can be re-written as:

$$\frac{1}{\frac{\mu_i^{1,k}\varepsilon^{1,k}}{\mu_i^{3,k}\varepsilon^{3,k}} + \frac{\mu_i^{2,k}\varepsilon^{2,k}}{\mu_i^{3,k}\varepsilon^{3,k}} + 1}$$

For the given beliefs to be consistent, this fraction would have to converge to 1 as k tends to infinity. This requires that the first two terms in the denominator converge to zero as k tends to infinity. Consider the second term. Note that the ratio $\varepsilon^{2,k}/\varepsilon^{3,k}$ tends to infinity as k tends to infinity. Therefore, the second term converging to zero requires that:

$$\lim_{k\to\infty} \frac{\mu_i^{2,k}}{\mu_i^{3,k}} = 0.$$

An analogous argument shows that the belief that player i is of type θ_i^2 if she is observed to choose $\hat{a}_i$ requires:

$$\lim_{k\to\infty} \frac{\mu_i^{3,k}}{\mu_i^{2,k}} = 0.$$

But the last two equalities contradict each other.

Fudenberg and Tirole then modified their definition of perfect Bayesian equilibrium with the objective of obtaining full equivalence of perfect Bayesian equilibrium and sequential equilibrium. It has, however, turned out that this project is more complicated than it appeared to be in Fudenberg and Tirole's original work. We therefore do not consider here alternative notions of perfect Bayesian equilibrium.

References

[1] Martin J. Osborne and Ariel Rubinstein, *A Course in Game Theory*, Cambridge and London: MIT Press, 1994.

[2] Kim C. Border, *Separating Hyperplane Theorems*, course notes available at: http://www.hss.caltech.edu/∼kcb/Notes/SeparatingHyperplane.pdf.

[3] Kim C. Border, *Alternative Linear Inequalities*, course notes available at: http://www.hss.caltech.edu/∼kcb/Notes/Alternative.pdf

[4] Rangarajan K. Sundaram, *A First Course in Optimization Theory*, Cambridge et. al.: Cambridge University Press, 1996.