


Advanced Microeconomics – Game Theory

1 – Preliminaries

This is a game theory course, and standard microeconomic knowledge is considered a prerequisite. In order to make sure that everybody understands the concepts used, the standard microeconomic model of consumer choice has to be recalled. The model boils down to the following picture familiar to economic students:



Picture explains how a rationally behaving consumer chooses between two goods: X_1 and X_2 . It is assumed that preferences with respect to these goods are reflected by utilities, that is some real functions. It would be too complicated to draw everything in a two-dimensional plane, so utilities are represented by so-called "indifference curves", that is lines linking points that give the consumer the same utility. Three such indifference curves – I1, I2, and I3 – are drafted above. I3 represents a higher utility than I2, and I2 represents a higher utility than I1. Hence a rationally behaving consumer would prefer to consume any combination of goods 1 and 2 lying on I3 rather than on I2, and anything lying on I2 rather than on I1.

But the consumer cannot afford any combination of goods (economists call these combinations "bundles"). He (or she) is constrained by the amount of available money. If the prices of goods 1 and 2 are p_1 and p_2 , respectively, the bundle (x_1, x_2) costs $p_1x_1 + p_2x_2$. This expenditure cannot be higher than the money available, m . The inequality $p_1x_1 + p_2x_2 \leq m$ is called "budgetary constraint" (BC in the picture above). Prices p_1 and p_2 determine the slope of BC. If they are equal to each other then BC has a (negative) slope of 45° . If p_1 is higher than p_2 then BC is steeper. If p_1 is lower than p_2 then BC is flatter. The proportion of prices determines the slope of BC, and the amount of money available – m – determines whether BC is farther or closer to the origin. If m is higher, then BC is farther from the origin which means that a consumer can afford more.

There are three indifference curves drawn in the picture: I1, I2, and I3. Bundles lying on I3 are more preferred by the consumer than those lying on I2 and I1. However he (or she) does not have money to afford anything that lays on I3. Only I2 and I1 are consistent with BC. Many bundles lying along I1 is affordable, but if the consumer is rational, then he (or she) will prefer to buy a bundle lying on I2, since this is the highest indifference curve intersecting

with BC. The combination (x_1^*, x_2^*) is the optimum choice. Students who recall the high school mathematics note that this optimum choice is where the BC is tangent to an indifference curve.

Picture above illustrates how a rational consumer chooses between two goods. This logic extends to the general choice problem when the number of goods is higher than two. In our classes we will make use of the following terminology.

Consumption bundle is a collection of goods $\mathbf{x}=(x_1, \dots, x_N)$ taken out of a consumption set X : $\mathbf{x} \in X \subset \mathbb{R}^N$. Normally it is also assumed that $\mathbf{x} \geq \mathbf{0}$. Let $\mathbf{p}=(p_1, \dots, p_N)$ be the vector of prices of goods from the consumer's bundle and let $m > 0$ be the consumer's income (money to be spent on the bundle). The consumer's budgetary constraint is then defined as: $p_1 x_1 + \dots + p_N x_N \leq m$, and $p_1 x_1 + \dots + p_N x_N = m$ is called budget line. The set $\{\mathbf{x} \in X: p_1 x_1 + \dots + p_N x_N \leq m\}$ is called budget set.

"Utility", analysed in economics in order to simplify relationships between bundles, was introduced in the 19th century. In mathematical language its definition reads: a utility function $u: X \rightarrow \mathbb{R}$ representing \geq is any function such that $\forall \mathbf{x}, \mathbf{y} \in X [\mathbf{x} \geq \mathbf{y} \Leftrightarrow u(\mathbf{x}) \geq u(\mathbf{y})]$. Preference relation (\geq) uses the same symbol as arithmetic relation (greater than or equal to). This, however, shall not lead to any misunderstandings, since it will always be clear from the context whether the formula " $\mathbf{x} \geq \mathbf{y}$ " means " \mathbf{x} is preferred over \mathbf{y} " or " \mathbf{x} is greater than or equal to \mathbf{y} ". If \mathbf{x} and \mathbf{y} are consumption alternatives (bundles), i.e. $\mathbf{x}, \mathbf{y} \in X$ then the formula reads " \mathbf{x} is preferred over \mathbf{y} ", but if \mathbf{x} and \mathbf{y} are numbers then the formula reads " \mathbf{x} is greater than or equal to \mathbf{y} ".

Utility Maximization Problem (UMP) is solving the problem of consumer's choice under budgetary constraint:

$$\text{Max}_{\mathbf{x} \in X} \{u(\mathbf{x}): p_1 x_1 + \dots + p_N x_N \leq m\}.$$

Several theorems reveal characteristics of the UMP.

Theorem 1: If $u: X \rightarrow \mathbb{R}$ is a continuous function, and all prices are positive ($p_i > 0$ for $i=1, \dots, N$) then UMP has a solution $(x_1^*, \dots, x_N^*) = \mathbf{x}^*(\mathbf{p}, m)$.

Theorem 2: In the theorem 1 above, if u is strictly concave then the solution is unique.

Theorem 3 (Kuhn-Tucker conditions):

In the theorem 1 above, under differentiability assumptions, a solution of the UMP satisfies:

$$\exists \lambda \geq 0 \quad \forall i=1, \dots, N [\partial u(\mathbf{x}^*) / \partial x_i \leq \lambda p_i \text{ \& } x_i^* > 0 \Rightarrow \partial u(\mathbf{x}^*) / \partial x_i = \lambda p_i].$$

Theorem 4 (Corollary): Theorem 3 implies that for a solution of the UMP such that $\forall i=1, \dots, N [x_i^* > 0]$ (called internal solution) the following equalities hold:

$$\partial u(\mathbf{x}^*) / \partial x_i : \partial u(\mathbf{x}^*) / \partial x_j = p_i : p_j$$

for any $i, j=1, \dots, N$.

Theorems 1-4 contain the basic microeconomic facts we are going to refer to in the class. The standard game theory material starts now. Definitions will be called D.1, D.2, and so on, and theorems will be called T.1, T.2, and so on.

The most important extension of the microeconomic consumer choice model outlined above is to interpret bundles as lotteries. It is easy to make this step if you imagine that whatever you do does not have a fully predictable outcome. If you buy a bottle of milk, you expect to enjoy a nice drink, but the milk can turn out to be bad, and you are surprised unpleasantly. If you decide to walk through a park in order to enjoy a nice weather, you can be attacked by a thief or by a drug addict and deprived of your wallet. If you kiss a frog, the frog may turn out to be a beautiful princess who wishes to marry you (perhaps in this case boys will be surprised pleasantly). These observations can be formalised using some mathematical definitions and theorems. Some of them are easy to prove (you are encouraged to check it). If they are very difficult, they will not be proved here; students are encouraged to look into more advanced game theory textbooks.

(D.1) Elements of the set X can be interpreted as lotteries $L=(p_1,...,p_N)$, where $p_1+...+p_N=1$ (non-negative $p_1,...,p_N$ can be interpreted as probabilities); \mathcal{L} is the set of such lotteries; their outcomes – numbered $1,...,N$ – are predetermined.

Convexity is a key assumption referred to in economic analyses. Your high school mathematics used this concept as well. In economics jargon a convex combination means simply that you take a sum of something with (positive) weights that add up to 1. Formally, a convex combination of $y_1,...,y_k$ is $\alpha_1 y_1 + ... + \alpha_k y_k$, where $\alpha_1 + ... + \alpha_k = 1$.

I try to be fair with respects to students, but I heard that some professors treat exams as lotteries: for instance, they toss a coin or throw a dice and based on the outcome they either pass a student or fail. Let us assume that lottery $L=(1/6, 5/6)$, and $L'=(1/2, 1/2)$. The first outcome is "fail", and the second outcome is "pass". L can be interpreted as throwing dice: if it is "1" you fail, and if it is "2", or "3", or "4", or "5" or "6" you pass. L' can be interpreted as tossing a coin: if it is "head" you fail, and if it is "tail" you pass. Of course, students prefer L over L' . How about a convex combination of L and L' with weights $\alpha_1=3/4$ and $\alpha_2=1/4$? The probability of failure is $1/4$ (because $3/4 \cdot 1/6 + 1/4 \cdot 1/2 = 1/4$), and the probability of passing is $3/4$ (because $3/4 \cdot 5/6 + 1/4 \cdot 1/2 = 3/4$). This combination is better than L' but worse than L for a typical student (I asked you how did you feel about these lotteries and most of you preferred – not surprisingly – the dice rather than coin).

The following theorem formalises the concept of convex combinations of lotteries.

(T.1) A convex combination of lotteries is also a lottery (with probabilities calculated as convex combinations of probabilities from original lotteries). Easy to prove.

The next theorem makes use of the continuity concept. This is one of the most important mathematical concepts. There are two definitions of continuity: by Cauchy and by Heine. The first one uses the "delta-epsilon" terminology ("for every $\epsilon > 0$ there exists $\delta > 0$ such that..."), while the second one uses the sequence terminology ("every infinite sequence $q_1, q_2, ...$ has a limit such that..."). Both definitions are equivalent and they can be used to describe the continuity of preferences. This concept is extended from standard (deterministic) bundles to lotteries.

(T.2) Preferences \geq are continuous on \mathcal{L} , if for all $L, L', L'' \in \mathcal{L}$ the following sets are closed:

$$\{\alpha \in [0,1]: \alpha L + (1-\alpha)L' \geq L''\},$$

and

$$\{\alpha \in [0,1]: L'' \geq \alpha L + (1-\alpha)L'\}.$$

Easy to prove

Note:

Continuity of preferences is defined in the "standard" (deterministic) consumer theory in a similar way.

There is one more important concept used in microeconomics. This is independence understood as an indifference with respect to adding the same lottery. If adding a third lottery changes our preferences with respect to the first and the second one, we say that they were not independent. Formally it can be stated as follows:

(D.2) Preferences \geq satisfy the independence axiom in \mathcal{L} , if for every $L, L', L'' \in \mathcal{L}$ and every number $\alpha \in (0,1)$ the relation $L \geq L'$ is satisfied if and only if

$$\alpha L + (1-\alpha)L'' \geq \alpha L' + (1-\alpha)L''.$$

We are now ready to introduce the main concept applied in consumer choice theory with respect to lotteries.

(D.3) The von Neumann-Morgenstern (vNM) expected utility function, U:

U: has the "expected utility form", when

$$\exists u_1, \dots, u_N \in \mathfrak{R} \quad \forall L = (p_1, \dots, p_N) \in \mathcal{L} \quad [U(L) = u_1 p_1 + \dots + u_N p_N].$$

It is assumed in this definition that a consumer can measure his (or her) preferences with respect to a lottery in the following way. Every outcome of the lottery (numbered 1, 2, ..., N) is characterised by the utility u_1, u_2, \dots, u_N . The utility of the lottery (defined by probabilities p_1, \dots, p_N that these outcomes happen) is then the average (expected) value of these numbers.

We usually think of lotteries as something uncertain. How about a lottery $L = (1, \dots, 0)$? The first outcome happens for sure, and the other ones do not happen at all. Is this a lottery? Despite doubts we may have, this is a lottery from a formal point of view (since there is no requirement for probabilities to be less than one or more than zero). But if a probability of something is one, then we know that this will happen for sure. Mathematicians call such lotteries "degenerated" ones.

(D.4) Bernoulli utilities

Numbers u_i from D.3 can be interpreted as utilities of "degenerated lotteries" $L^1 = (1, 0, \dots, 0), \dots, L^N = (0, \dots, 0, 1)$.

The following theorem says that the "expected utility form" is satisfied when convex combinations of lotteries are taken into account.

(T.3) A utility function $U: \mathcal{L} \rightarrow \mathfrak{R}$ has the "expected utility form" if and only if

$$\forall K=1,2,\dots \quad \forall L_1,\dots,L_K \in \mathcal{L} \quad \forall \alpha_1,\dots,\alpha_K > 0 \quad [\alpha_1 + \dots + \alpha_K = 1 \Rightarrow$$

$$U(\alpha_1 L_1 + \dots + \alpha_K L_K) = \alpha_1 U(L_1) + \dots + \alpha_K U(L_K).$$

Proof:

The theorem has the "if and only if" (equivalence) form, and it will be proved separately for "only if" (\Leftarrow) and "if" (\Rightarrow) parts.

\Leftarrow

Let $L = (p_1, \dots, p_N)$. We define degenerated lotteries L^1, \dots, L^N such that $L^i = (0, \dots, 0, 1, 0, \dots, 0)$; the i th probability is equal to 1. Then $L = p_1 L^1 + \dots + p_N L^N$ and $U(L) = U(p_1 L^1 + \dots + p_N L^N) = p_1 U(L^1) + \dots + p_N U(L^N) = p_1 u_1 + \dots + p_N u_N$, where the second to the last equality holds by the assumption.

\Rightarrow

Let us consider a combination of lotteries $(L_1, \dots, L_K; \alpha_1, \dots, \alpha_K)$, where $L_k = (p_1^k, \dots, p_N^k)$. Let $L' = \alpha_1 L_1 + \dots + \alpha_K L_K$. Hence it can be calculated that: $U(L') = U(\alpha_1 L_1 + \dots + \alpha_K L_K) = u_1(\alpha_1 p_1^1 + \dots + \alpha_K p_1^K) + \dots + u_N(\alpha_1 p_N^1 + \dots + \alpha_K p_N^K) = \alpha_1(u_1 p_1^1 + \dots + u_N p_N^1) + \dots + \alpha_K(u_1 p_1^K + \dots + u_N p_N^K) = \alpha_1 U(L_1) + \dots + \alpha_K U(L_K)$, where the second to the last equality follows from the assumption.

So far all theorems were fairly obvious and easy to prove. The next one is even more obvious, but very difficult to prove. It will be quoted here without a proof.

(T.4) The expected utility theorem

If a preference relation \geq in \mathcal{L} is rational (i.e. complete and transitive), and satisfies independence and continuity axioms then it can be represented by a vNM function, i.e. numbers u_n can be attributed to each outcome $n=1, \dots, N$ such that:

$$\forall L = (p_1, \dots, p_N), L' = (p'_1, \dots, p'_N) \in \mathcal{L} [L \geq L' \Leftrightarrow u_1 p_1 + \dots + u_N p_N \geq u_1 p'_1 + \dots + u_N p'_N].$$

From now on it will be assumed that all lotteries have monetary payoffs. A formal definition reads:

(D.5) Lottery with monetary payoffs

The lottery has monetary payments x_1, \dots, x_N , and Bernoulli utilities are a function $u: \mathfrak{R} \rightarrow \mathfrak{R}$ of these payments: $u(x_1), \dots, u(x_N)$.

The von Neumann-Morgenstern theory of expected utility (vNM) has been convincing to a lot of people, and from the very beginning (i.e. from the middle of the 20th century) it has become a standard reference for decision making in the presence of uncertainty. Nevertheless, from the very beginning too, several paradoxes were demonstrated in order to reveal that it does not explain everything; it fails to explain how people choose in some circumstances.

The first warning was formulated by Maurice Allais in 1953. He analysed lotteries with the following (monetary) outcomes:

- $x_1 = 0$,
- $x_2 = 1$ million USD, and
- $x_3 = 5$ million USD.

It is important that the first outcome is zero, and the second and third ones represent large amounts of money. He ran two experiments. In the first one he asked people to choose between two lotteries defined as $L_1=(0,1,0)$ and $L_2=(0.01,0.89,0.1)$. Most people preferred L_1 . Then he asked the same people to choose between other two lotteries: $L_3=(0.89,0.11,0)$ and $L_4=(0.9,0,0.1)$. Most people preferred L_4 . Allais' experiments were replicated many times over the last 67 years – always with the same conclusion: if confronted with L_1 and L_2 people choose L_1 , and if confronted with L_3 and L_4 they choose L_4 .

These results are hardly surprising. In the first lottery (which happens to be a degenerated one) people are offered 1 million USD for sure. The second one gives a higher reward on average; 1.39 million USD rather than 1 million USD. But there is a small probability (just 1%) to get nothing. The small risk of getting nothing *vis-à-vis* getting 1 million USD for sure makes the first lottery more attractive. Also the second experiment has an easy intuitive interpretation. The lottery L_4 differs from L_3 by two things: the zero amount is to be received either with 89% probability (in L_3), or 90% probability (in L_4). These numbers are almost the same; it is very difficult to see the difference between 89% and 90%. Another difference between L_3 and L_4 is that the former attaches zero probability to the amount of 5 million USD, while the latter attaches zero probability to 1 million USD. In other words, in L_4 one can win 5 million USD, and in L_3 one can win 1 million USD; probabilities are almost the same (10%, and 11% respectively). Hence it is not surprising that almost everybody prefers L_4 over L_3 . The paradox discovered by Allais is that people who make such choices do not comply with vNM, as the following theorem states.

Before proving the theorem, let me note, that one of the students observed (correctly) that our attitude towards L_1 and L_2 should take into account risk aversion. Indeed, our preference for L_1 rather than L_2 (despite the fact that L_2 has a higher expected payoff) reflects the fact that L_1 is risk-free, whereas L_2 is not. Yet there is nothing paradoxical in our risk aversion (next lecture will address this problem in greater detail). The paradox is when you compare the results from the first experiment (L_1 versus L_2) and from the second one (L_3 versus L_4).

Theorem

People who choose L_1 in the first experiment and L_4 in the second one do not comply with the vNM theory.

Proof:

If the vNM theory was followed, then some Bernoulli's utilities would have been applied:

- $u(x_1)=u_1$,
- $u(x_2)=u_2$, and
- $u(x_3)=u_3$.

Please note that we do not know these numbers, and the proof works irrespective of what they are. The outcome of the first experiment implies that:

$$u_2 > 0.01u_1 + 0.89u_2 + 0.1u_3.$$

The outcome of the second one implies that:

$$0.9u_1 + 0.1u_3 > 0.89u_1 + 0.11u_2.$$

These two inequalities contradict each other, since the second one can be rewritten as:

$$0.01u_1 + 0.1u_3 > 0.11u_2, \text{ and consequently}$$

$$0.01u_1 + 0.1u_3 > u_2 - 0.89u_2, \text{ or } 0.01u_1 + 0.89u_2 + 0.1u_3 > u_2$$

which contradicts the first one.

Another warning was provided by Mark Machina. As before, let there be three outcomes of lotteries. This time it is important that the third one is much more attractive than the first two ones. In some textbooks these outcomes are identified as providing a luxurious trip to Venice for two people (x_3), and watching a movie of someone else's trip to Venice (x_2). The first outcome does not need any explanation.

- $x_1=0$,
- $x_2=10$ USD, and
- $x_3=10,000$ USD.

It is obvious that $u(x_1) < u(x_2) < u(x_3)$; hence $L_1 \leq L_2 \leq L_3$, where $L_1=(1,0,0)$, $L_2=(0,1,0)$, and $L_3=(0,0,1)$. Thus, by the independence axiom, the lottery $0.01L_2+0.99L_3$ should be preferred over $0.01L_1+0.99L_3$. And yet experiments show that most people choose otherwise. The mechanism of their choice is perhaps a psychological one. In both lotteries, $0.01L_2+0.99L_3$ and $0.01L_1+0.99L_3$, you win a trip to Venice almost surely (with the same very high probability of 99%). However in the first case if you lose, you get nothing, while in the second case you get the video of what you lost. You are so angry, that you prefer to get nothing rather than a dismal surrogate. The conclusion is a negative one for the vNM theory once again.

One of the students suggested why people prefer winning nothing rather than 10 dollars. Some of us think in terms "everything" or "nothing". From this point of view, the lottery L_2 does not exist (it is neither "everything" nor "nothing"). Consequently, we ignore it. This may provide some explanation for how we behave, but violating the "independence of irrelevant alternatives" is clear. When comparing L_1 and L_2 , we should disregard L_3 , but we do not. We are so angry about not going to Venice, that – contrary to the vNM theory – we prefer 0 than 10.

Conclusion: The expected utility theory may be insufficient to model people's behaviour in some applications.

Questions and answers to lecture 1

1.1 What is the slope of the budget line (see picture on page 1) if one of the goods does not cost anything (say, $p_1=0$)?

It is either horizontal (when $p_1=0$) or vertical (when $p_2=0$). If $p_1=0$ and $p_2>0$ then the only expenditure the consumer needs to make is to buy the second good. If he (or she) spends all the money m , he (or she) can buy m/p_2 units of the second good. The horizontal line with such a coordinate is the budget line (any amount of the first good can be combined with m/p_2 units of the second good. If $p_1>0$ and $p_2=0$ it is the other way around. The only expenditure the consumer needs to make is to buy the first good. If he (or she) spends all the money m , he (or she) can buy m/p_1 units of the first good. The vertical line with such a coordinate is the budget line (any amount of the second good can be combined with m/p_1 units of the first good. If both prices are positive then the budget line determines a triangle (like in the picture on page 1).

1.2 Rationality of preferences is defined by two axioms: completeness ($\forall x, y \in X [x \leq y \text{ or } y \leq x]$) and transitivity ($\forall x, y, z \in X [(x \leq y \text{ \& } y \leq z) \Rightarrow x \leq z]$). Some textbooks add the third axiom called reflexivity ($\forall x \in X [x \leq x]$). Is it necessary to add it?

No. Reflexivity is implied by completeness. To see this, please look at the completeness axiom: $\forall x, y \in X [x \leq y \text{ or } y \leq x]$. If you substitute x for y , it will read: $\forall x \in X [x \leq x \text{ or } x \leq x]$. The alternative is true if and if one of its components is true. Thus either $x \leq x$ (the first component) or $x \leq x$ (the second component). But they say exactly the same. Hence it can be concluded that $\forall x \in X [x \leq x]$.

1.3 One of the assumptions of theorem 4 reads $\forall i=1, \dots, N [x_i^* > 0]$ (the solution is an internal one). What happens if this assumption is violated?

If the solution is not an internal one (for one coordinate i , $x_i^* = 0$), the proportion of prices does not have to be the same as the proportion of marginal utilities: the equality

$$\partial u(\mathbf{x}^*) / \partial x_i : \partial u(\mathbf{x}^*) / \partial x_j = p_i : p_j$$

is not necessarily true.

1.4 Can indifference curves cross each other?

No. If they have one point of intersection they must be the same. To see this, let us assume that I_1 and I_2 are two indifference curves and $x \in I_1$, and $x \in I_2$ (they have at least this common point). In other words, I_1 consists of all bundles y such that y is indifferent with x , and I_2 consists of all bundles y such that y is indifferent with x . Hence they must be the same.

1.5 Can vNM theory be applied to so-called ordinal utility?

No. The definition of utility refers to one condition only: $u(x) \leq u(y)$ if and only if $x \leq y$. The utility function u is called "ordinal" (indicating ordering only) if no additional assumptions are adopted. If $u(x) = 2u(y)$ (the utility enjoyed by consuming x is twice as high as when consuming y) one cannot say that "the bundle x is preferred twice as much as the bundle y "; one can only say that the bundle x is more preferred than the bundle y . If some additional assumptions are adopted about the function u (u becomes a "cardinal" utility), one can make statements not only about the ordering of bundles, but also about the strength of preferences. Adding utilities and multiplying them by some numbers becomes justified. In the vNM theory utilities, are multiplied and added (in order to calculate average values). Thus these numbers have to be understood as cardinal utilities.

1.6 Do lotteries exhaust uncertainty linked to non-deterministic consumption bundles?

No. The definition of a lottery refers to probabilities that some outcomes happen. If one cannot calculate probabilities of some outcomes, the uncertainty cannot be assumed to resemble a lottery.

1.7 Is the expected payoff of a lottery a sufficient indicator of its utility for a consumer?

No. Not only the average, but also the distribution of specific outcomes can be relevant for our attitudes with respect to lotteries. The next lecture (GT-2) will address these problems in more detail.

1.8 Can Bernoulli utilities be defined for lotteries with non-monetary outcomes?

Yes. In the course of the lecture, we emphasise lotteries with monetary outcomes (monetary valuation is important in economics), but the definition does not require that outcomes have to be monetary. The definition of Bernoulli utilities states that these are numbers from the equation $U(L) = u_1p_1 + \dots + u_Np_N$, and if degenerated lotteries are considered $L^1 = (1, 0, \dots, 0)$, ..., $L^N = (0, \dots, 0, 1)$, then $U(L^i) = u_i$. There is no requirement that in the definition $u_i = u(x_i)$ x_i must be denominated in money.

2 – Risk aversion

Our attitudes towards lotteries are affected not only by expected utilities, but also by a specific distribution of various outcomes. Let me define the following lottery. You win 10 € with the probability 50%, or you lose 8 € with the probability of 50%. On average you win 1 €, so some of us will probably agree to take part in such a lottery. Yet some of us will not. Even though the expected payoff of 1 € provides an incentive to take part, there is a probability of 50%, that a participant loses and has to pay 8 €. If a prospective loss of 8 € is more meaningful than a prospective gain of 10 €, then the lottery is not attractive.

Let me define another lottery. You win 8 € with the probability of 50%, or you lose 10 € with the probability of 50%. On average you lose 1 €, so probably many of us will refuse to take part in such a lottery. Yet for some of us the lottery can be attractive. If someone considers winning 8 € more meaningful than losing 10 €, the lottery is attractive.

Finally, let me define yet another lottery. You either win 10 € or lose 10 € with the same probability of 50%. On average you receive zero, that is you neither win nor lose. Some people will be indifferent with respect to participating in such a lottery: they can either take part or not, and this makes no difference for them. In the class today some students refused to take part in such a lottery. Those who refused to take part in first lottery will probably refuse to take part in this one too. Those who wanted to take part in the second lottery will probably want to take part in this one too. Our attitudes depend on how we look at risk. The following definition formalises this idea.

(D.6) Risk aversion and risk neutrality implied by Bernoulli utilities

Aversion: $\forall L = (p_1, \dots, p_N) \in L [u(x_1)p_1 + \dots + u(x_N)p_N \leq u(x_1p_1 + \dots + x_Np_N)]$,

Neutrality: $\forall L = (p_1, \dots, p_N) \in L [u(x_1)p_1 + \dots + u(x_N)p_N = u(x_1p_1 + \dots + x_Np_N)]$.

The definition says that we need to compare two numbers: $u(x_1)p_1 + \dots + u(x_N)p_N$, and $u(x_1p_1 + \dots + x_Np_N)$. They look almost the same. The first one calculates the utility of every outcome separately ($u(x_i)$), and then takes the average of them. The second one calculates the average outcome $x_1p_1 + \dots + x_Np_N$ and then takes the utility of it. The difference is thus what comes first, and what comes after it; in the definition D.6 we take the utilities first and then


their average on left. On the right, we take the average first, and then calculate its utility. In other words, for risk neutral people the sequence does not make any difference. For risk averse people, the average utility of outcomes can turn out to be lower than the utility of the average outcome. People for whom the inequality reads $u(x_1)p_1 + \dots + u(x_N)p_N \geq u(x_1p_1 + \dots + x_Np_N)$ are called risk loving or risk prone.

This terminology allows us to characterise choices made with respect to the third lottery (with outcomes -10€ and +10€). Risk neutral consumers are indifferent (they can take part or not). Those who are risk averse will avoid such a lottery, and those who are risk loving may consider this lottery as attractive.

The same idea is captured by the following definition.

(D.7) Certainty equivalent of the lottery $L=(p_1, \dots, p_N)$ is a number $c(L, u)$ such that $u(c(L, u)) = u(x_1)p_1 + \dots + u(x_N)p_N$

Once again let us look at examples of lotteries we referred to before. The first one was defined in the following way. You win 10 € with the probability 50%, or you lose 8 € with the probability of 50% (on average you win 1 €). Let us assume that prospective participants are offered cash which makes them indifferent with respect to participating or not participating in such a lottery. This is certainty equivalent. For a risk neutral person this will be the amount of money giving him (or her) the same utility as 1 € (the average outcome of the lottery). For a risk averse person (who did not want to participate in this lottery) the certainty equivalent is apparently less than that (less than 1 €). For a risk loving person it is the other way around: more than 1 €.



Picture above illustrates how a risk averse and a risk prone person feels. The graph on the left is characteristic for a risk averse consumer. The utility of loss of -10 is more meaningful than the utility of gain of 10:

$$|u(-10)| > u(10).$$

The graph on the right is characteristic for a risk prone consumer. The utility of loss of -10 is less meaningful than the utility of gain of 10:

$$|u(-10)| < u(10).$$

We find very often that consumers are risk averse. This is how many people feel. Such people should not agree to take part in lotteries which – on average – make them lose more than gain. Yet this is not how they choose sometimes. In many countries popular lotteries announce that

what they pay in prizes is less than what they collect for tickets. In other words, the average gain is negative (people pay more for tickets than what they receive in prizes). It can be argued that those who participate are not rational, and they do not understand probabilities and expected values. Of course some of them may behave irrationally. Nevertheless some of them are rational and consider themselves risk averse. Once again, economists have to admit that the vNM theory does not explain fully how we behave.

Picture above can be summarised in the following theorem (easy to prove):

(T.5) The following conditions are equivalent:

1. A consumer is risk averse
2. Function u is concave
3. $\forall L=(p_1, \dots, p_N) \in \mathcal{L}$ [$c(L, u) \leq x_1 p_1 + \dots + x_N p_N$]

Assuming that $u(\cdot)$ is a twice differentiable Bernoulli utility function of money, coefficients of risk aversion can be defined. One of them is called Arrow-Pratt coefficient of absolute risk aversion, and the other one – Arrow-Pratt coefficient of relative risk aversion. Their definitions read:

(D.8) Arrow-Pratt coefficient of absolute risk aversion

$$r_A(x, u) = -u''(x)/u'(x).$$

Please note that r_A is a positive number (assuming that $u' > 0$ and $u'' < 0$; as in the left picture on page 10; in the right picture the coefficient is negative).

(D.9) Arrow-Pratt coefficient of relative risk aversion

$$r_R(x, u) = -xu''(x)/u'(x).$$

Please note that $r_R(x, u) = xr_A(x, u)$, and $r_A(x, u) = r_R(x, u)/x$. The corollary is that a risk averse individual with fixed (constant with respect to x) coefficient of absolute risk aversion will reveal an increasing (with respect to x) coefficient of relative risk aversion.

The Arrow-Pratt coefficient of absolute risk aversion lets compare risk aversion across individuals.

(T.6) Let u_1, u_2 be Bernoulli utility functions characterizing two individuals. Then the following conditions are equivalent:

1. $r_A(x, u_2) \geq r_A(x, u_1)$ for every $x \in X$
2. There exists an increasing concave function $\psi: \mathcal{R} \rightarrow \mathcal{R}$ such that $\psi(u_1(x)) = u_2(x)$ for every $x \in X$
3. $c(L, u_2) \leq c(L, u_1)$ for every $L \in \mathcal{L}$.

Proof:

There will be only proof of (1) \leftrightarrow (2). The proof of (2) \leftrightarrow (3) is more difficult (students can look into more advanced game theory textbooks). It will be assumed that $u_1', u_2' > 0$ (otherwise Arrow-Pratt coefficients cannot be defined).

The most tricky step is to observe that there exists an increasing function $\psi: \mathcal{R} \rightarrow \mathcal{R}$ such that $\psi(u_1(x)) = u_2(x)$ for every $x \in X$ (in fact, there also exists an increasing function $\psi: \mathcal{R} \rightarrow \mathcal{R}$ such

that $\psi(u_2(x))=u_1(x)$ for every $x \in X$, but we will use the first statement only). The existence of such a function results from the fact that u_2 and u_1 are increasing functions (as utilities). As both u_2 and u_1 are differentiable, it can be assumed that ψ is differentiable too.

Let us differentiate the equation $\psi(u_1(x))=u_2(x)$, yielding $u_2'=\psi'(u_1)u_1'$; and then again, yielding $u_2''=\psi''(u_1)u_1'+\psi'(u_1)u_1''$. Dividing u_2'' into u_2' we get (after cancellations): $u_2''/u_2'=\psi''(u_1)/\psi'(u_1)+u_1''/u_1'$. In other words, $r_A(x,u_2)=-\psi''(u_1)/\psi'(u_1)+r_A(x,u_1)$. Thus we will have $r_A(x,u_2) \geq r_A(x,u_1)$ if and only if $-\psi''(u_1)/\psi'(u_1) \geq 0$, i.e. if ψ is concave.

The last idea to be introduced is a fair price of a lottery ticket. If we refer to the three examples introduced at the beginning of this lecture, we can ask the question how much should a lottery cost to break even (i.e. to leave – on average – its operator without a profit, but also without a loss). All three lotteries had just two outcomes each, and these outcomes happened with the probability of 50%. The first lottery had outcomes -8 and 10; the second one: -10 and 8; and the third one: -10 and 10. The expected outcome of the first lottery was 1; the expected outcome of the second one: -1; and the expected outcome of the third one: 0. These numbers are fair prices of the respective lottery tickets. In the first case the operator charges the price of 1, and – on average – has to pay the prize of 1. In the second case the operator has to pay each prospective participant 1 in order to receive – on average – 1. In the third case participation in the lottery is for free, and nobody – on average – loses and gains. This concept is formalised in the following definition.

(D.10) A fair price of a lottery ticket:

The price $t(L)$ equal to the expected payoff, i.e. if $L=(p_1, \dots, p_N)$, and the corresponding payoffs are x_1, \dots, x_N , then $t(L) = x_1 p_1 + \dots + x_N p_N$.

Corollary:

For a risk neutral consumer $c(L,u)=t(L)$.

Please note that the definition of a fair price of a lottery ticket depends on probabilities, not on how consumers perceive risks. In contrast, the certainty equivalent depends not only on what the lottery offers, but also on how a consumer perceives its risk.

Questions and answers to lecture 2

2.1 In microeconomic theory consumers may be characterised by different attitudes towards risk (for instance they can be risk averse, or risk prone). Yet it is assumed that firms are risk neutral. Why?

According to economic theories, consumers maximise utility while firms maximise profits. Of course, we can provide examples that sometimes they do not, but – as a rule – we assume an optimising behaviour, because otherwise little can be predicted. Firms' profits are denominated in money. Thus if a firm makes twice as high a profit, it doubled its objective. At the same time, if a consumer gets twice as much money, his (or her) utility goes up, although not necessarily in the same proportion (see graphs on page 10). The fact that firms measure things in monetary units, and consumers in utility units results in firms' risk neutrality; the loss and gain have the same meaning for a firm, while not necessarily they mean the same for a consumer.

2.2 Can various consumers have different certainty equivalents of the same lottery?

Yes. You can refer to the example used in the class (the lottery giving -10 or 10 with same probability of 50%). For a risk neutral consumer its certainty equivalent will be the amount of money giving the same utility as 0. For a risk averse person it will be less, and for a risk prone person – more. Please note that we look at the same lottery in each case.

2.3 Has the utility function of a given consumer to be always convex or always concave?

No. It may have a not constant curvature. In fact, one of the explanations of the Allais paradox is that consumers who choose among lotteries in a different way than vNM theory predicts, may have a different attitude towards risk when they lose and different when they win.

2.4 Does the Arrow-Pratt coefficient of absolute risk aversion change if the utility function is transformed linearly?

No. Let us assume that $u_2(x) = \alpha u_1(x)$. Then $r_A(x, u_2) = -u_2''(x)/u_2'(x) = -\alpha u_1''(x)/(\alpha u_1'(x)) = -u_1''(x)/u_1'(x) = r_A(x, u_1)$. In fact the same result holds if $u_2(x) = \alpha u_1(x) + \beta$ (the utility function can be multiplied by a constant number, and shifted upwards or downwards, and the coefficient does not change).

2.5 Let us assume that a consumer's utility function is defined as $u(x) = (x+100)^{1/2} - 10$ for $x > -100$ (it resembles the left picture on page 10). Please calculate the Arrow-Pratt coefficients of absolute and relative risk aversion.

$u'(x) = (x+100)^{-1/2}/2$, and $u''(x) = -(x+100)^{-3/2}/4$. Hence $r_A(x, u) = (x+100)^{-1/2}/2$, and $r_R(x, u) = x(x+100)^{-1/2}/2$. The absolute coefficient decreases when the consumer becomes wealthier (i.e. when x increases). To see this, please check that $r'_A(x, u) < 0$. The interpretation of the relative coefficient is more difficult. If one confines to $x > 0$ (please disregard the case of "negative" wealth which can be interpreted as indebtedness), then $r'_R(x, u) > 0$. To see this, one needs to calculate $(x(x+100)^{-1/2}/2)' = ((1(x+100)^{1/2}) - x(x+100)^{-1/2})/(4(x+100)) = ((x+100)(x+100)^{1/2} - x)/(4(x+100)) = 100/(4(x+100)) > 0$. Thus the relative coefficient increases.

2.6 Let there be two consumers: one characterised by the utility function like in 2.5, that is $u_1(x) = (x+100)^{1/2} - 10$; and the second by the utility function $u_2(x) = \ln(x+1)$. Let us also assume that $x \geq 0$ (negative wealth is disregarded). Both utility functions look like the left picture on page 10. Which of the consumers has higher certainty equivalents of lotteries?

By the theorem T.6, $c(L, u_2) \leq c(L, u_1)$ for every $L \in \mathcal{L}$ is equivalent to $r_A(x, u_2) \geq r_A(x, u_1)$ for every $x \in X$ (or *vice versa*, if the second consumer is less risk averse than the first one). We will demonstrate that $r_A(x, u_2) > r_A(x, u_1)$ holds for some $x \in X$, and $r_A(x, u_2) < r_A(x, u_1)$ holds for some other $x \in X$; thus $c(L, u_2) \leq c(L, u_1)$ does not hold for every lottery. In 2.5 it was calculated that $r_A(x, u_1) = (x+100)^{-1/2}/2$. It is easy to calculate that $r_A(x, u_2) = 1/(x+1)$. $r_A(0, u_1) = 1/20$, and $r_A(0, u_2) = 1$. Thus $r_A(0, u_1) < r_A(0, u_2)$, but we will demonstrate that this inequality is reverted for sufficiently high x . To this end, we will solve the equation $r_A(x, u_1) = r_A(x, u_2)$, that is

$(x+100)^{-1/2}/2=1/(x+1)$, or – equivalently – $2(x+100)^{1/2}=x+1$, and finally $4(x+100)=(x+1)^2$. This is a quadratic equation $x^2-2x-399=0$. Its positive root is $x^*=21$. For $0<x<21$, $r_A(x,u_1)<r_A(x,u_2)$, and for $x>21$, $r_A(x,u_1)>r_A(x,u_2)$.

2.7 There is a lottery $L=(1/3,1/3,1/3)$ with three outcomes: -9 €, 0 €, +12 €. What is a fair price of this lottery ticket?

According to the definition D.10, $t(L)=1$ €.

2.8 The price of the ticket to the lottery from 2.7 is 2 €. Who is likely to participate in it?

For a risk neutral consumer, the acceptable price of the ticket is 1 €. Thus a risk neutral consumer is not likely to participate in such a lottery. Risk averse consumers are even less likely to participate in it. The only consumers who may buy such a lottery ticket are those risk prone. But they need to love risk sufficiently strongly. Their certainty equivalent must satisfy $c(L,u) \geq x_1p_1+\dots+x_Np_N=-1$ (please note that the outcomes of the lottery are -11,-2,+10 rather than -9, 0,12, because from every prize one needs to subtract the price of the ticket, i.e. 2).

3 – Return-Risk Comparisons

This lecture is devoted to comparing various lotteries in order to judge which one is "riskier". For instance, there may be two lotteries with the same expected payoff, $L_1=(1/2,1/2)$ with outcomes -1 €, and 1 €, and the other one $L_2=(1/3,1/3,1/3)$ with outcomes -1 €, 0, +1 €. In a sense they can be considered equivalent, because the expected outcome is the same, that is 0. However, if one wants to minimise the risk of losing 1 €, they are not equivalent; the first one has a larger probability of the loss (1/2 instead of 1/3). So-called stochastic dominance formalises this idea. But before we introduce the relevant definitions, we have to adopt an ordering convention which is fairly simple. The set of outcomes is finite: $\{x_1,\dots,x_N\}=X$; and the outcomes are ordered from the lowest one to the highest one: $x_1<\dots<x_N$. They are strictly increasing, because if some outcomes were the same, they could have been combined (with probabilities simply added).

(D.11) First-Order Stochastic Dominance

Let $L=(p_1,\dots,p_N)$, $L'=(p'_1,\dots,p'_N) \in L$. L first-order stochastically dominates L' if and only if, for every nondecreasing function $u:X \rightarrow \mathbb{R}$: $p_1u(x_1)+\dots+p_Nu(x_N) \geq p'_1u(x_1)+\dots+p'_Nu(x_N)$.

The expression "for every nondecreasing function u " can be understood by economists as "for every consumer having a utility function u " (since the only condition imposed on utilities is that they do not decrease; if a consumer is offered more money, then the utility may not increase, but it will not decrease). The following example explains the idea of first order stochastic dominance.

Example I

Let there be 6 outcomes of throwing dice: $x_1=1$, $x_2=2$, $x_3=3$, $x_4=4$, $x_5=5$, $x_6=6$. Let L be a lottery that we are used to, i.e. $L=(1/6, 1/6, 1/6, 1/6, 1/6, 1/6)$. Its average outcome is $E(x) = 3.5$. We will define alternative lotteries L' (with the same outcomes), and analyse how they

are perceived by consumers with different attitudes towards risk (demonstrated by different utility functions u).

Now we will look at how various consumer types are likely to see these lotteries. Examples of alternative utility functions (corresponding to alternative attitudes towards risk; please note that they are nondecreasing, i.e. they satisfy D.11) are:

- $u(x) = 0 = \text{const}$ (a consumer ignoring money)
- $u(x) = x$ (a risk-neutral consumer)
- $u(x) = x^2$ (a risk-loving consumer)
- $u(x) = x^{1/2}$ (a risk-averse consumer)

The values of the left-hand-side (LHS) of the inequality in D.11 (as if calculated by these four types of consumers) therefore are:

- $(0+0+0+0+0+0)/6 = 0$
- $(1+2+3+4+5+6)/6 = 3.5$
- $(1+4+9+16+25+36)/6 = 91/6 = 15.17$
- $(1+1.41+1.73+2+2.24+2.45)/6 = 1.81$

Now we define 3 alternative lotteries L' that will be compared with L . The first and the second are degenerated ones ($E(x)$ stands for the expected value):

- $L' = (1, 0, 0, 0, 0, 0)$ with $E(x) = 1$
- $L' = (0, 0, 0, 0, 0, 1)$ with $E(x) = 6$
- $L' = (0, 0, \frac{1}{2}, \frac{1}{2}, 0, 0)$ with $E(x) = 3.5$

For each of the three lotteries, we will calculate the value of the right-hand-side (RHS) of the inequality (in D.11) for the utility functions defined earlier.

For the lottery $L' = (1, 0, 0, 0, 0, 0)$:

- $u(x)=0=\text{const} \Rightarrow \text{RHS} = 0$
- $u(x)=x \Rightarrow \text{RHS} = 1$
- $u(x)=x^2 \Rightarrow \text{RHS} = 1$
- $u(x)=x^{1/2} \Rightarrow \text{RHS} = 1$

For the lottery $L' = (0, 0, 0, 0, 0, 1)$:

- $u(x)=0=\text{const} \Rightarrow \text{RHS} = 0$
- $u(x)=x \Rightarrow \text{RHS} = 6$
- $u(x)=x^2 \Rightarrow \text{RHS} = 36$
- $u(x)=x^{1/2} \Rightarrow \text{RHS} = 2.45$

For the lottery $L' = (0, 0, \frac{1}{2}, \frac{1}{2}, 0, 0)$:

- $u(x)=0=\text{const} \Rightarrow \text{RHS} = 0$
- $u(x)=x \Rightarrow \text{RHS} = 3.5$
- $u(x)=x^2 \Rightarrow \text{RHS} = 12.5$

- $u(x)=x^{1/2} \Rightarrow \text{RHS} = 3.73$

By summing up these calculations, we can see that for the lotteries:

- $L=(1/6,1/6,1/6,1/6,1/6,1/6)$ and $L'=(1, 0, 0, 0, 0, 0)$: $\text{LHS} \geq \text{RHS}$ for all four utility function analysed
- $L=(1/6,1/6,1/6,1/6,1/6,1/6)$ and $L'=(0, 0, 0, 0, 0, 1)$: the inequality does not hold for utility functions #2, #3, and #4
- $L=(1/6,1/6,1/6,1/6,1/6,1/6)$ and $L'=(0, 0, 1/2, 1/2, 0, 0)$: the inequality does not hold for utility function #4

The example does not prove the dominance, but it demonstrates that the second and third lotteries cannot satisfy the definition D.11, since at least for some utility functions the inequality does not hold. The first lottery seems different. The following theorem explains why it is.

(T.7)

L first-order stochastically dominates L' if and only if

$$p_1 \leq p'_1, p_1 + p_2 \leq p'_1 + p'_2, \dots, p_1 + \dots + p_{N-1} \leq p'_1 + \dots + p'_{N-1}.$$

Please note that the next (unwritten) inequality would read: $p_1 + \dots + p_N \leq p'_1 + \dots + p'_N$. This is always satisfied obviously, since both at the left hand side and at the right hand side there is 1 (the sum of all probabilities). The proof of the theorem given by most game theory textbooks – using advanced mathematical analysis methods – requires an assumption that the lotteries are infinite (contrary to what we assumed earlier, that the number of lottery outcomes is finite). Yet one implication ('only if') does not require any sophisticated mathematics. We will confine the proof to this direction.

Proof \Rightarrow only:

From the definition of probabilities we obviously know that $p_1 + \dots + p_N = p'_1 + \dots + p'_N = 1$. Thus, for any $i=2, \dots, N-1$ we have $p_1 + \dots + p_{i-1} = 1 - (p_i + \dots + p_N)$ and $p_1 + \dots + p_{N-1} = 1 - p_N$ (this is for L), and $p'_1 + \dots + p'_{i-1} = 1 - (p'_i + \dots + p'_N)$ and $p'_1 + \dots + p'_{N-1} = 1 - p'_N$ (this is for L'). This makes the proof easier. The argument will be based on the fact that any implication $p \Rightarrow q$ is equivalent to $\neg q \Rightarrow \neg p$.

Let us assume that $p_1 + \dots + p_{i-1} > p'_1 + \dots + p'_{i-1}$ (for some i), i.e. $p_i + \dots + p_N < p'_i + \dots + p'_N$. A function $u: X \rightarrow \mathbb{R}$ defined as $u(x)=0$ for $x \leq x_{i-1}$ and $u(x)=1$ for $x \geq x_i$ is non-decreasing. At the same time $p_1 u(x_1) + \dots + p_N u(x_N) < p'_1 u(x_1) + \dots + p'_N u(x_N)$ (because $p_1 u(x_1) + \dots + p_{i-1} u(x_{i-1}) = p'_1 u(x_1) + \dots + p'_{i-1} u(x_{i-1}) = 0$) which means that L does not first-order stochastically dominate L' which ends the proof.

Now we can confirm our earlier conjecture that the lottery L first-order stochastically dominates over $L'=(1,0,0,0,0,0)$. As well we can see why L did not first-order stochastically dominate over other lotteries.

Corollary

In the Example I only the first lottery – $L'=(1,0,0,0,0,0)$ – satisfies the inequalities from T.7 ($1/6 < 1$, $1/3 < 1$, $1/2 < 1$, $2/3 < 1$, and $5/6 < 1$). Thus L first-order stochastically dominates over L' . For other pairs of lotteries (L and L'), some inequalities do not hold.

The concept of first-order stochastic dominance referred to comparing lotteries irrespective of what attitudes with respect to risk people may have (we assumed that certain inequalities must hold for any nondecreasing function u – interpreted as a utility function). In many circumstances, people are assumed to be risk averse. Therefore an interesting extension of the concept of first-order stochastic dominance is to see what happens if inequalities hold for every nondecreasing concave function u – interpreted as a utility function of a risk averse person; see T.5).

(D.12) Second-Order Stochastic Dominance

For any two lotteries L and L' with the same mean μ (i.e. when $p_1x_1 + \dots + p_Nx_N = p'_1x_1 + \dots + p'_Kx_K = \mu$) L second-order stochastically dominates L' (or L is less risky than L') if for every nondecreasing concave function $u: X \rightarrow \mathbb{R}$:

$$p_1u(x_1) + \dots + p_Nu(x_N) \geq p'_1u(x_1) + \dots + p'_Ku(x_K)$$

The condition $p_1x_1 + \dots + p_Nx_N = p'_1x_1 + \dots + p'_Kx_K = \mu$ is so important that there is a separate name for lotteries that comply with it.

(D.13) Mean-preserving spread


Let L be a lottery with mean μ , and let L_x'' be a family of lotteries indexed with x (outcomes of the lottery L) such that means (expected values) of each lottery L_x'' are equal to zero. Thus $L = (p_1, \dots, p_N)$, $p_1x_1 + \dots + p_Nx_N = \mu$, and for every outcome x , $L_x'' = (p^x_1, \dots, p^x_K)$, $p^x_1z^x_1 + \dots + p^x_Kz^x_K = 0$. Let L' be a compound lottery with L_x'' superimposed on L , where its outcomes are $x+z$ with appropriate probabilities. The mean of L' is thus μ . L' is called a mean preserving spread of L .

It turns out that second-order stochastic dominance is equivalent to being a mean-preserving spread, but the theorem is a difficult one. It will be quoted here without a proof.

(T.8) If L and L' have the same mean then the following statements are equivalent:

1. L second-order stochastically dominates L'
2. L' is a mean-preserving spread of L

Instead we will look at examples of a mean-preserving spread.



Example II


Let $L = (1/2, 1/2)$, with $x_1 = 2$, and $x_2 = 3$; and let $L_2'' = (1/2, 1/2)$, with $z^2_1 = -1$, $z^2_2 = 1$, and let $L_3'' = (1/2, 1/2)$, with $z^3_1 = -1$, $z^3_2 = 1$. Thus we obtain the following compound lottery:

$L=(1/4,1/4,1/4,1/4)$ with outcomes: 1, 2, 3, and 4. As L_2'' and L_3'' have mean zero, L' is a mean preserving spread of L .

The example demonstrates that a mean preserving spread may have more outcomes than the original lottery (four instead of two). But this is not a necessary result. To see why, please consider a lottery defined in the next example.

Example III

Let $L=(1/3,1/3,1/3)$ have outcomes $x_1=1$, $x_2=2$, and $x_3=3$. Its mean is $\mu=2$. The superimposed lotteries are defined as follows. $L_1''=L_3''=(1)$, $z^1=z^3=0$ (if 1 or 3 result from L , nothing is added). $L_2''=(1/2,1/2)$, with $z^2_1=-1$, and $z^2_2=1$. If the outcome of the first lottery is 2, then the superimposed lottery adds 1 to, or subtracts 1 from, the original result. In the end the compound lottery will have two results only: 1 and 3 (formally it will have three results – 1, 2, and 3 – but one of them will have the probability of 0). Of course the mean of this compound lottery is 2 as for the original one.



T.8 states that the original lottery second-order stochastically dominates over its spread, that is, it is "less risky" (preferred by people who are risk averse). Do you see the higher riskiness of mean preserving spreads? The expected outcome (the mean) must be the same by definition. Let us look at the example II. The expected utility of L is $U(L)=u(2)/2+u(3)/2$. The expected utility of L' is $U(L)=u(1)/4+u(2)/4+u(3)/4+u(4)/4$. The weight of the utility enjoyed from 2 and 3 is reduced, but utilities enjoyed from 1 and 4 are added. A question arises whether these compensations increase the total or decrease it. If the consumer was risk averse then $u(4)-u(3)$ must be lower than $u(2)-u(1)$.

The example III illustrates the same pattern, even though one can think that less outcomes (i.e. more "predictability") reduces the risk. The expected utility of the original lottery is $U(L)=u(1)/3+u(2)/3+u(3)/3$, and the expected utility of the compound lottery is $U(L')=u(1)/2+u(3)/2$. Here $u(2)$ disappears, and $u(1)$ and $u(3)$ are given more weight. Again, by the concavity of u , the "loss" from $u(2)$ to $u(1)$ is more meaningful than the "gain" from $u(2)$ to $u(3)$. Despite what people may think, reducing the number of outcomes, may be not attractive for a risk averse person.

Questions and answers to lecture 3

3.1 If L turns out more attractive for people with utility functions representing risk aversion and risk loving than L', does this prove that L first-order stochastically dominates over L'?

No. The definition D.11 requires that L must be more attractive than L' for any utility function, not just for two functions – one illustrating risk aversion, and one illustrating risk loving. In example I utility functions #3 and #4 (representing risk loving and risk aversion) confirmed the relevant inequalities indicating that L was more attractive than L', but this was not sufficient to prove the first-order stochastic dominance.

3.2 First-order stochastic dominance indicates that one lottery is preferred over another by every rational consumer. Can the price of lottery tickets change this conclusion?

Yes. In the definition D.11 it is tacitly assumed that outcomes from the lotteries are calculated in "net" terms, that is assuming that ticket prices are subtracted from prizes or (which implies the same) that ticket prices are zero ($t(L)=t(L')=0$). For instance, T.7 indicated that $L=(1/6,1/6,1/6,1/6,1/6,1/6)$ first-order stochastically dominated over $L'=(1,0,0,0,0,0)$. This is the case assuming that we do not have to pay for the participation. If the tickets to these lotteries were priced fairly – that is $t(L)=3.5$, and $t(L')=1$ – then the second lottery would turn out more attractive for some people.

3.3 First-order stochastic dominance implies the second-order stochastic dominance, but not conversely. Why?

The first-order stochastic dominance of L over L' means that for every utility function u the inequality holds: $p_1u(x_1)+...+p_Nu(x_N) \geq p'_1u(x_1)+...+p'_Nu(x_N)$. The second order stochastic dominance implies that $p_1u(x_1)+...+p_Nu(x_N) \geq p'_1u(x_1)+...+p'_Ku(x_K)$. One difference between these inequalities is the number of outcomes. This is not an important difference, because if $N \neq K$, then missing outcomes can be added either to the left hand side or to the right hand side multiplied by zero probabilities. The crucial difference between the two concepts is included in what we assume about the function u. In the first case u is any nondecreasing function. In the second case it is any concave nondecreasing function. The concavity assumption is missing in D.11, but it shows up in D.12. In other words, if something holds for any function, it also holds for concave functions, but not necessarily *vice versa*.

3.4 Does a mean preserving spread have a higher variance of outcomes than its original lottery?

Not necessarily. In example II the original lottery $L=(1/2,1/2)$, had two outcomes: 2, and 3. The compound lottery was $L'=(1/4,1/4,1/4,1/4)$ with four outcomes: 1, 2, 3, and 4. In either case the mean was 2.5. The variance of the first lottery was $\text{var}(L)=((-0.5)^2+0.5^2)/2=0.25$. The variance of the second lottery was $\text{var}(L')=(-1.5)^2+(-0.5)^2+0.5^2+1.5^2)/4=1.25$. Thus $\text{var}(L)<\text{var}(L')$. However the compound lottery can be exactly the same as the original one. Then it has the same variance, so a general conclusion that the variance increases is not justified. As the example III explained the increased riskiness of mean preserving spreads is not implied by the number of outcomes, but rather by their distribution.

4 – Strategic form games I

Previous classes reviewed basic microeconomic concepts applied in game theory. This one starts the "regular" material of the course. Games are used in order to analyse human behaviour found in situations when nobody can fully control the outcome; what happens results from independent decisions of several entities. The following definition captures this idea.

(D.14) Non-zero sum two-person game

A representation of a decision situation with a table of pairs of numbers (P_{ij}, D_{ij}) . Index $i=1, \dots, m$, where m is the number of strategies (decision variants) for the first player, and $j=1, \dots, n$, where n is the number of strategies (decision variants) for the second player. The numbers P_{ij} are payments to the first player, and D_{ij} – to the second player, if the first chose the i th strategy, and the second – the j th one. Payments can be interpreted as utilities. The table including pairs of numbers (P_{ij}, D_{ij}) is called payoff matrix.

The terminology has to be explained. The term 'two-person game' does not need any explanation. The D.14 can be generalized into n -person games. The notation becomes more complex, because simultaneous decisions of three or more players cannot be represented by a single matrix. The 'non-zero sum' expression allows for 'zero sum' games as well. 'Non-zero sum' refers to the fact that $P_{ij} + D_{ij}$ is not necessarily zero. If incidentally $P_{ij} + D_{ij} = 0$ (i.e. $P_{ij} = -D_{ij}$) the entire framework works, even though certain additional facts can be observed. The next definition introduces the concept of domination.

(D.15)

The strategy i_0 of the first player is called (strictly) dominant, if for any strategy i of the first player, and any strategy j of the second player, $P_{i_0j} > P_{ij}$; likewise, strategy i_0 is (strictly) dominated, if there exists strategy i such that for any strategy j , $P_{i_0j} < P_{ij}$. (analogously for strategies of the second player). Please note that this definition refers to strict inequality ($>$).

Example I

		Second	
		L	R
First	U	(1,3)	(4,1)
	D	(0,2)	(3,4)

U is (strictly) dominant for the First. D is not likely to be played. Therefore the game reduces to:

		Second	
		L	R
First	U	(1,3)	(4,1)

Now it turns out that R is (strictly) dominated for the Second player, and therefore it is not likely to be played. Thus the game is iteratively reduced to what you see on the next page.

In other words, the solution is (U,L) the First player will get 1, and the Second will get 3. This is what will happen if the players behave rationally. Please note that the same result would be

obtained if (strictly) dominated strategies of the Second player were eliminated first. If the domination is somewhat weaker (the next definition will explain this concept formally), the result may depend on the sequence.

		Second
		L
First	U	(1,3)

(D.16)

The strategy i_0 of the first player is called weakly dominant, if for any strategy i of the first player, and any strategy j of the second player, $P_{i_0j} \geq P_{ij}$ with strict inequality for at least one strategy of the second player; likewise, strategy i_0 is weakly dominated, if there exists strategy i such that for any strategy j , $P_{i_0j} \leq P_{ij}$ with strict inequality for at least one strategy of the second player (analogously for strategies of the second player). Please note that this definition refers to non-strict inequality (\geq).

Let us look at the following example.

Example II

		Second	
		L	R
First	U	(5,1)	(4,0)
	M	(6,0)	(3,1)
	D	(6,4)	(4,4)

Strategies U and M of the first player are weakly dominated by D. If U is eliminated, the game reduces to:

		Second	
		L	R
First	M	(6,0)	(3,1)
	D	(6,4)	(4,4)

Now the Second player can eliminate L as weakly dominated by R. Then the game reduces to:

		Second
		R
First	M	(3,1)
	D	(4,4)

Finally the First player can eliminate M, and the game reduces to:

		Second
		R
First	D	(4,4)

However strategies can be eliminated in a different order. If M is eliminated first, the game reduces to:

		Second	
		L	R
First	U	(5,1)	(4,0)
	D	(6,4)	(4,4)

Now R becomes weakly dominated, and the game reduces to:

		Second
		L
First	U	(5,1)
	D	(6,4)

The First player eliminates U now, and the game reduces to:

		Second
		L
First	D	(6,4)

Please note that this is a different outcome than in the previous elimination sequence. Weak dominance does not justify a reasonable iterative elimination procedure, since the choice of a specific sequence is arbitrary. In the first case we started with U, and in the second one we started with M.

We will define the most important concept in game theory, namely what we call now Nash Equilibrium – called also Nash strategy. John Nash (1928-2015; Nobel prize in 1994) invented a new type of equilibrium in 1950 (he did not call it "Nash Equilibrium"). A traditional equilibrium was understood as a situation when people do not have a motivation to move away. Nash changed this understanding slightly by assuming that nobody has a motivation to move away unilaterally. This slight change turned out to be extremely important, and forced economists to look at their discipline differently. The difference between moving away and moving away unilaterally is essential. Moving away unilaterally assumes, that other decision makers will keep their choices (and will not move). It precludes moving away jointly, even though this is what economists understood earlier when they argued that an equilibrium is synonymous with an optimum. Nash equilibrium does not have to be an optimum, as we will see soon.

(D.17) Nash strategy (equilibrium)

Any pair of strategies (i_0, j_0) such that $P_{i_0 j_0} = \max_i \{P_{ij_0}\}$ and $D_{i_0 j_0} = \max_j \{D_{i_0 j}\}$.

The following theorem links the concept of domination and Nash equilibrium. However the implication leads in one direction only (i.e. there may be Nash equilibria consisting of strategies that are not dominant).

(T.9) Corollary of D.15, D.16, and D.17

If players have (either strictly or weakly) dominant strategies, then their pair is a Nash equilibrium.

Proof:

If there is i_0 such that $P_{i_0j} \geq P_{ij}$ for any j , and there is j_0 such that $D_{ij_0} \geq D_{ij}$ for any i , then (i_0, j_0) is a Nash strategy. Please also note that if both inequalities are strict then the equilibrium is unique.

(T.10) Alternative definition of Nash equilibrium

If players are in a Nash equilibrium, then – if they wish to maximize their payoffs – neither has a motivation to unilaterally change his or her strategy

This alternative definition of Nash equilibrium is sometimes adopted as the main one. Indeed it is very convenient to apply, as you will see in a moment. Nevertheless only one definition can be adopted formally, and anything else needs to be proved to be equivalent. That is why the criterion T.10 is quoted as a theorem (even though its proof is a trivial one).

Now we will look at the best known game called "Prisoner's Dilemma". The story behind is that two criminals are caught at a minor offence – like stealing a purse from an old lady in a supermarket. They were caught by the local security, everything is recorded, so they cannot deny. You go to jail for 1 month for such an offence. However this is not the only crime they committed. Sometime earlier they committed a heavier crime – something you go to jail (when convicted) for 1 year (12 months). But they were not caught then and nobody saw them, so unless at least one of them confesses, they cannot be convicted. They promised each other not to confess about what they did.

The police interrogates them in two separate rooms. The police do not have any doubts about their minor offence (everything is recorded and certain), but whenever criminals are arrested, they are interrogated about all previous unsolved mysteries. When the police asked them "did you do it", they could deny without any hesitation many times. But when they are asked about what they actually did, how will they react? They did it, but should they confess? They promised each other to deny. Yet would you trust the word of honour given by a criminal? Each of the thieves has a dilemma. Indeed I promised my partner to deny, and he promised me the same. But shall I trust him? What if he breaks the promise and confesses?

Example (Prisoner's Dilemma)

		Second	
		C	D
First	C	(-12,-12)	(0,-18)
	D	(-18,0)	(-1,-1)

The numbers you see in the payoff matrix above refer to their predicament. "C" stands for "Confess", and "D" stands for "Deny". If both of them confess, they both go to jail for 12 months. If both of them deny then they go to jail for the minor offence they were caught at, and the major crime stays unsolved (their participation is undiscovered). If the first confesses but the second denies, then the first is freed and the second goes to jail for 18 months. The asymmetry is caused by the fact, that the police rewards the first criminal for his collaboration; both crimes are forgiven (the major crime and the minor offence). The second criminal is punished not only for the crime that he did but also for the fact that he lied to the police (by denying).

In order to identify a Nash equilibrium, payoffs given in every cell need to be analysed. Let us start with the (D,D) cell. Both have a motivation to unilaterally move away. If the first player switches from D to C, then he will be given 0 rather than -1. If the second player switches from D to C, then he will be given 0 rather than -1 too. Hence this is not a Nash equilibrium. If (C,D) is checked, then the first player does not have a motivation to unilaterally switch (he would be given -1 rather than 0), but the second player by switching from D to C will be given -12 rather than -18. Hence (C,D) is not a Nash equilibrium. Likewise (D,C) is not a Nash equilibrium. Even though the second player does not have a motivation to unilaterally switch (he would be given -1 rather than 0), the first player by switching from D to C will be given -12 rather than -18. The only outcome where neither of the players has a motivation to unilaterally move away (from C to D) is (C,C).

Nash equilibrium is thus in (C,C) which is the very worst outcome for both players; i.e. Nash equilibrium does not necessarily 'optimise' the global outcome (which – in this case – would be (D,D) giving the payoffs (-1,-1)). The 'Prisoner's Dilemma' demonstrates that Nash equilibrium is not an optimum, but in fact it can be the very worst outcome: 24 months in jail, *versus* 18 or 2 for other combinations.

That is why economists were shocked when John Nash introduced this strange equilibrium concept. Nevertheless the awareness grew gradually that this is a kind of equilibrium that was analysed by them earlier as well. For instance Antoine Augustin Cournot (1801–1877) studied oligopolistic competition and calculated what we now call 'Cournot equilibrium'. If oligopolists wanted to maximise their joint profit by charging the monopolistic price, they should calculate the joint supply corresponding to that price, and to allocate the total supply to participating firms. But having done this, every oligopolist would have a motivation to unilaterally cheat (by increasing its supply somewhat, spoiling the price a little bit, and enjoying a higher profit at the expense of others). If all of them do so, all will lose. Cournot calculated a solution such that nobody has a motivation to unilaterally move away. This is a Nash equilibrium. It also demonstrates the difference between moving away jointly, and moving away unilaterally (that is assuming that others will not move).

Many social arrangements turn out to be Nash equilibria that are not optima. For instance, Windows operating system is used by almost all students, but it has a number of shortcomings (and besides, it is expensive). Some time ago I contested it, using an alternative one. But my students complained that they could not read my electronic letters. And *vice versa* – I had problems with what my students sent me. Finally I gave up, and I switched to Windows; I do not have a motivation to unilaterally move away. Likewise, other people do not have a motivation to unilaterally move away. If everybody switched to something else, all could have gained, but this requires coordination that this lecture does not study. Windows is a Nash equilibrium – in a sense that nobody has a motivation to unilaterally move away – even though perhaps it is not optimal.

Questions and answers to lecture 4

4.1 Can a three-person game be described by payoff tables?

Yes, but not by a single payoff table. In the next lecture (GT-5) there will be an example of such a game; two players choose rows and columns, respectively, as in two-person games.

However, they do not know which payoff matrix is applied. The third player chooses the payoff matrix.

4.2 Can a zero sum game be considered a non-zero sum game as in D.14?

Yes. The payoff matrix of a zero sum game has the following form:

		Second		
		1	...	n
First	1	$(x_{11}, -x_{11})$...	$(x_{1n}, -x_{1n})$

	m	$(x_{m1}, -x_{m1})$...	$(x_{mn}, -x_{mn})$

Payments received by the first player and by the second player have the same absolute values, but opposite signs (hence their sums are zeroes in every cell).

4.3 Can a game have several Nash equilibria?

Yes. Please note that both solutions obtained by eliminating weakly dominated strategies in example II are Nash equilibria. Both (6,4) and (4,4) comply with D.17 (and T.10):

		Second	
		L	R
First	U	(5,1)	(4,0)
	M	(6,0)	(3,1)
	D	(6,4)	(4,4)

If (D,L) is chosen, then the First will get 6; otherwise the payoff will be 5 if he (or she) switches to M or 4 if he (or she) switches to U. The Second will get 4; he (or she) will get 1 if he (or she) switches to R. Similarly, if (D,R) is chosen, then the First will get 4; otherwise the payoff will be 3 if he (or she) switches to M or 5 if he (or she) switches to U. The Second will get 4 if he (or she) switches to L. Hence both (6,4) and (4,4) are Nash equilibria.

4.4 In prisoner's dilemma, why do not they take the joint decision to deny? If they deny, then the police cannot convict them.

They cannot agree on a decision, because they are being interrogated in two separate rooms. The First player does not know the decision taken by the Second one and *vice versa*. Professional interrogators do not reveal such an information. Even if they know that one confessed or denied they would not tell this to the other one. The criminals promised to deny, but they do not know if the denial is 100% sure.

4.5 Martin Shubik (an excellent economist and one of the founders of game theory) said that "you need to know John Nash in order to understand the Nash equilibrium". Please comment on this opinion.

This is not a nice opinion. Nash equilibrium is about taking decisions for your own sake. You check whether something can improve your situation, and you do not ask what would be the result for anybody else. In other words, Nash equilibrium is about egoistic decision making. You do not try to achieve a situation which can be considered an optimum for a wider group (the society at large, or perhaps a group consisting of yourself and the other player). According to Martin Shubik, John Nash took decisions with his own well-being in mind only. Have you seen the movie 'Beautiful mind'?

4.6 Why is the Cournot equilibrium considered a Nash equilibrium?

In the Cournot model, sellers decide about the supply they are going to offer. They would like to maximise their profits. They understand that the price will depend on the total supply available in the market. They understand that the more they offer, the lower the price. Thus – in order to maximise the profit – they would like to sell more, but not too much in order to enjoy a decent price. They cannot control what others will supply, but they assume that the others have similar objectives. Everybody makes certain assumptions regarding the supply offered by the others and adjusts their own supply in order to maximise the profit (please keep in mind that the price depends on the supply). If the decisions actually taken are consistent with what oligopolists assumed, they do not have a motivation to change these decisions. This is what Cournot supposed. But this is exactly what the Nash equilibrium states: nobody has a motivation to unilaterally move away. If the oligopolists colluded, they could have taken the advantage of the monopolistic price. But – as game theory asserts – they may not collude when they take decisions independently.

4.7 Please discuss whether the right hand side traffic can be interpreted as a Nash equilibrium.

This is a social convention. In some countries cars drive on the left and in some countries on the right, and both solutions are good, but once something has been agreed upon people should obey the rules they adopted. Actually some people say that the left hand side traffic – prevalent in Europe before Napoleon – is a more natural one (given the fact that most of us are right-handed, if you rode your horse on the left you could have your sword kept in the right hand ready to hit an enemy who approached you). Nevertheless, let us assume that we have the right hand side traffic as a rule, and in principle everybody drives on the right. Does anybody have an incentive to switch to the left? The price to be paid for such a diversion would be very high. Therefore, even though driving on the left is equally good (some people say that it is even better), we do not have a motivation to change unilaterally. That is why driving on the right can be interpreted as a Nash equilibrium.

5 – Strategic form games II

Nash equilibrium was the main concept of the previous lecture. We learnt how to identify Nash equilibria. But there are also games without them. Please check that the game below does not have a Nash equilibrium.

We have to check all the cells. (4,2) is not a Nash equilibrium, because player M has a motivation to unilaterally switch from Y to N. (-2,3) is not a Nash equilibrium, because player F has a motivation to unilaterally switch from Y to N. (-1,0) is not a Nash equilibrium, because M has a motivation to switch from N to Y. Finally (2,1) is not a Nash equilibrium,

because F has a motivation to switch from N to Y. All four cells were checked, and always one player had a motivation to unilaterally move away.

		M	
		Y	N
F	Y	(4,2)	(-2,3)
	N	(2,1)	(-1,0)

Strategies chosen by players resemble bundles chosen by consumers. Some bundles are preferred to others if they satisfy consumer's needs better. Likewise, some strategies are preferred to others if they provide players with higher payoffs. In the first lecture (GT-1) we observed that preferences with respect to lotteries take into account the fact that certain outcomes happen with certain probabilities. Also in games we can think of a probability that a player applies a given strategy with. The definitions below formalise this idea.

(D.18) Mixed strategies

Strategies defined so far are called *pure*. A game can be defined where pure strategies are selected by players randomly with certain probabilities $\pi=(\pi_1,\dots,\pi_m)$ and $\delta=(\delta_1,\dots,\delta_n)$, respectively (for the first and the second player), where $\pi_1,\dots,\pi_m\geq 0$, $\pi_1+\dots+\pi_m=1$ and $\delta_1,\dots,\delta_n\geq 0$, $\delta_1+\dots+\delta_n=1$. The pair (π,δ) is called a mixed strategy selection.

(D.19) Payoffs in games with mixed strategies

If the players select mixed strategies, then the payoffs are understood as expected payoffs from their pure strategies. In other words, the payoff for the first is $\sum_{ij}\pi_i\delta_jP_{ij}$, and for the second is $\sum_{ij}\pi_i\delta_jD_{ij}$.

Once again let us compare the concept of mixed strategies to lotteries. Lotteries with probabilities equal to 1 were called 'degenerate' ones. Formally they are lotteries, but their outcomes are certain rather than random. A 'traditional' game (with pure strategies) can be interpreted as a mixed-strategy game where probabilities are either 0 or 1 (they are 'degenerate'). In other words, everything we learnt about 'traditional' games applies to 'degenerate' games with mixed strategies directly. For non-degenerate games we need to introduce new definitions. In particular the Nash equilibrium can be defined for such games. Please check that it will coincide with D.17 for 'degenerate' games. A pair of mixed strategies (π^0,δ^0) is a Nash equilibrium, if

- $\sum_{ij}\pi_i^0\delta_j^0P_{ij}=\max_{\pi}\{\sum_{ij}\pi_i\delta_j^0P_{ij}\}$, and
- $\sum_{ij}\pi_i^0\delta_j^0D_{ij}=\max_{\delta}\{\sum_{ij}\pi_i^0\delta_jD_{ij}\}$.

It has the same economic interpretation like the Nash equilibrium in 'traditional' games: it is a pair of strategies such that nobody has a motivation to unilaterally move away. "Moving away" in this case means changing probabilities unilaterally, that is assuming that the other player sticks to the same probabilities as before. Unlike games with pure strategies, games with mixed strategies have always Nash equilibria, as the theorem below states.

(T.11)

For every non-zero sum two-person game there exists a Nash equilibrium in mixed strategies.

The proof can be derived from the Brouwer's fixed-point theorem, but it is difficult, and it will not be provided here. Instead we will check that the game this lecture started from (see page 26-27) has a Nash equilibrium in mixed strategies.

Example

If p is the probability of choosing Y for the player F, and q is the probability of choosing Y for the player M, then $(1/2, 1/2; 1/3, 2/3)$ can be proved to be the Nash equilibrium in mixed strategies (when $p=1/2$, and $q=1/3$ the players do not have a motivation to unilaterally change these probabilities).

You may check that if the player F switches from probabilities $(1/2, 1/2)$ while M sticks to $(1/3, 2/3)$, or if M switches from $(1/3, 2/3)$ while F sticks to $(1/2, 1/2)$ their average payoffs will not increase. In what follows we will see how these numbers can be calculated.

The expected payoff for F is

$$E(P) = 4pq - 2p(1-q) + 2(1-p)q - (1-p)(1-q) = 4pq - 2p + 2pq + 2q - 2pq - 1 + q + p - pq = -p + 3q + 3pq + 1.$$

The expected payoff for M is

$$E(D) = 2pq + 3p(1-q) + (1-p)q = 2pq + 3p - 3pq + q - pq = 3p + q - 2pq.$$

F would like to maximise $E(P)$, and M would like to maximise $E(D)$, therefore we need to solve system of equations

$$\partial E(P) / \partial p = 0,$$

$$\partial E(D) / \partial q = 0;$$

that is

$$-1 + 3q = 0; \quad q = 1/3,$$

$$1 - 2p = 0; \quad p = 1/2.$$

This method has to be used carefully. Please note that vanishing derivatives can indicate a minimum, a maximum or an inflexion point of a function. Higher derivatives need to be calculated in order to make sure that we found what we looked for. In this particular case the second derivatives vanish as well, so determining what we found is difficult.

In some game theory textbooks F stands for a female, M for a male, Y stands for "yes, I am loyal to my spouse", and N for "no, I am not loyal to my spouse". If payoffs resemble benefits from being loyal or not loyal, the probabilities calculated as Nash equilibrium explain why men and women may be characterised by different loyalty patterns.

The meaning of the word 'strategy' used by us was always obvious and meant 'choosing a decision' (in two-person games it was synonymous with choosing a row or a column; in mixed strategy games it was synonymous with choosing probabilities used to select pure strategies). In the following definition the word 'strategy' is used in an abstract (mathematical) way, although it is consistent with our earlier terminology.

(D.20) Game in a normal (condensed, strategic) form is $\Gamma_N = [I, S_1 \times \dots \times S_I, (u_1, \dots, u_I)]$, where

1. I – the number of players,
2. S_i – set of strategies of player i ($s_i \in S_i$),
3. $u_i(s_1, \dots, s_I)$ – a payoff function whose values can be interpreted as expected utilities (in the von Neumann-Morgenstern sense) of outcomes (perhaps probabilistic ones)

Note

D.20 generalizes D.14 by letting the number of players be I (instead of 2), introducing abstract "strategies" (instead of "rows" and "columns"), and utilities instead of payoffs.

(D.21) Notational convention: $s_{-i} = (s_1, \dots, s_{i-1}, s_{i+1}, \dots, s_I)$

For example s_{-4} stands for a vector of strategies for players 1, 2, 3, 5, 6, ..., I , that is all the players except the 4th one. Set side by side, s_i and s_{-i} list all strategies applied by the players. This notational convention proves useful whenever we want to look at one player separately and at the rest of them – jointly. The following definition explains the usefulness of the idea.

(D.22) Correlated equilibrium

Probability distribution p over the set of strategies $S = S_1 \times \dots \times S_I$ such that for every player i and every permutation $d_i : S_i \rightarrow S_i$ the following holds: $\sum_{s \in S} p(s) u_i(s_i, s_{-i}) \geq \sum_{s \in S} p(s) u_i(d_i(s_i), s_{-i})$.

The definition needs an explanation. The difference between the left hand side and the right hand side of the inequality is $d_i(s_i)$ on the right instead of s_i on the left. At the left there is the expected payoff resulting from the strategy s_i , and at the right – the expected payoff from other strategies; the strategy on the left should provide a better (\geq) outcome than anything else (this 'anything else' is hidden in d_i). Permutations mean simply reordering: $d_i(s_i)$ is another strategy of the i th player. Hence the inequality means that – given the probability distribution p – if the i th player applies his (or her) strategy s_i while other players apply s_{-i} , the expected outcome will be better. The following theorem is extremely easy to prove.

(T.12)

Nash equilibrium is a correlated equilibrium for a degenerate distribution p . More precisely: $p(i_0, j_0) = 1$, and for every $(i, j) \neq (i_0, j_0)$, $p(i, j) = 0$.

Several examples will clarify what the correlated equilibrium means. The first one – called coordination game – can be interpreted as an attempt to meet in a distant town. There are two places where the First and the Second can meet, but for whatever reason they cannot communicate directly (there are no mobile phones). If the First chooses U, and the Second chooses L, they meet. Likewise they succeed if the First chooses D and the Second chooses R. But if the First chooses U while the Second chooses R, or the First chooses D while the Second chooses L they miss each other. If the players choose their strategies randomly and independently (with probabilities 1/2), their average payoffs will be 0. In 50% of cases they will meet (so their payoffs will be 1), but in the other 50% of cases they will not (and their payoffs will be -1). If the First knew what the Second was going to choose or *vice versa*, there would be no confusion, but if they cannot communicate directly confusions happen in 50% of cases.

The concept of the correlated equilibrium asserts that players (even though they do not communicate directly) can observe some signals which make choosing some strategies more likely and they know their reaction patterns.

Example I

There are two Nash equilibria in the following game: (U,L), and (D,R), but if the players make uncorrelated choices they will end up in (U,R) and (D,L) cells sometimes:

	Second		
		L	R
First	U	(1,1)	(-1,-1)
	D	(-1,-1)	(1,1)

They can correlate their strategies, if they choose them based on a random but publicly observed signal s , such like, say, s =Heads or s =Tails; and if s =Heads then the First plays U and the Second plays L, while if s =Tails then the First plays D and the Second plays R. This is a correlated equilibrium with payoffs equal to 1 enjoyed by both players. Their choices are correlated in 100%, and they can avoid the 'bad' outcomes (-1,-1) always. This is not always possible, as the next example demonstrates.

Example II

In the following game

	Second		
		L	R
First	U	(5,1)	(0,0)
	D	(4,4)	(1,5)

there are two Nash equilibria: (U,L) and (D,R). If the players choose their strategies randomly and independently (with probabilities 1/2), then – on average – they will get 2.5 each. Their joint payoffs can be 6 (if they play (U,L) or (D,R)), 8 (if they play (D,L)), or 0 (if they play (U,R)). The best joint outcome is 8 when they play (D,L). But please note that no publicly available signal can motivate them to play (D,L), i.e. the combination yielding the maximum sum of payoffs. If the First knows that the Second will play L, then he (or she) will play U, and if the Second knows that the First will play D, then he (or she) will play R.

If there is no correlation, the expected joint payoff is 5. The maximum is 8, but it is impossible to achieve. Is there any correlation device allowing the players to enjoy better payoffs than 2.5 on average? Exercise E-5 attached to the overhead for this lecture develops a mechanism required.

Now we will analyse a three person game. It looks complicated, but the correlation mechanism is fairly simple. As noted in the earlier lecture (GT-4), three person games cannot be represented by single payoff matrices. Nevertheless in the answer to question 4.1 it was explained that two players can choose matrix rows and columns, respectively, as in two-person games. However, they do not know which payoff matrix is applied. The third player chooses the payoff matrix.

Example III

Let us consider the following 3-person game:

First chooses U or D, Second chooses L or R, Third chooses the payoff matrix A, B, or C. This game has the unique Nash equilibrium is (D,L,A) with payoffs 1 to each of the players. To see that (D,L,A) is a Nash equilibrium indeed, let us check that nobody has a motivation to

move away unilaterally. If the First was to switch from D to U, the payoff would drop from 1 to 0. If the Second was to switch from L to R, the payoff would drop from 1 to 0. If the Third was to switch from A to B or C, the payoff would drop from 1 to 0 too. It is easy to check that there are no other Nash equilibria. If you look at the other 11 cells, you will always find that somebody has a motivation to move away unilaterally.

	L	R
U	0,1,3	0,0,0
D	1,1,1	1,0,0

A

	L	R
U	2,2,2	0,0,0
D	2,2,0	2,2,2

B

	L	R
U	0,1,0	0,0,0
D	1,1,0	1,0,3

C

In the (D,L,A) cell players receive the payoff of 1 each – that is jointly 3. There are several other combinations with higher joint payoffs. In particular, there are two outcomes – (U,L,B) and (D,R,B) – which provide them with the payoff of 2 each. This is not a Nash equilibrium. If the Third knows that the First chooses U and the Second – L, then he (or she) will choose A. And if the Third knows that the First chooses D and the Second – R, then he (or she) will choose C.

Is it possible to design a correlation device which allows the players to enjoy (U,L,B) or (D,R,B)? The device is based on using information in a very careful way. It can be interpreted as tossing a coin. Its outcome is s =Heads or s =Tails, and the signal is revealed to First and Second, but not to the Third player. If it is Heads then the First chooses U and the Second chooses L. If it is Tails, they choose otherwise. A correlated equilibrium is then

$$(U/2, D/2, L/2, R/2, 0A, 1B, 0C)$$

with expected payoffs (2,2,2). The numbers are interpreted as in D.22. $U/2, D/2$ means that the First chooses U and D with 50% probability. $L/2, R/2$ means that the Second chooses L and R with 50% probability. $0A, 1B, 0C$ means that the Third never chooses A or C, and always chooses B. It works in the following way. If s =Heads then the First plays U, otherwise D. If s =Heads then the Second plays L, otherwise R. Please note that if the outcome of flipping the coin was revealed to the Third player, there would be no correlated equilibrium (as explained in the previous paragraph).

The somewhat strange corollary from this example is that less information rather than more can be beneficial. If the result of tossing a coin (s =Heads or s =Tails) was known to the third player, the joint payoff would be 4 rather than 6. People often believe that more information is better than less. This is not always true.

The last two lectures were filled with examples of Nash equilibria. The Nash equilibrium has a precise definition in the mathematical game theory. Nevertheless it turns out inspiring in social sciences. The condition that nobody has a motivation to unilaterally move away is found in many applications. In the previous lecture I offered an interpretation of the Windows operating system as a stable solution. In one of the questions (4.7) there was a discussion of whether the right hand side traffic can be considered an arbitrary yet useful social convention. In this lecture we looked at the problem of mixing strategies with fixed probabilities. In all these cases Nash equilibria could be interpreted as solutions making everybody feel that moving away unilaterally is not attractive. Thus they can be seen as stable social conventions.

In international negotiations analysts define so-called self-enforcing agreements. An agreement of this sort provides incentives for outsiders to join, and incentives for insiders to

stay. Montreal Protocol, i.e. a settlement to protect the global ozone layer, serves as the best known example of a self-enforcing agreement. Montreal Protocol states that a signatory cannot trade freons (ozone-depleting substances) with a non-signatory. Once USA joined the Protocol, and every country wants to sell freons to, or buy them from, USA, the agreement became self-enforcing. The participation turned out to be full, and no country has an incentive to move away unilaterally.

In life sciences Nash equilibrium proved useful in interpreting (and predicting) evolutionary stable behaviour. The last lecture (GT-15) will analyse this problem in greater detail. Here it is sufficient to indicate that behaviour of individuals may prove useful for their breeding success or not. Those who survive apparently behaved in a way which helped them to survive. Therefore they have no incentives to unilaterally move away from what their earlier generations practised which means that the behaviour was evolutionary stable.

Questions and answers to lecture 5

5.1 Why can pure strategies be interpreted as a special case of mixed strategies?

Because pure strategies can be considered mixed with degenerate probability distribution. A pure strategy i can be considered mixed with probabilities $(0, \dots, 0, 1, 0, \dots, 0)$, where strategies $1, \dots, i-1$, and $i+1, \dots, m$ are never applied (that is they are applied with the probability 0), and the i th strategy is applied for sure (that is with the probability 1).

5.2 Please try to calculate the Nash equilibrium in mixed strategies for the following game

		M	
		C	B
F	C	(3,1)	(0,0)
	B	(0,0)	(1,3)

This game is called "the battle of sexes" (but it has nothing to do with the game defined on page 26-27 of GT-5). The story is that two friends – a girl (F), and a boy (M) – would like to go out. The girl prefers a concert (C), and the boy prefers a baseball game (B). The girl enjoys the baseball game somewhat, and the boy enjoys the concert, but this is not what they prefer. Yet the worst outcome is when the girl insists on a concert, and the boy insists on the ball. Then they do not go together anywhere. The girl would like the boy to accompany her for the musical performance, and the boy wants the girl to accompany him for the sport event. It is easy to check that there are two Nash equilibria in pure strategies: (C,C), and (B,B). You can suspect that if they agree sometimes to what they do not prefer, both can be satisfied to some extent. If they choose randomly between C and B with 50% probabilities they will get 1 each (on average). If they choose either of the Nash equilibria all the time (that is they avoid (C,B), and (B,C)) they will get 2 each (on average). What will happen if they switch from C to B randomly, with different probabilities? Let us try to calculate the Nash equilibrium in mixed strategies.

By applying the method explained in the class, we can calculate that $E(F) = 3pq + (1-p)(1-q) = 3pq + 1 - p - q + pq = 4pq - p - q + 1$, and $E(M) = pq + 3(1-p)(1-q) = pq + 3 - 3p - 3q + pq = 4pq - 3p - 3q + 3$. By taking partial derivatives and equating them to zero, we get the system of equations

$$4q - 1 = 0, \text{ and } 4p - 3 = 0 \text{ which yields } p = 3/4, \text{ and } q = 1/4.$$

In other words, the girl should randomly choose C in 75% cases, and the boy should randomly choose C in 25% of cases. The solution has an intuitive appeal, because – after all – the girl prefers the concert, and the boy prefers the ball. But if they stick to such mixed strategies they will achieve even less (3/4 each on average) than by switching from C to B with 50% probabilities.

If there is a Nash equilibrium in pure strategies then calculating it in mixed strategies does not make sense (even without a Nash equilibrium in pure strategies a caution is necessary, as explained on page 28 in GT-5). Besides, please note that Nash equilibria do not have to maximise joint payoffs (prisoner's dilemma demonstrates this). Let us recall that the game has two Nash equilibria in pure strategies. If the players stick to these equilibria they can get the payoff of 4 jointly. One Nash equilibrium (C,C) favours the girl, and the other one (B,B) favours the boy. An obvious way to introduce a sort of symmetry is to design a simple correlation device. The players should toss a coin. If it is Heads, then they choose (C,C), if it is Tails, then they choose (B,B). This is a similar correlation device to the one identified in the example I in this lecture.

5.3 In example II players enjoy the payoff 2.5 each (on average) if they choose strategies randomly with 50% probabilities. In E-5 (the exercise listed in overheads to this lecture) a correlation device was defined which allows them to enjoy the average payoff of $3 \frac{1}{3}$. Can you improve the result?

It can be improved theoretically. The average payoff of $3 \frac{1}{3}$ was achieved thanks to a random signalling device sampling three states: A, B, and C with probabilities $\frac{1}{3}$ each. Thanks to the device the players avoid the worst outcome (0,0), because they never choose U and R simultaneously. If the signal is B the First chooses D, the Second chooses L, and they enjoy the highest payoff (4,4). Even though (4,4) is not a Nash equilibrium, they do not try to unilaterally move away; the First does not know whether the signal was B or C, and the Second does not know whether it was A or B. They receive (5,1), (4,4) and (1,5) with probabilities $\frac{1}{3}$ each. On average it makes $(3 \frac{1}{3}, 3 \frac{1}{3})$. This average can be increased if the probability of B is increased. For instance, if it grows from $\frac{1}{3}$ to $\frac{1}{2}$, and the probabilities of A and C are $\frac{1}{4}$ each, then they receive (4,4) in 50% cases, and on average they receive (3.5, 3.5). This is what the definition of correlated equilibrium suggests. It also suggests that the average payoffs can be increased further if the probability of B goes up. If it reaches 1, then the distribution becomes 'degenerate', and players should understand that when they are told B or C, and A or B in fact it means B. And if they are sure that the signal was B then they may take advantage of moving away unilaterally. If both of them do the same, they will end up in (0,0) – something they wanted to avoid. But if the probability of B is very high (even if it is not 1, but it is much higher than $\frac{1}{3}$), they may try to take advantage of moving away unilaterally (when learning B or C and A or B, but being almost sure that it was B), and the attempt to avoid (0,0) will fail. This is an empirical issue. After all, when von Neumann and Morgenstern assumed that people maximise expected utility, the Allais experiments have not been carried out yet.

5.4 Examples II and III demonstrate that sometimes it is better not to know something. Why is it so?

A common wisdom suggests that it is always better to know (and perhaps disregard information we have) than not to know. It may be true if the other person who plays the game with us does not know that we know. Please note that in both examples referred to, players improved the average outcome because they knew that certain information was not disclosed to everybody.

In every country intelligence service tries to possess information that is confidential or secret. But when they succeed they pretend that they do not have it. Only when your rival does not know that you know, your success gives you a real advantage.

Let us assume that in the three-person game from example III, the Third player acquired information that $s = \text{Heads}$. Thus he (or she) knows, that the First player chose U, and the Second chose L. Matrix B will give the payoffs (2,2,2). However, if the Third player chose the matrix A, the payoffs would be (0,1,3). What is he (or she) likely to do? The decision may depend on whether this is a one-shot game or whether it is going to be played over and over again. In lecture 8 (GT-8) a 'backward induction' principle will be analysed; it will be explained that our decisions may depend on whether we expect a follow-up or not.

If the game is considered a one-shot event, then the Third player is likely to choose the matrix A. By doing so, he (or she) reveals knowing that the signal was as it was. But if the game is going to be played over and over again, the Third player is more likely to hide the information about the signal (just like intelligence officers do) and choose the matrix B. His (or her) payoff will be 2 instead of 3, but this may be more attractive than leaving the correlated equilibrium (which is based on the secrecy of the signal) and slipping into the Nash equilibrium of (1,1,1) for ever.

5.5 Can the QWERTY keyboard be considered 'a stable social convention'?

Yes. In many countries we use the QWERTY keyboard (the name comes from the first letter keys in the upper left of your keyboard). Not everybody applies the same standard. For instance in French speaking countries they use the QWERTZ keyboard. So when I asked a colleague in her office in Brussels to use her computer, when I tried to type my name, instead of Zylicz I saw Yzlicy on the computer screen. If you are used to a QWERTY keyboard, then you lose a lot of time when you type on another one.

Likewise, some of the French-speaking visitors who come to my office have to spend extra effort if they wish to use my computer. QWERTY keyboard was adopted some time ago in order to allow easy typewriting, and soon it became a standard. In the middle of the 20th century studies were carried out in order to optimise keys' positions. It turned out that QWERTY keyboard is not a good solution (at least not good for English language users). There were attempts to introduce these better keyboards, but they were abandoned. The reason behind this was that the cost to be paid when switching to a better standard would be excessive. Please imagine a situation that the standard is QWERTY, and you want to unilaterally abandon it (even if there are thousands of users of the new keyboard there are millions of users of the old standard). If everybody left the old standard, and moved to the new one, the situation would be different.

QWERTY keyboard is a social convention. It proved to be stable, because attempts to move to something different (better) failed. They did, because the cost of a unilateral – that is not

universal – change would be excessive. This is what we check when we look for Nash equilibria. QWERTY keyboard should be considered a sort of a Nash equilibrium.

5.6 What kind of a 'stable social convention' does the game on page 31 allude to?

The Nash equilibrium in mixed strategies was calculated as $(1/2, 1/2; 1/3, 2/3)$. It means that women are loyal in 50% cases, and men in 33% cases only. If a woman is found not to be loyal, then many people are outraged. They are less outraged if a man is found not to be loyal. The social convention which supports such opinions has been stable over the last centuries (if not millennia).

6 – Bayesian games

Thomas Bayes (1701-1761) was the first to introduce the concept of conditional probability and to suggest how acquiring information helps to improve estimates of likelihood. Pierre-Simon Laplace (1749-1827) obtained similar results independently, but Bayes was the first one and his name is referred to whenever we want to emphasise that using *a priori* beliefs can influence our reasoning.

You should remember the definition of conditional probability from your high school mathematics classes. It reads:

$$P(A \mid B) = P(A \cap B) / P(B)$$

The simplest version of the Bayes formula reads:

$$P(A \mid B) = P(B \mid A) \cdot P(A) / P(B)$$

The following Example I explains what the conditional probability is, and what the Bayes formula allows us to calculate. The example is about sampling four balls (one white and three green) from two boxes (left and right). The left box (L) contains two green balls (and no white balls). The right box (R) contains one green ball and one white ball. The green balls can be identified as G_1 , G_2 , and G_3 , and the white ball can be identified as W_1 .

Example I

When sampling takes place, either the left box is selected or the right one. We assume that they are selected with the same probabilities, $P(L) = P(R) = 1/2$. Since there are four balls, the obvious arithmetic implies that the probability of selecting a ball is


$$P(G_1) = P(G_2) = P(G_3) = P(W_1) = 1/4.$$

However, the probability of sampling a green ball is $P(G) = 3/4$. Similarly the probability of sampling a white ball is $P(W) = 1/4$.

Picture below illustrates the story, and the formulae below calculate conditional probabilities.

- $P(G \mid L) = P(G \cap L) / P(L) = (1/2) / (1/2) = 1,$
- $P(G \mid R) = P(G \cap R) / P(R) = (1/4) / (1/2) = 1/2,$

- $P(W | L) = P(W \cap L) / P(L) = 0 / (1/2) = 0$,
- $P(W | R) = P(W \cap R) / P(R) = 1/4 / (1/2) = 1/2$.



These numbers are obvious, but let them be explained anyway. The conditional probability that a green ball was sampled, if we know that it came from the left box is 1. This is a sure thing, because the left box does not contain white balls at all. The conditional probability that a green ball was sampled, if we know that it came from the right box is 1/2; the information that it was the right box changed the unconditional probability $P(G)=3/4$ to the conditional probability $P(G | R)=1/2$. The conditional probability that a white ball was sampled, if we know that it came from the left box is 0; it is impossible to observe it, because the left box does not contain white balls at all. Finally the conditional probability that a white ball was sampled, if we know that it came from the right box is 1/2; the information that it was the right box changed the unconditional probability $P(W)=1/4$ to the conditional probability $P(W | R)=1/2$. The corollary from these calculations is that *a priori* information (or even simply beliefs) may influence decisions. In this case these were decisions about probability estimates. When playing games the *a priori* information may have impact on how people choose.

The following example illustrates how *a priori* beliefs influence players' strategies. It is a slightly modified prisoner's dilemma.

Example II (modified prisoner's dilemma)

Let there be two games

"I"		Second	
		C	D
First	C	(-5,-5)	(-1,-10)
	D	(-10,-1)	(0,-2)

and

"II"		Second	
		C	D
First	C	(-5,-11)	(-1,-10)
	D	(-10,-7)	(0,-2)

The difference between "I" and "II" is in payoffs for the Second player in his "C" strategy. They read -11 or -7 instead of -5 or -1 (they are lowered by 6 in both cases). From the

psychological point of view we can think of these two variants as games played by a ruthless guy in I, and a sensitive guy in II. Please recall from the discussion of the prisoner's dilemma in lecture #4 (GT-4), that criminals promised each other not to confess. The Second player in I does not feel bad if he confesses, even though he breaks a promise given to the First one. In contrast, in II he feels bad if he confesses, as demonstrated by his payoffs (-11 instead of -5, and -7 instead of -1).

In other words, the difference between "I" and "II" is in the type of the Second player. The First player does not know what is the type of the Second player, but he has certain beliefs. Based on their acquaintance, the First player thinks of the Second player either as a ruthless guy, or a sensitive guy. Namely, the First player believes that the game is of the type "I" with probability μ , or of the type "II" with probability $1-\mu$. The Second player knows his type (knows whether the game is of the I or II type).

If the game is of the type "I" then the strictly dominant strategy of the Second player is C, but it is D, if the game is of the type "II". Thus the First player – who believes that the game is of the "I" type with probability μ – should choose C if $-10\mu + 0(1-\mu) < -5\mu - (1-\mu)$, or – alternatively – choose D if $-10\mu + 0(1-\mu) > -5\mu - (1-\mu)$; if the equality holds, the First player is indifferent between choosing C and D. Both inequalities are based on expected payoffs in I and II.

Thus the solution (the best strategy) for the First player is C if $\mu > 1/6$, and D if $\mu < 1/6$. The solution depends on *a priori* beliefs or perhaps on some other information that he has.

Example II demonstrates how a Bayesian game can be solved, but its general definition is more complicated.

(D.23)

A Bayesian non-zero sum two-person game is a game where payoffs $P_{ij}(\theta_P)$ and $D_{ij}(\theta_D)$ depend on realisation of random variables θ_P and θ_D . It is assumed that: (1) distributions of θ_P and θ_D are known to both players; and (2) realisation of θ_P is known to the First player only, and of θ_D – to the Second one only.

In order to understand D.23, one has to recall the difference between distributions and realisations of random variables. Let us assume that, for instance, the random variable considered is the prize in a lottery. It can be either -2 € or 2 €. It has a two-point distribution (consisting of two states: -2, and 2) characterised by probabilities of ending up in either of the states, say, 60% and 40%. The numbers -2 and 2 are realisations of this random variable (specific outcomes of the lottery).

In the example above it was assumed that θ_P had a degenerate distribution, that is P_{ij} values were known with certainty (they are identical in I and II: $P_{CC}=-5$, $P_{DC}=-10$, $P_{CD}=-1$, $P_{DD}=0$). Thus – knowing just the distribution – in fact the Second knew the realisation of the variable too. In contrast, at the same time, θ_D had a two-point (not degenerate) distribution such that $D_{CC}=-11$ and $D_{DC}=-7$ with the probability of μ , and $D_{CC}=-5$ and $D_{DC}=-1$ with the probability of $1-\mu$. ($D_{CD}=-10$, and $D_{DD}=-2$ for any realisation of the random variable θ_D). Knowing just the distribution, the First did not know the realisation of the variable. As the realisation of this variable (whether the Second player is ruthless or sensitive) determined which game (I or II) was played, the First had to choose his strategies conditionally.

D.23 is general, but numerical solutions can be easily found in special cases only. The special case that our example II was based on requires that one distribution is degenerate. In such a case, if one knows the distribution of a random variable, one knows its realisation as well. D.23 can be generalized so as to apply to n-person games. In plain language it says that players look at expected payoffs taking into account that actual payoffs are random variables with known distributions (but not necessarily realisations).

(D.24)

In a Bayesian game D.23, a (pure) strategy (decision rule) for the First player is $S_P(\theta_P)$ and for the Second player – $S_D(\theta_D)$ with payoffs $E_{\theta}P_{S_P(\theta_P)S_D(\theta_D)}(\theta_P)$ and $E_{\theta}D_{S_P(\theta_P)S_D(\theta_D)}(\theta_D)$, where E_{θ} is the expected value given the distribution of θ_P and θ_D , $S_P \in \{1, \dots, m\}$ and $S_D \in \{1, \dots, n\}$.

D.24 lets define a Nash equilibrium for Bayesian games that is for games with payoffs being random variables rather than concrete numbers (or functions) resulting from specific decisions of players. The general logic of Nash applies: players are in equilibrium if none has a motivation to unilaterally move away.

(D.25)

In a Bayesian game D.23, a (pure strategy) Bayesian Nash equilibrium is a pair $(S_P(\theta_P), S_D(\theta_D))$ of D.24 decision rules such that for any other strategies $S'_P(\theta_P), S'_D(\theta_D)$:

$$E_{\theta}P_{S_P(\theta_P)S_D(\theta_D)}(\theta_P) \geq E_{\theta}P_{S'_P(\theta_P)S_D(\theta_D)}(\theta_P)$$

and

$$E_{\theta}D_{S_P(\theta_P)S_D(\theta_D)}(\theta_D) \geq E_{\theta}D_{S_P(\theta_P)S'_D(\theta_D)}(\theta_D)$$

Please note that these inequalities state that strategies $S_P(\theta_P)$, and $S_D(\theta_D)$ give the players an opportunity to maximise their expected payoffs understood as random variables. Thus before taking decisions they need to check all the possible realisations of these variables.

In the example II above the Bayesian Nash equilibrium was a pair of strategies $(S_P(\theta_P), S_D(\theta_D))$ where $S_P(\theta_P) = S_P = (C \text{ if } \mu > 1/6, D \text{ if } \mu < 1/6)$, and $S_D(\theta_D) = (C \text{ if the realisation of } \theta_D \text{ is type "I", D if the realisation of } \theta_D \text{ is type "II"})$. Please note that S_D depends on θ_D . Both players know that the distribution of θ_P is degenerate (there is only one realisation of θ_P), and the distribution of θ_D has two values (attained with probabilities μ and $1-\mu$). The First player does not know the realisation of θ_D , but the Second one does. In other words, the Second is privileged by knowing which game (I or II) he plays.

It is important to interpret μ and $1/6$ in this example. The number $1/6$ is calculated from the payoff tables. In contrast μ is an *a priori* belief of the First player about the character of the Second player. Namely the first player may feel sure that the Second player is a ruthless guy. In this case $\mu=1$ ($\mu > 1/6$), and the First player should choose C. But the First player may feel sure that the second player is a sensitive guy. In this case $\mu=0$ ($\mu < 1/6$), and the First player should choose D. Of course, the First player can be uncertain about the character of his companion. If he believes that there is a 50% chance that the Second player is a ruthless guy ($\mu=1/2$), then $\mu > 1/6$, and his decision should be C. If he believes that there is only a 10% chance that the Second player is a ruthless guy ($\mu=0.1$), then $\mu < 1/6$, and his decision should be D.

Bayesian games are sometimes called games of imperfect information. This name emphasises the fact that players do not know everything about their rivals. In example II the First player

did not know whether the Second player takes decisions by looking at the payoff table I or II. If information is imperfect, then the choice of strategies may depend on our *a priori* beliefs.

Questions and answers to lecture 6

6.1 If $P(A | B) = P(A)$, A and B are called statistically independent. Can you give an intuitive meaning of the term?

A better definition of statistical independence reads $P(A \cap B) = P(A) P(B)$: the probability of $A \cap B$ can be calculated as the product of probabilities of A and B. The equality is not always true. For instance in the example I, if A means sampling a white ball, and B means sampling from the right box, then $P(A \cap B) = 1/4$, but $P(A) = 1/4$, and $P(B) = 1/2$. The reason that the left hand side (1/4) is different from the right hand side (1/8) is that A and B are not statistically independent. By combining the criterion $P(A \cap B) = P(A) P(B)$ and the definition of a conditional probability $P(A | B) = P(A \cap B) / P(B)$ we get $P(A | B) = P(A)$ as in the question, but this requires an assumption that $P(B) \neq 0$ (which does not hold in general).

Assuming that $P(B) \neq 0$ (otherwise the conditional probability cannot be defined), the equality $P(A | B) = P(A)$ says that conditional and unconditional probabilities are the same. Intuitively it means that knowing B does not change our estimate of the likelihood of A. Coming back to the example I, the information that sampling was from the right box does change our estimate of the likelihood of sampling a white ball. Without this information we estimate it at 1/4, but with the information it is 1/2. Similarly information that sampling was from the left box makes the likelihood of sampling a white ball 0, and the likelihood of sampling a green ball – 1. Statistical independence of any two events means intuitively, that knowing that the second happened does not change our estimate of the likelihood of the first one.

6.2 Statistical evidence shows that most car accidents (say, 80% of them), happen along a straight part of a highway (say, only 20% of them happen where the highway turns). Can one infer from this that driving along a straight part of a highway is more dangerous?

No. You have to check the distribution of straight parts and curvy parts of a highway. Let us assume that 90% of them belong to the first category, and 10% – to the second one. If 20% of the accidents (rather than 10%) happen in locations that belong to the second category, then this second category must be more dangerous. Please note that 80% is not the conditional probability of having a car accident if you drive along a straight part of a highway; this probability is much lower!

6.3 The conditional probability of acquiring a cancer if you smoke cigarettes is 3%. 20% of the population smoke cigarettes, and 80% of them do not smoke. The overall occurrence of a cancer is 1%. Assuming that a patient is diagnosed with cancer, what is the likelihood that he (or she) smoked?

According to the Bayes formula, $P(A | B) = P(B | A)P(A) / P(B)$. If B is having a cancer, and A is smoking, then the information that 3 % of smokers acquire a cancer, 20% people smoke, and 1% suffer of cancer allows us to calculate the right hand side: $0.03 \times 0.2 / 0.01 = 0.6$. Hence if a patient is diagnosed with cancer, the probability that he (or she) smoked is 60%.

6.4 Example II took advantage of the fact that in both variants (I and II) the Second player had a strictly dominant strategy (C in I, and D in II). Thus the First player knew what his rival was likely to choose. How would the calculations of μ change if there were no dominant strategies?

They would have been much more complicated. The First player would have to guess what strategy (C or D) the Second player is likely to choose in alternative circumstances. The choice of the Second player would depend on what he thinks that the First player thinks about himself (i.e. about the Second player). If the First player considers the Second player to be a sensitive guy, and the Second player expects that this is the case (irrespective of whether this is true), he may behave in a certain way. If the First player considers the Second player to be a sensitive guy, but the Second player is not aware of this, he may behave in a different way. There are too many arbitrary assumptions to be made to try to solve such a difficult problem.

6.5 If μ calculated for a Bayesian game like in the example II is a number close to zero, then what strategy (C or D?) should the First player choose?

If the number is close to zero, then unless the First player is almost sure that the rival is a sensitive guy, the choice should be C (that is "confess"; because it is almost certain that the rival confesses despite what he promised before the arrest). If the First player is almost sure that the rival is a sensitive guy, the choice should be D (that is "deny"; because it is almost certain that the rival will deny as promised before the arrest; their joint denial will let them enjoy a fairly light sentence).

6.6 Can the knowledge of the specific realisation of a random variable change a decision?

Yes, it can. Let us assume that a certain gene implies that people who have it are less productive; on average – because of absenteeism – their output is only 80% of what can be expected from people who are free of this gene. But people who have it can be as productive as those who do not have it. Having this gene is a random variable. The expected value implied by its distribution is 80% of what can be expected otherwise, but some people are 100% productive; the realisation of that random variable can result in no productivity deficit at all.

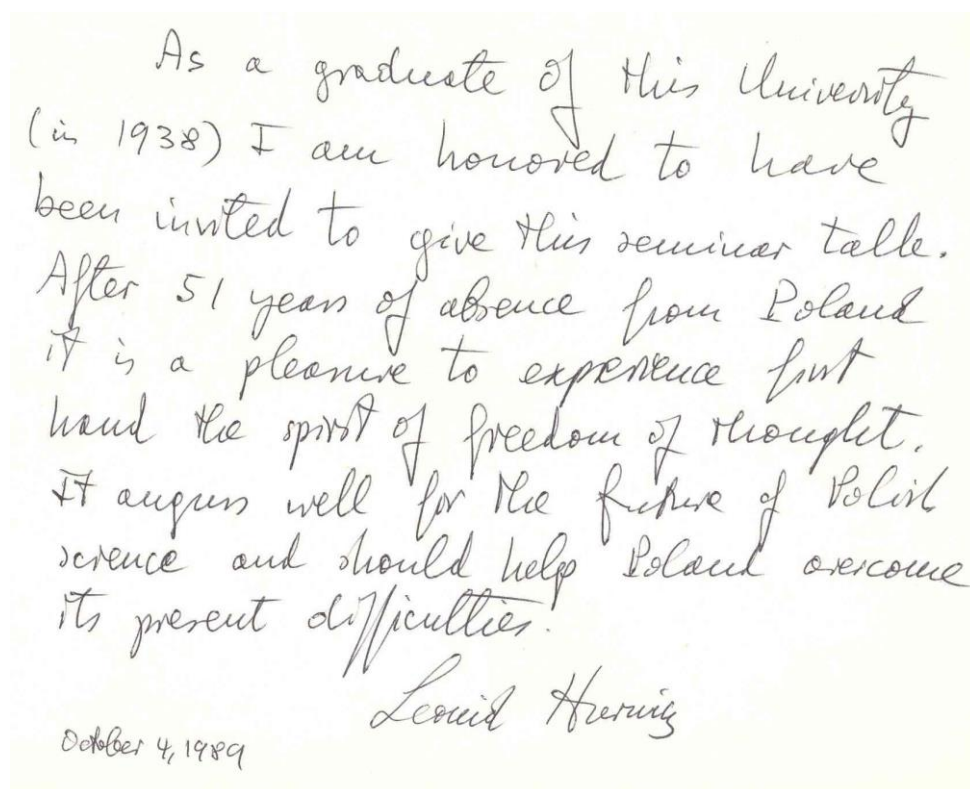
A potential employer has a method of checking this gene. Based on a test, he (or she) is likely not to employ a person that was found to have it (because the expected productivity is too low). However, in the case of this particular person the productivity is the same as without the gene. If the employer knew the realisation of this random variable he (or she) would have employed this person.

This is like deciding to take part in the lottery with two outcomes: -2 €, and 2 €. Based on the distribution of payoffs some people may consider it unattractive. But the realisation can be 2 € (attractive). If they knew the realisation rather than distribution, the decision to take part would be different (positive).

7 – Mechanism design

Leonid Hurwicz (1917-2008) was awarded Nobel Prize in economics in 2007 for his role in *Mechanism Design* (among other things). The topic of this lecture – mechanism design – was a difficult area at least for me.

Leonid Hurwicz was a graduate of the University of Warsaw. However at that time (in 1938) he attended a so-called economics seminar in the Department of Law, since there was no separate economics faculty. In 1989 he talked about mechanism design in our Department (see picture below). I understood every mathematical detail of his argument, but I could not comprehend the entire idea. I started to grasp it once I learnt that the famous Solomon's judgement (see Example I below) can be interpreted as an early application of the mechanism design.



As a graduate of this University
(in 1938) I am honored to have
been invited to give this seminar talk.
After 51 years of absence from Poland
it is a pleasure to experience first
hand the spirit of freedom of thought.
It augurs well for the future of Polish
science and should help Poland overcome
its present difficulties.

Leonid Hurwicz

October 4, 1989

'Mechanism design' means constructing a game which reveals something by overcoming asymmetric information. The definition of asymmetric information is straightforward: the buyer has less information about the commodity than the seller or *vice versa*; acquiring information is possible, but costly.

There are examples when the buyer knows less than the seller, like in the second-hand car market. But there are also examples when the seller knows less than the buyer. Insurance policy markets have to deal with such a problem. If you want to buy insurance against medical care expenditures, the seller may ask you whether any of your ancestors suffered from a certain disease. You may have such an information, but you do not have to disclose it. As a result, the company cannot estimate a fair price of your policy. It may be too high if the company is afraid that you may require an expensive treatment but you will not, or *vice versa*. It may be too low, if the company does not expect you to require an expensive treatment when in fact you will. Consequently insurance policy markets cannot function well.

In the present lecture we will refer to a labour market where this asymmetry takes place too. The following *principal-agent* model explains the idea of how asymmetric information can be dealt with. The model is about designing a contract between an employer (called principal) and an employee (called agent) when the employer does not know exactly what the employee does. The 'product' the employer is interested in depends on the 'effort' the employee puts into the job. The problem is caused by the fact that the agent's effort cannot be observed perfectly by the principal, since only the agent knows it precisely. Likewise only the agent knows his (or her) aspiration level, that is the minimum net outcome of his (or her) effort. The model is formalised as follows.

- x – employee's effort
- $y=f(x)$ – product (we assume that its price is equal to 1)
- $s(y)$ or $s(x)$ – employee's salary
- $c(x)$ – cost born by the employee
- u^0 – employee's aspiration level: $s(f(x))-c(x) \geq u^0$ (*participation constraint*)

As it is, the model can be seen as an abstract one. Let me explain what it means by looking at my contract. The Rector of the University of Warsaw is my 'principal'. Being professor of economics, I am an 'agent'. Teaching is what I am supposed to do. I can put a lot of effort into what I am doing, if I start lectures on time, if I answer students' questions competently, if I prepare course materials, if I spend a lot of time to make exams, and so on. Or – on the contrary – I put little effort when I come late to the class, when I ignore what students ask me about, when my course materials are prepared poorly, if I am not very serious when I give the exams, and so on. Putting adequate effort into what I am supposed to do requires some cost. For instance, if I woke up late, then I need to take a taxi instead of a bus. If my students' questions are to be answered competently, I must read scientific literature. If my course materials are to be useful, I need to spend a lot of time preparing them. Likewise, if the exam results are to be reliable, I need to spend a lot of time preparing them.

The cost of the effort can be understood as an out-of-pocket expenditure, for example, if I take a taxi instead of a bus. But economists understand it also in terms of opportunities lost. For instance, if I spend my time on preparing notes or exams, I cannot earn money otherwise. If I have an opportunity to earn money by, say, cleaning windows, I lose this opportunity if I prepare my lecture or if I prepare an exam.

The Rector wants me to put a lot of effort into what I am supposed to do, but he cannot control it easily. He can ask the students, if I come on time. In this case acquiring information is fairly simple. Students can answer such a question easily. But if the Rector asked you about the quality of my lecture, some students would probably say that that it was good, but others would say, that it was bad. Are students well prepared to answer such questions competently? Despite of what some people suggested today, I am afraid that these are difficult questions. Sometimes students like professors who – irrespective of whether they are clever or not – are polite, smiling, tell jokes, do not require a lot of work, and so on. Such professors can fare fairly well in students' valuations.

But is this what the Rector cares for? The Rector would like to have students graduating with state-of-the-art knowledge and successful in the labour market. If the University is considered as a good starting point for a professional career, then high school graduates are likely to

enrol. This is what the Rector cares for. He would like me to put adequate effort into what I am doing, because he is interested in the 'product' of my job, that is in the success of the University.

Asking students is easy, but – as noted above – information obtained in this way about how much effort I put into the job is not always reliable. In order to know more, the Rector can send his spies to check the quality of my work. But spies are not always reliable (I can bribe them), so to learn better, the Rector can send second order spies (spies who spy on the first order spies) in order to verify information collected from the first order spies. To be at the safe side, the Rector can send third order spies (who spy on the spies who spy on the first order spies) and so on. Spies do not work for free usually – they need some salary. Therefore acquiring information – even though possible – can be rather costly. It may turn out that by financing spies the Rector motivates me to improve my performance somewhat, but it is reflected by improved enrolment to a disappointingly low degree only.

I would like to maximise my salary (net of cost), and the Rector would like to minimise it. If he pays me too little – less than required by my *participation constraint* – I will quit the job. My aspiration level is not known to the Rector, and if asked, I would declare it very high, hoping the Rector to pay me more. A more likely situation is that my *participation constraint* is satisfied, and I will stay at the University. Nevertheless, my salary may motivate me to put too much effort into what I am doing or too little. In the first case, I am paid so well that I spend a lot of time to prepare course materials, but my effort does not translate into better effects to the extent expected by the Rector. In the second case my salary does not motivate me to spend a lot of time to prepare course materials, and – as a result – our graduates are not satisfied with what they learn in the University.

Let us come back to the main topic of the lecture. The principal would like to minimise the salary. On the other hand, the agent would like to maximise it. The salary cannot be too low, because the principal risks to get the profit lower than possible. At the same time, the salary requested cannot be too high, because the agent risks losing the opportunity to be employed at all, and thus losing the possibility to satisfy his (or her) aspiration level. The following theorem states how these conflicting tendencies can be reconciled by designing an *incentive compatible* contract (i.e. a contract which satisfies the so-called *incentive compatibility constraint*).

(T.13)

The incentive compatibility constraint is: $s(f(x^*)) - c(x^*) \geq s(f(x)) - c(x)$ for all x , where

- x^* maximizes $f(x) - s(f(x))$, i.e. $f(x) - c(x) - u^0$ that is (by conventional assumptions):
- $MP(x^*) = MC(x^*)$

Proof:

The inequality $s(f(x^*)) - c(x^*) \geq s(f(x)) - c(x)$ means that x^* solves the problem

$$\text{Max}_{\text{for all } x} \{s(f(x)) - c(x)\}$$

That is the agent's problem. Thus *incentive compatibility constraint* means that x^* maximizes the net benefit for the agent. The theorem (first bullet) states that the same x^* maximizes

$$f(x) - s(f(x)),$$

i.e. the net benefit for the principal. By the rationality of the principal, $s(f(x)) = c(x) + u^0$ (the principal should offer the employee exactly what is expected, nothing more). The expression

$f(x)-c(x)-u^0$ is maximized (under the concavity assumption of the function $f(\cdot)-c(\cdot)$) for x^* such that $f'(x^*)=c'(x^*)$ (the second bullet, because $MP=f'$, and $MC=c'$).

Corollary

Incentive compatibility constraint is satisfied if the employee is a so-called *residual claimant* (i.e. has the right to get the entire marginal product of his (or her) effort when it is close to x^*).

In our example of the Rector-Professor problem, the *residual claimancy* condition means that the professor's salary should be increased by the incremental success of the professor's extra effort. In other words, if my extra effort translates into the University success measured by additional 1,000 €, then these additional 1,000 € should be added to my salary. Of course, this is not practical, but it shows the idea of *incentive compatibility*. Some traces of this idea can be found in prize systems. If a professor is recognised as somebody standing behind the success of the University, he (or she) can be awarded a special prize (in addition to the regular salary). They can be found also in employment regulations: if a professor is recognised as somebody contributing negatively to the success of the University, he (or she) can be fired.

The following employment arrangements provide examples of incentive-compatible contracts. The numbers R , K , and B are calculated in such a way, that the agent is a *residual claimant*, that is $MP(x^*) = MC(x^*)$.

Examples (of incentive-compatible contracts)

- Rental payment, R : $s(f(x)) = f(x)-R$, where R is derived from the participation constraint: $f(x^*)-c(x^*)-R = u^0$
- Hourly (daily) salary rate w plus flat rate K so that: $s(x) = wx+K$, where $w=MP(x^*)$, and K is derived from the participation constraint: $wx+K-c(x) = u^0$
- Threshold condition (*take-it-or-leave-it*), payment B , if $x \geq x^*$ (alternatively: if $y \geq f(x^*)$): the amount B is calculated from the participation constraint: $B-c(x^*) = u^0$ (assuming that $B \leq MP(x^*)$); otherwise the agent is paid nothing


In all examples the contract is based on the agent's participation constraint, and – indirectly – on a relevant aspiration level which is unknown to the principal. Auctions are a mechanism of revealing this unknown aspiration level. Please refer to the first of examples listed above. The story can be about a landlord who has a piece of land, but does not want to cultivate it. He (or she) invites farmers to make offers of how much they are willing to pay him (or her) in order to cultivate this land and enjoy the benefits of its crops. The landlord orders a so-called sealed-bid auction where everybody interested writes the amount of the rent to be paid for this piece of land. When all the offers (bids) are delivered in sealed envelopes, the landlord opens them, and selects the bid which offered the highest rent. The land goes to the farmer who offered this largest bid.

How does this relate to the aspiration level? Every farmer made the following calculation. If I put the amount of effort x into the cultivation, the crop will be $f(x)$. If I offer the rent R , then I will be left with $f(x)-R$. If what I am left with meets my aspiration level u^0 , I will take the arrangement. The farmer had to solve the inequality: $f(x)-R \geq u^0$, or $R \leq f(x)-u^0$. Of course the lower the rent the better for the farmer. From that point of view the best offer would be to declare the rent $R=0$, or even a negative one (that is to expect the landlord to pay for renting the land). But at the same time the farmer knows that his (or her) bid will lose if somebody

else offers a higher rent. Thus he (or she) offers the highest rent the inequality allows, i.e. $R=f(x)-u^0$.

So far the rent depends on the amount of effort x put into the cultivation. The rent that the farmer finally offers is based on x^* , that is the amount of effort which maximises the expression $s(f(x))-c(x)$. In other words, by organising the auction the landlord acquires information on what is the value of his (or her) land. This value depends (indirectly) on the unknown to him (or her) aspiration level of the farmer who can make the best use of the land. An incentive compatible contract makes the agent the *residual claimant*, i.e. the agent knows that if he (or she) puts more effort, the entire product of this can be claimed by him (or her).

I trust that the examples explain what an incentive compatible contract looks like. But it seems to me that an even more instructive example is to analyse a contract which fails to be incentive compatible. My colleagues economic historians tell me that the most popular type of contract in the world over the last 10,000 years is renting land for a share in crops. Instead of the arrangement explained above, the landlord announces that he (or she) collects offers indicating what share of crops tenants are willing to pay as a remuneration for leasing the land. Let us look at a picture which demonstrates what is the 'optimum' effort that the farmer contemplates.



The farmer would like to get as much as possible from the land. If he (or she) is the residual claimant he (or she) looks at MB and MC in order to determine x^* as the optimum effort. If the effort is higher ($x > x^*$) then what the farmer enjoys (MB) is lower than what the farmer has to pay (MC). Hence this additional effort does not make sense. If the effort is lower ($x < x^*$) then what the farmer loses (MB) is higher than the cost avoided (MC). Hence lowering the effort does not make sense either. Indeed x^* is the optimum effort.

Now let us analyse the contract for, say, 40% of the crop (40% of whatever is produced goes to the landlord, and 60% remains for the farmer). As before, all the costs are to be borne by the farmer. If this is the case then the farmer looks at the 0.6MB curve (the red one rather than the black MB) in order to determine x^* (printed in red) as the optimum effort. If the effort is higher ($x > x^*$) then what the farmer enjoys (0.6MB) is lower than what the farmer has to pay (MC). Hence this additional effort does not make sense. If the effort is lower ($x < x^*$) then what the farmer loses (0.6MB) is higher than the cost avoided (MC). Hence lowering the effort does not make sense either. Indeed x^* (printed in red) is the optimum effort.

One clarification on the cost of effort. It is understood not only as the physical effort the farmer has to exert, but also as expenditures the farmer has to bear. These expenditures are

understood in economics as an 'opportunity cost', for example when the farmer uses his (or her) child labour. They can be also understood as 'out-of-pocket' expenditures, when the farmer hires additional workers, buys fertilisers, pesticides or has to purchase any other input. Everything is his (or her) responsibility.

Picture on page 45 demonstrates that the rental for, say, 40% of the crops is not incentive compatible. It motivates farmer to less effort than justified by interests of the society (consisting of landlords and farmers). If it is not incentive compatible (it motivates for exerting less effort than optimum), why was it applied for hundreds of generations worldwide? One answer that some students would like to hear is that people are stupid; they do not understand economic concepts, and they behave irrationally. I am afraid that this is not a satisfactory answer. The rental contract where the rent is not fixed, but it depends on what the farmer produced has important advantages. Namely it is considered safe and fair unlike the contract including a fixed rent.

Under the fixed-rent contract all the risk is borne by the farmer, while the landlord does not risk. I expected the students to note this, but apparently it turned out to be not that obvious. If the weather is fine and the crops are abundant, the farmer is happy because everything above the rent is his (or her). If the weather is bad, and the crops are poor, the rent may exhaust whatever was produced, and nothing is left for the farmer. Apparently both landlords and farmers understood that a fixed-rent contract is unfair, and they agreed to the share-in-the-crops principle. Everybody understood that it is more difficult to enforce (because the landlord has to measure the crop somehow in order to calculate the 40%). Not everybody understood the concept of incentive compatibility, so not everybody saw the contract's inefficiency. Nevertheless there has been an overwhelming agreement as to its fairness.

The *principal-agent* approach (explained here by referring to the model with a single agent) can be extended to situations with many agents. The *incentive compatibility* condition should integrate interests of the principal and agents, and reveal information about agents' preferences (thus reducing the information asymmetry).

Now we are ready to look at the example of the Solomon's judgement. King Solomon has been known of unprecedented wisdom. He knew how to solve even the most difficult puzzles. His most famous judgement (see *1 Kings* 3:16-28) was when two women came to him with a newly born baby and each of them claimed that it was her child (there were no DNA tests available at that time). They knew who was the real mother, and who lied, but nobody else knew this. Also the king Solomon did not have this information. So he proposed to cut the baby into halves. He suspected that the real mother would object desperately, while the fake one would be outraged less. Indeed this is how the women reacted, so he acquired information he lacked. In game theory language the trick can be characterised as designing a game not to be won or lost, but just to observe what is its outcome. The outcome is supposed to be a Nash equilibrium. If one knows the Nash equilibrium then one can guess what hypothetical payoff functions could have been.

Example I: Judgement of Solomon

Two women quarrelled about who was the mother of a baby. Each of them knew the truth, but nobody else did. Solomon was to judge. He designed a game. Each woman was to say 'Yes' or 'No' to his proposal to cut the baby into halves. These answers can be summarised in the payoff matrix of a game.

The payoffs of -100 to the first if the baby is killed should be understood as something very large in absolute value and negative in sign. Solomon assumed that the real mother would say N, and the fake one would say Y. Based on the outcome (N,Y) he identified the First as the real mother thus solving his objective which was to make a fair judgement. The trick can be considered a successful implementation of 'mechanism design'. Of course, if the game were to be played repeatedly, its 'payoffs' would have to change.

		Second	
		Y	N
First	Y	(-100,1)	(-100,0)
	N	(0,1)	(0,0)

Sometimes students ask me which woman won and which lost. This is not the issue. As observed earlier, the game was not to win or lose; it was just to reveal information. The next example demonstrates how the 'mechanism design' works in less drastic problems of asymmetric information.

Example II: auctions

The principal has an object for sale. It can be purchased by either of two agents which value the object v_1 and v_2 , respectively. Each valuation is known only to the agent who keeps it secret. Nevertheless the principal and both agents know the distributions of both values treated as random variables: θ_1 for the first valuation and θ_2 – for the second. Agents are supposed to indicate their valuations as sealed bids s_1 and s_2 (which do not have to coincide with v_1 and v_2). The object goes to the agent who offered the higher bid. If both bidders offer the same bid then the winner is selected randomly with probabilities 1/2 and 1/2. Auctions vary with respect to what is the amount paid by the winner.

1. First-Price Auction

The winner pays the amount from his/her own bid

2. Second-Price Auction

The winner pays the amount from the loser's bid

It looks as if the Second-Price Auction was less attractive for the principal. This is not the case though, as the following theorems demonstrate. On average the principal gets exactly the same amount, that is $v_1=v_2$.

(T.14)

Let us assume that θ_1 and θ_2 have the same distributions. In both variants of auctions the truthful revelation of preferences makes a Nash equilibrium in the 'bidding game' (the payoff of the loser is 0, and the payoff of the winner is v_i-s_i in the first-price auction or v_i-s_{-i} in the second-price auction).

Proof:

Let us consider the first-price auction first. The winner pays s_i , and gets the net benefit v_i-s_i . If $s_i > v_i$ then the net benefit would be negative and the winner would have a motivation to be the loser rather than the winner (in other words, he does not have a motivation to overstate the valuation). If $s_i < v_i$ then the winner risks losing the auction and hence $s_i=v_i$ makes the

preferred bid. Now let us consider the second-price auction. The motivation for the winner not to overstate the bid is like before. A likely motivation to understate in order to pay less disappears, because the winner pays the bid of the loser (rather than his/her own). Hence – like in the first-price auction – $s_i = v_i$ makes the preferred bid.

(T.15) (Revenue Equivalence Theorem)

In both types of auctions the expected revenue of the principal (who organizes an auction) is the same.

Proof:

The seller receives s_i of the winner (in the first-price auction) or s_i of the loser (in the second-price auction). By T.14, bidders indicate their true valuations (i.e. $s_i = v_i$, for $i=1,2$). As v_i are sampled from the same distributions, then their expected values, $Ev_1 = Ev_2$ are the same.

The last theorem demonstrates why sellers should not be afraid of organising second-price auctions. On average they provide them with the same revenues as first-price auctions. Economists prefer the second-price auctions, so-called Vickrey's auctions (William Vickrey, 1914-1996, got the Nobel Prize in 1996 for his studies on asymmetric information). Even though they provide the seller with the same revenue theoretically, they provide stronger incentives not to underestimate bids, because the winner does not have to pay his (or her) bid, but rather the loser's bid. Thus the motivation to underestimate one's own bid ("If I win the auction I will have to pay what I declared") disappears. Hence one can argue that second-price auctions provide sellers with even higher revenues.

Auctions seem to be an excellent mechanism to reveal information. Nevertheless T.16 demonstrates that they are not perfect. I will quote this theorem without a proof for two reasons. First, the proof is very difficult. Second this is a game theory course, and we do not need to analyse auctions in such a detailed way.

(T.16) (Inefficiency Theorem)

Let us assume that θ_1 and θ_2 have strictly positive distributions on intervals that overlap at least partially. Neither of the auctions can – in general – provide an outcome that satisfies three conditions:

- individual rationality,
- incentive compatibility, and
- budget balance (any payoffs for the principal need to be financed from agents' payments).

The first two bullets are self-explanatory. The third one is more complicated. Under some conditions, in order to provide incentives for agents to truthfully reveal preferences, the principal needs to set aside certain funds to finance them. This subtracts from what the agents pay, and violates the budget balance understood as the requirement that the principal's revenue should be equal to the value of what he (or she) sells.

Probably an example from work-of-art auctions explains the idea. Each of the prospective buyers of an attractive painting are afraid of buying a fake copy, so they would like to pay for an expert assessment of the painting's authorship. As a result, the price they offer in the auction is decreased by this amount. Knowing this, the seller may prefer to pay for such an expert assessment and to share it with auction participants for free. But this violates the

'budget balance', because the cost of this assessment is paid from what the participants of the auction pay in their bids.

Auctions are a very popular area of economic research. The 2020 Nobel prize went to Paul Milgrom and Robert Wilson for their contribution to organising auctions. Their findings proved instrumental in establishing auctions for many difficult transactions like selling concessions for mobile phone operators. Lawrence Ausubel patented a complicated design for so-called multi-unit auctions where buyers have to offer not only prices, but also quantities to be bought. Sellers of raw materials (like timber or minerals) rely on such complicated auctions.

Questions and answers to lecture 7

7.1 Why is there asymmetric information in the second-hand car market?

Because the sellers know more than the buyers. The seller of the car (who drove it for several years) knows it better than the buyer (who takes a test drive for several minutes). If a prospective buyer asks the seller "Is this car a good one?" the seller is likely to answer "Yes, of course!" irrespective of whether this is true or not. Everybody knows this, so buyers do not even ask questions of that sort. Instead they try to acquire as much technical information about the car as they can, but they will never know what the seller knows (but perhaps does not want to share).

7.2 Why is the asymmetric information in the second-hand car market harmful for economic efficiency?

Because good cars are likely to disappear from the market, as the following example explains. Let us assume that there are two kinds of cars: good ones and bad ones. The owner of a good car is willing to sell it at the price of 8,000 € – this is what the car is worth to him (or her). A prospective buyer is willing to pay at most 10,000 € – this is what the car is worth to him (or her). They can agree to the price of 9,000 €. Irrespective of the price agreed, as a result of transaction, economic surplus increases by 2,000 €: the seller receives 9,000 € for something he (or she) valued at 8,000 €, so he (or she) is left with 1,000 € more, and the buyer paid 9,000 € for something that is worth 10,000 €, so he (or she) enjoys the surplus of 1,000 €.

At the same time there are bad cars. The owner of a bad car is willing to sell it at the price of 3,000 € – this is what the car is worth to him (or her). A prospective buyer is willing to pay at most 5,000 € – this is what the car is worth to him (or her). They can agree to the price of 4,000 €. Irrespective of the price agreed, as a result of transaction, economic surplus increases by 2,000 €: the seller receives 4,000 € for something he (or she) valued at 3,000 €, so he (or she) is left with 1,000 € more, and the buyer paid 4,000 € for something that is worth 5,000 €, so he (or she) enjoys the surplus of 1,000 €.

In the second-hand market both good and bad cars can be found. Buying a car in such circumstances is like taking part in a lottery. With 50% probability you will get a good car, and with 50% probability you will get a bad one (if good cars account for 50% of the total fleet). By looking at the expected value of what you get, you are willing to pay at most 7,500 € for a car. Owners of bad cars (who were willing to sell them for 3,000 €) will be happy to

accept such a generous price. But owners of good cars will not agree to such a price, and they are likely to withdraw. This is what economists call "adverse selection" resulting from asymmetric information. Hence only bad cars are likely to be traded.

The market failure is caused by the fact that good cars should change owners as well. It is a social loss if a car stays with a person who values it at 8,000 € instead of being transferred to a person who values it at 10,000 €. There is a potential improvement of 2,000 €, but the market cannot realise this. The efficiency of the economy is impaired.

7.3 How can the economic efficiency of the second-hand market be improved?

There are several ways to bring the good cars back to the market. Some of them have little to do with economics. For instance, you can pay an honest price of a good car if you buy it from a friend or a cousin you trust. They do not lie when they assure you that the car is good. Apart from that there are at least two economic mechanisms aimed at reducing the asymmetry of information. One of them uses guarantees. They can be either obligatory or voluntary. If the government requires that the seller gives a guarantee against defects of the car, buyers do not risk buying a car likely to break down (because any repairs are to be paid by the seller), and mutually beneficial transactions can take place. If guarantees are voluntary, buyers of good cars may request that their sellers give them a guarantee against any defects. If the owner refuses to give such a guarantee, information asymmetry is reduced, because people may suspect that the car is a bad one.

In addition to guarantees, an inspection of a professional mechanic can reduce the asymmetry. Some time ago when I wanted to buy a beautiful Chevrolet station wagon for a very attractive price of \$400, an American colleague advised me to pay \$20 to a mechanic for an inspection. This proved to be the best investment in my life. The mechanic revealed that the engine required at least \$1,000 in order to fix defects that were likely to show up after driving 50-100 kilometres. My expenditure of \$20 saved me \$1,000. Thanks to this moderate expenditure, I acquired information that turned out to be very valuable.

7.4 How do insurance companies try to reduce the asymmetric information before selling a health insurance policy?

Insurance companies face a different kind of asymmetric information. They know less than a potential buyer. The risk they have to cope with is that medical treatment required by the person insured is more expensive than what they expected. The first question they ask is about hereditary diseases and smoking habits. People can cheat, so the company typically requires professional blood and urine tests, and a visit to a doctor who works for the company. In addition, they exclude injuries caused by certain risky activities like skiing, parachuting, etc. If an insured person tries to cheat and wants the company to pay for a treatment that results from a forbidden activity, the company can check what was the cause of the injury and refuse to pay when they find that it was exempt from the insurance. They cannot reduce the asymmetry completely, but they try to keep it below what they can cope with.

7.5 How do employers try to reduce the asymmetric information resulting from hiring an employee?

In many cases hiring a new employee involves a complex procedure. The person should pass a number of tests in order to demonstrate his (or her) aptitude to play the roles the employer

expects. This is a relatively easy part of the procedure. The more difficult part is to agree on an incentive compatible contract. Its main objective is to provide the employee with incentives to put adequate effort into what he (or she) is doing. According to the principal-agent model, the contract should make the employee the 'residual claimant'. This is rather impractical, and a typical contract is based on a fixed salary. Employers prefer to offer short-term contracts to be extended only when the employee is found to put adequate effort into the job. After several months the employer has a much better knowledge of the employee's aptitude than at the time of recruitment.

7.6 Two conditions are stated in the principal-agent model: *incentive compatibility constraint*, and *participation constraint*. What are their roles?

The principal does not know the effort put by the agent into the job. The *incentive compatibility constraint* is to provide the agent with motivation to put the 'right' amount of effort into the job: to do what maximises the principal's net gain (product minus salary). The participation constraint is to make sure that the agent – if he (or she) puts the 'right' effort – is left with net gain (salary minus cost) sufficient to reach the aspiration level.

7.7 How can the Rector acquire information about a professor's effort put into teaching?

Surveying students is the most obvious way to check if classes are meeting their expectations. However they do not always reflect the professor's effort adequately. Students may be angry that the professor tries to enforce difficult requirements and consequently they evaluate his (or her) lectures low. Or *vice versa*, they may be amused by funny stories he (or she) tells and evaluate the class high. Such evaluations do not reflect the amount of effort well, the Rector is aware of this, and professors' performance is also checked by other methods. Whether the classes last 90 minutes can be checked fairly easily. Likewise course materials and exams are retrieved, so they can be evaluated by experts, if necessary. Classes are visited sometimes by other specialists. Nevertheless professors' salaries are usually fixed and they do not depend on the amount of effort directly. Professors' performance is evaluated by special committees periodically, and if found unsatisfactory it can result in their firing or other consequences.

7.8 Can Solomon's judgement become a routine method of solving mysteries?

No. The expected outcome can reveal missing information if participants of the 'game' are not aware of its objective. If they are, then they can act strategically, in order to create an impression that the truth is where it is not. In the case of the most famous Solomon's judgement, if there were more such puzzles (women quarrel over whose baby this is), then a woman could beg the king not to kill the baby not because it was her child, but because she suspected that based on her reaction the king would acquire information he looked for.

7.9 Can an unsuccessful Vickrey auction result in adopting the first-price auction next time?

It should not. An average Dutch is taller than an average Greek. But there are tall Greeks, as well as there are short Dutch. The height of an individual from a population has some statistical distribution. The height of a specific person is a realisation of the variable with this distribution. Meeting a tall Greek and a short Dutch does not imply that an average Greek is taller than an average Dutch.

Auctions are statistical events. It may be the case that the winner bade high and the loser bade disappointingly low. The owner of the object for sale could be upset by the fact that he (or she) gets the second price instead of the first one. Nevertheless, rather than taking a decision based on one realisation, the principal should make sure that next auctions reduce the probability of collusion of agents, or any other circumstances resulting in bids far from what they are expected (far from average).

7.10 Can you envisage a negative value of B in the contract defined on page 44 (third bullet)?

The contract can be for a person employed to watch over clothes left in a restaurant's cloak room. An average client who leaves the coat pays a tip. If the restaurant is visited by poor people, the sum of tips will be low probably. On the contrary, if the restaurant is visited by wealthy people, the sum of tips can be a very attractive source of income. In the first case, B is likely to be positive – the owner of the restaurant should offer an adequate remuneration for somebody who is obliged to watch over the clothes. In the second case, B is likely to be negative – the owner of the restaurant should expect to be paid in order to allow somebody to collect tips left by wealthy clients.

8 – Extensive form games

The most important feature of such games is that players do not have to move simultaneously, and they remember all the previous moves. Before we define an extensive form game, we will play a game that has been familiar to everybody. This is 'Tic Tac Toe'. Two players take moves in turn. They put either a cross (x) or a circle (o). The winner is who managed to place three signs in a row. Let us start. The guy who uses the cross puts it in a blank matrix:

x		

The other guy puts the circle somewhere.

x	o	

It is the first guy's turn now.

x	o	
x		

The second guy has to place the circle in the lower left corner in order to prevent the first guy from winning.

x	o	
x		
o		

The first guy places the cross in the upper right corner (this is a mistake).

x	o	x
x		
o		

Now the second guy places the circle in the lower right corner.

x	o	x
x		
o		o

The first guy has to prevent the second from winning.

x	o	x
x		
o	x	o

The second guy reacts somehow, but they probably see that the game cannot be won by anybody.

x	o	x
x	o	
o	x	o

The first guy places x in the last empty cell just to do something.

x	o	x
x	o	x
o	x	o

The game is familiar. The first has a winning strategy, but the second can win if the first makes a mistake (like in the scenario above). In general, when the game is over, a tie is announced (0,0), or the first has to pay to the second (-1,1), or the second has to pay to the first (1,-1). We will see now, how this trivial game can be defined using a very sophisticated mathematical terminology. But before we do it, we can draw it in the form of a 'tree' (which grows upside down).


It represents the tables from previous pages. In the beginning the table was empty and the cross-guy could select any of the 9 cells to put the "x" mark. Next it was the circle-guy's turn. He (or she) could choose any of the remaining 8 cells to put the "o" mark. Then it was the cross-guy's turn again. He (or she) could choose any of the remaining 7 cells to put the "x" mark. And so on. The number of branches of the tree is $9 \times 8 \times 7 \times 6 \times 5 \times 4 \times 3 \times 2 \times 1 = 9! = 362,880$ – a very large number. In the picture below I drew just those branches that were going to be extended. Otherwise picture (with all the 362,880 branches) would have been impossible to present. The place where new branches originate is called a node. The "x" and "o" letters identify the player who takes a decision in a given node. In the very last node, there is no real decision to be made by anybody. The numbers in brackets state the payoffs which go to the players. In the case of our Tic-Tac-Toe game the result was a tie (0,0), that is neither the first player nor the second one were paid anything.

Now please appreciate how clever the definition D.26 is. It defines even trivial games – such as Tic-Tac-Toe – as algebraic structures. We use the Greek letter Γ to identify a game. The subscript E refers to the so-called extensive form (the games we analysed in previous lectures were described in a "normal" or "condensed" form).

(D.26) Extensive Form Game

$\Gamma_E = [\mathcal{X}, \mathcal{A}, n, p, \alpha, \mathcal{H}, H, \iota, \rho, u]$, where 1-7 contain explanations of symbols. Additional definitions are underlined:

1. A finite set of nodes \mathcal{X} , a finite set of possible actions \mathcal{A} , and a finite set of players $\{1, \dots, n\}$.



The game you see in the picture above satisfies these requirements. The total number of nodes is very large (362,880), but nevertheless it is finite. So is the number of players – 2. The set of possible actions is also finite. There are 9 cells that players can put their marks in. With every stage of the game (layer of the "tree") the set of possible actions shrinks, but D.26 says nothing about it. The first set of conditions explains \mathcal{X} , \mathcal{A} , and n .

The second set of conditions explains p , and introduces concepts of initial (x_0), successor (s), terminal (T) and decision nodes.

2. A function $p: \mathcal{X} \rightarrow \mathcal{X} \cup \{\emptyset\}$ specifying a single immediate predecessor of each node x ; for every $x \in \mathcal{X}$ $p(x)$ is non-empty except for one, designated as the initial node, x_0 . The immediate successor nodes of x are thus $s(x) = p^{-1}(x)$. Please note that s does not have to be a function, because a node x may have many successors. All predecessors and all successors of x can be found by iterating operations p and s . It is assumed that for any x and for any

$k=1,2,\dots$ $p(x) \cap s^k(x) = \emptyset$ (if $s^k(x)$ is defined). The set of terminal nodes $T = \{x \in \mathcal{X} : s(x) = \emptyset\}$. All other nodes ($\mathcal{X} \setminus T$) are called decision nodes.


The third set of conditions explains α , and the concept of choice.

3. A function $\alpha: \mathcal{X} \setminus \{x_0\} \rightarrow \mathcal{A}$ giving the action that leads to any non-initial node x from its immediate predecessor $p(x)$ and satisfying the condition that: if $x', x'' \in s(x)$ and $x' \neq x''$, then $\alpha(x') \neq \alpha(x'')$ (different nodes are reached through different actions). The set of choices available at decision node x is $c(x) = \{a \in \mathcal{A} : a = \alpha(x') \text{ for some } x' \in s(x)\}$.

Information sets are probably the most complicated concept in extensive form games. You will find their formal definition below, but perhaps an example will clarify what they are.

4. A family of information sets \mathcal{H} and a function $H: \mathcal{X} \rightarrow \mathcal{H}$ assigning each decision node to an information set $H(x) \in \mathcal{H}$. Thus, the information sets in \mathcal{H} form a partition of \mathcal{X} . It is required that all decision nodes assigned to the same information set have the same choices available, i.e.: $H(x) = H(x') \Rightarrow c(x) = c(x')$. Hence choices available at information set H can be defined as $C(H) = \{a \in \mathcal{A} : a \in c(x) \text{ for } x \in H\}$.

The 'matching pennies' game is quoted in many textbooks as an illustration of what information sets are. The game is not very sophisticated. One player chooses Heads or Tails. The other one independently (perhaps simultaneously) chooses Heads or Tails as well. They disclose their choices. In case there is HH or TT, the second pays 1 to the first; otherwise the first pays 1 to the second. Independence (perhaps simultaneity) of both choices is reflected by the fact that the nodes where the second player makes the choice belong to the same information set. In the diagram ('game tree') this is reflected by the dotted (or broken) line.



I have no idea why they are called 'information sets'. To me a more adequate name would be 'disinformation sets', since the player who makes the move in nodes x_1 or x_2 that belong to the same information set $H(x_1) = H(x_2)$ does not know, whether the game goes through x_1 or x_2 . Items 5-7 are self-explanatory.

5. A function $i: \mathcal{H} \rightarrow \{0, 1, \dots, n\}$ assigning each information set in \mathcal{H} to a player (or to nature formally identified as the player 0), who moves at the decision nodes in that set. We can define the family of player's i information sets as $\mathcal{H}_i = \{H \in \mathcal{H} : i = i(H)\}$.

6. A function $p: \mathcal{H}_0 \times \mathcal{A} \rightarrow [0, 1]$ assigning probabilities to actions at information sets where nature moves and satisfying $p(H, a) = 0$ if $a \notin C(H)$, and $\sum_{a \in C(H)} p(H, a) = 1$ for all $H \in \mathcal{H}_0$.

7. A family of payoff functions $u=(u_1,\dots,u_n)$ assigning utilities to the players for each terminal node that can be reached, $u_i:T \rightarrow \mathfrak{R}$. In order to be consistent with the theory of expected utility, values of the functions u_i should be interpreted as Bernoulli utilities.

The terminology adopted for extensive form games allows to define a strategy. Please note that this is consistent with the same concept defined for normal (condensed) form games.

(D.27)

Let \mathcal{H}_i denote the family of information sets of the player i , \mathcal{A} – set of possible actions, and $C(H) \subset \mathcal{A}$ – the set of actions available at information set H . A strategy for the player i is a function $s_i: \mathcal{H}_i \rightarrow \mathcal{A}$ such that $s_i(H) \in C(H)$ for all $H \in \mathcal{H}_i$.

The following theorem establishes a relationship between both types of games analysed in the lecture.

(T.17)

For every extensive form game (D.26) there is a unique normal form game (D.20), but not *vice versa* (i.e. the same normal form game may correspond to several extensive form games).

Proof:

The payoff matrix is defined by payments taken from the terminal nodes. 'Paths' leading to the terminal nodes define strategies of the players. The not '*vice versa*' part is implied by the fact that the same outcome can be reached by applying the same strategies although in a different sequence.

The 'Matching pennies' game depicted in the tree above (page 55) corresponds to the following payoff matrix:

		Second	
		H	T
First	H	(1,-1)	(-1,1)
	T	(-1,1)	(1,-1)

Please note that the payoff matrix above contains all crucial information about the 'matching pennies' game. Its extensive form (on page 55 includes additional details (like who made the first move) that cannot be deduced from the matrix.

The next concepts show how to predict the results of finite games. The idea is basically to look at the last stage of the game, and to ask questions, what decisions must have been made in the earlier stages in order to get what is encountered in the later one. The idea starts with the 'sequential rationality principle' which states that at any stage of the game a player is not likely to choose a strategy leading to a lower payoff, if a better strategy is available.

(D.28) Sequential rationality principle

Each player's strategy should contain actions that are optimal for every node (taking into account other players' strategies).

(D.29) Finite game with perfect information

Every information set contains only one node and the number of nodes is finite.

Here perfect information means that at every stage of the game players know which node they are in (unlike in the 'matching pennies' game).


(T.18) Backward induction

Sequential rationality principle is satisfied if an optimum action for $p(x)$ is determined once an optimum action for x is determined (i.e. anticipating an optimum solution at x).


Proof:

For finite games with perfect information backward induction boils down to determining outcomes for each terminal node x , determining optimum actions for preceding nodes $p(x)$, assigning them payoffs that result from these optimum actions, and eliminating remaining strategies. This procedure is then iterated for earlier nodes $p(p(x))$ and so on, until all the nodes are exhausted (the initial node is achieved).


The following example is likely to explain the idea of backward induction. We look at a finite game consisting of three stages. There are three players. The first player has to choose between L and R, the second between A and B, and the third player – between C and D. In the second stage, player two moves if the first player chose R, or player three moves if the first player chose L.



By applying the sequential rationality principle the game can be reduced to the following one.



The third player (who moves as the last one) is not likely to play C, because his (or her) payoff would be 1 instead of 6, 0 instead of 1, and 0 instead of 7 (depending on what others chose). This simpler game can be further reduced to the following one, since the second player is not likely to play B, because his (or her) payoff would be -1 instead of 4.



When the first player is to choose between L and R then he (or she) is not likely to play L, because the payoff would be -1 instead of 5. Thus by applying the sequential rationality principle three times iteratively, one can predict the outcome (5,4,1). Strategies identified by the sequential rationality principle (and *backward induction*) are:

- $s_1 = R$,
- $s_2 = A$ if '1 plays R',
- $s_3 = D$ if '1 plays L' or if '1 plays R and 2 plays A',

Hence the Nash equilibrium is (R,A,D).

This procedure can be formalised in the following Zermelo theorem. Even though its proof is trivial, and requires nothing more than applying definitions, it is one of the most important results in game theory, and it suggests how certain games can be analysed.

(T.19) Zermelo Theorem

Every finite game with perfect information Γ_E has a Nash equilibrium in pure strategies which can be found using backward induction. Moreover if none of the players has identical payoffs in two different terminal nodes, this is the only Nash equilibrium that can be found in this way.

Questions and answers to lecture 8

8.1 The last table of the 'Tic-Tac-Toe' game (like on page 53) provides some information on what players did. What information included in the extensive form (like on page 54) cannot be deduced from this table?

The table reads

x	o	x
x	o	x
o	x	o

It contains useful information. First of all it shows that it was a tie (nobody won). Besides, by calculating the number of "x" marks (5) and the number of "o" marks (4) you know, that the cross-guy started and apparently made a mistake. But you do not know the sequence of choosing the cells. For instance the first move of the cross-guy could have been

x		

as it actually was, but it could have been also

	x	

This is not a very important information (the game ended as the final table indicates), but it cannot be deduced from the condensed form. On the contrary, the tree (that is the game in its extensive form) contains this information as well.

8.2 Does the 'matching pennies' game have a Nash equilibrium in pure strategies?

No. No matter which cell was chosen (HH, HT, TH, or TT) one player has a motivation to unilaterally change the decision and get 1 instead of -1. In mixed strategies the Nash equilibrium is $(1/2, 1/2; 1/2, 1/2)$ and the average payoff is 0.

8.3 How many nodes can an information set include?

Any number. In perfect information games all information sets consist of single nodes (there are no information sets linking two or more nodes). In the 'matching pennies' game (page 55) there is an information set consisting of two nodes. In more complicated games there can be information sets consisting of three or more nodes. The only requirement is that all decision nodes assigned to the same information set have the same choices available.

8.4 How does the nature choose (please refer to the condition 6 of D.26)?

Game theory is mainly about how people choose. According to the vNM theory of expected utility, we choose so as to maximise something. In D.26 we assume that choices are made by nature in some nodes. For example, the nature chooses whether it is going to rain or to be sunny; the nature chooses whether the coin tossed shows heads or tails; the nature chooses what symptoms a patient reveals; and so on. Unlike people who act so as to accomplish something, nature in D.26 acts randomly. Only probabilities can be attached to its "choices", no intentions.

8.5 The condition 7 in D.26 states that payoffs defined in terminal nodes, u_i , have to be understood as Bernoulli utilities. Yet in many examples that we analyse payoffs are understood as monetary outcomes. Does this introduce an important difference in applying the sequential rationality principle?

Not necessarily. The sequential rationality principle states that people choose a higher payoff rather than a lower one. Utility functions are monotonic with respect to money. Thus if something is more attractive in terms of utilities it is also more attractive in terms of money. It does not work in the opposite direction though. If the utility function has flat portions (more money does not increase the utility) there may be some problems. In backward induction (the most important application of the sequential rationality principle), reductions of games are based on the fact that decisions leading to lower outcomes are less likely to be chosen. This procedure fails, if payoffs recorded in terminal nodes do not allow for such an elimination. However, please see the exercise attached to the overheads to this lecture (E-8) in order to observe that even if some of the payoffs are the same, backward induction may still work.

8.6 Is the sequential rationality principle consistent with common sense?

I think it is. It basically says that – irrespective of what one planned (irrespective of any long-term strategy adopted) – if there is an opportunity to increase the payoff at any stage, players take advantage of this. However, D.28 includes the condition "taking into account other players' strategies". Any rational player cannot be expected to be blind and deaf to what others

are likely to do. On the contrary, he (or she) should assume that others are rational too, and they chose what is attractive for them.

8.7 Did Sherlock Holmes apply "backward induction?"

Yes, he did. He always started with the very final stage of a mystery. He tried to understand what it was caused by, and to guess what happened earlier. Having found the last detail, he repeated the procedure in order to understand what happened at earlier stages, until he reached the beginning and identified the person responsible for the entire sequence.

8.8 Does chess comply with the Zermelo theorem?

Chess is perfect information game; players do not move simultaneously, and they remember all the previous moves. The number of strategies is very large, but it is finite. Hence it seems that assumptions of the Zermelo theorem are satisfied, but in fact they are not. The theorem refers to the definition of an extensive form game Γ_E (D.26). The definition includes the condition that for any x and for any $k=1,2,\dots$ $p(x) \cap s^k(x) = \emptyset$ (if $s^k(x)$ is defined) (see page 55). It requires that there are no "loops" (after some moves you are exactly where you were). If a game has a "loop", even if it has a finite set of nodes, it cannot be concluded actually. Thus backward induction cannot be applied. But even if loops are excluded, chess is so complicated that there are no machine algorithms to win always.

9 – Subgame perfection

Everybody has an intuitive understanding of a subgame – it must be a part (maybe a stage?) in some other game. It turns out, that its definition has to be a bit more formal.

(D.30)

Subgame of an extensive form game Γ_E – a subset satisfying the following two conditions:

1. It starts with an information set consisting of a single decision node x , contains all subsequent nodes, $s(x)$, $s(s(x))$ and so on, and it does not contain other nodes;
2. If the node x is in the subgame then every $x' \in H(x)$ is in it as well (i.e. a subgame does not contain 'incomplete' information sets).

T.20 and T.21 are extremely easy to prove, when you look at the requirements of D.30.

(T.20)

The entire game Γ_E is also a subgame.


(T.21)

Every decision node in a finite game with perfect information (that is all information sets consist of single nodes) can 'start' a subgame

It would be helpful to refer to an example popular in economic textbooks.

Example I

A market is served by one firm called Incumbent (I). Another firm (F) contemplates entering the market. It has two strategies: "Enter" (E) and "Do Not Enter" (D). If it chooses D, the payoffs are (0,4) (its payoff is 0, and the payoff of the Incumbent is 4 as before). If it chooses E, the two firms will play a duopolistic game which boils down either to attacking (A) e.g. by a price war or to cooperating (C), e.g. *à la* Cournot or Stackelberg (or by sharing the market if they make a monopolistic collusion). If both attack then the payoffs are (-3,-1), if they cooperate then the payoffs are (3,1). If I attacks and F cooperates the payoff is (-2,-1), and if I cooperates and F attacks it is (1,-2). The story can be reflected by a game Γ_E described by the following tree:



If F chooses D, then the game ends, and there are no changes. If F chooses E, then I has two strategies: A or C. Choosing its own strategy, F does not know what strategy has been chosen by I. This is reflected by the two nodes where F moves belonging to the same information set (the two nodes are linked by a horizontal broken line). The game Γ_E has two subgames: the entire game and, say Γ_O , the oligopolistic rivalry part following the decision E taken by F. By T.21, the number of subgames is the difference between the number of non-terminal nodes (here 4) and the number of nodes 'hidden' in non-single node information sets (here 2). Hence the number of subgames is 2. Since by T.20 Γ_E is also its subgame, then Γ_O is the only other subgame (it can be called a 'proper' subgame).

The game Γ_E has the following strategic (normal) form Γ_N :

		Incumbent (I)	
		C if F enters	A if F enters
Entering Firm (F)	D&C	0,4	0,4
	D&A	0,4	0,4
	E&C	3,1	-2,-1
	E&A	1,-2	-3,-1

This strategic form game has three Nash equilibria (D&C,A), (D&A,A), and (E&C,C) with payoffs (0,4), (0,4), and (3,1) printed in bold, respectively. In order to check that (0,4) is a Nash equilibrium it is sufficient to observe that if F changes its decision (either D&C, or D&A) while I sticks to its decision (A), the payoff for F will be -2, or -1; at the same time if I moves from A to C unilaterally, then its payoff will not change. Hence neither of the firms has an incentive to move away unilaterally. Also (3,1) is a Nash equilibrium. To see this, please note that if F moves away unilaterally then its payoff goes down from 3 to 1 or 0. If I moves away unilaterally, then its payoff goes down from 1 to -1.

The subgame Γ_O has the following strategic form (payoff matrix):

		Incumbent (I)	
		C	A
Entering Firm (F)	C	3,1	-2,-1
	A	1,-2	-3,-1

The only Nash equilibrium in this subgame is (C,C). We see that there is a difference between (3,1) and (0,4). The former shows up in the subgame, while the strategy A of the incumbent does not show up in (C,C).

Let us look at the tree on page 61. The first two Nash equilibria found in Γ_E do not satisfy the sequential rationality principle. All terminal nodes are achieved after a decision of the entrant. By comparing -3, -2, 0, 1, and 3, the entrant will take the decisions E and C. Knowing this, the incumbent will take the decision C too. Thus its threat that it will attack if F enters is not credible.

Additional definitions need to be introduced in order to analyse various Nash equilibria, some of which are more likely to be observed than others.

(D.31)

A strategy profile $s=(s_1,...,s_n)$ in an n-person game Γ_E induces a Nash equilibrium in a subgame of this game, if actions defined by s for information sets of this subgame (understood as a separate game) make a Nash equilibrium in it.

(D.32) A strategy profile $s=(s_1,...,s_n)$ in an n-person game Γ_E is called Subgame Perfect Nash Equilibrium, SPNE, if it induces a Nash equilibrium in its every subgame.

These definitions allow us to state T.22. This is an obvious corollary of the Zermelo theorem (T.19).

(T.22)

Every finite game with full information Γ_E has an SPNE in pure strategies. Moreover, if none of the players has identical payoffs in two different terminal nodes, this is the only SPNE.

This theorem explains why (E&C,C) was the most likely result expected in example I. Unlike other Nash equilibria, this is an SPNE.

Neither of the Nash equilibria that were not SPNEs in the example I above seemed realistic, since the threat of choosing A by I (if F enters) was not credible. The potential entrant may suspect that when confronted with its presence in the market, the incumbent will cooperate despite threats. Hence such an entry deterrence strategy of the incumbent is not effective.

Note:

To be credible, a threat must include a decision that will be taken irrespectively of the 'sequential rationality principle'. This is called a pre-commitment. Examples of pre-commitment include: (in military planning) automatic destruction of one's own infrastructure in the case of attack, and (in economics) lack of possibility of changing a business decision before a certain date.

Let me elaborate on pre-commitments starting with the military example.

In Switzerland everything is undermined. There are explosives under every tunnel, every highway, and every bridge. This should deter from attacking the country. Any aggressor knows that if attacked, the Swiss will damage all the country's infrastructure. They will hurt themselves, but the aggressor will be hit as well. Thus attacking Switzerland does not make sense.

Is this threat credible? This is questionable; some students doubted whether it is. Given the sequential rationality principle it is doubtful whether the Swiss – when attacked – will decide to hurt themselves. They would like to prevent the attack, of course, but if it happened, then it would be better to cooperate with the aggressor rather than to take desperate steps. Hence the threat can be considered not credible. In order to be credible, the threat should be based on a pre-commitment. If the action to explode everything is to be taken automatically rather than following a decision, the threat must have been considered credible.

Now the economic example. There are negotiations between a seller and a buyer. The seller offers something for the price of 80€ per piece, and the potential buyer says that 60€ is the maximum acceptable price. The seller warns that unless the buyer pays 80€ today, there will be no opportunity to buy over the next 30 days. Is this threat credible? It depends on whether it is backed by a pre-commitment. The potential buyer may consider the threat a negotiation trick: "the seller insists on 80€, but if I do not yield to this pressure, soon the price will go down somewhat".


But the seller can demonstrate a decision of the company's supervisory board which authorises the sale for the price not lower than 80€. If the price is to be lowered, then a new decision of the supervisory board needs to be taken. But according to the company's by-laws, the supervisory board meetings have to be announced at least 30 days ahead. Therefore, indeed, there is no possibility of lowering the price over the next 30 days. The threat is credible.

Game theory developed a number of concepts which are helpful in our analyses of economic behaviour. At the same time, experiments demonstrate that some of these concepts are not confirmed in actual behaviour. The sequential rationality principle and backward induction serve as examples. They provide convincing arguments that rational players take decisions based on their payoffs rather than on cooperation while taking into account the payoffs of others. But experiments show that they trust each other and they cooperate. The following example illustrates the problem.

Example II (Centipede game)

This 2-person game lasts $2k$ periods (where $k=1,2,3,\dots$), it has $2k$ decision nodes, and $2k+1$ terminal nodes. The players move in turns, starting with the player number 1. In every node either of the two decisions can be taken: Stop (S) or Continue (C). If the game is stopped at node $i=1,2,3,\dots,2k$ then the payoffs are: $P(i)=i-(1+(-1)^i)$, $D(1)=0$, and for $i>1$ $D(i)=i-(1-(-1)^i)$. If the game is not stopped by any player at any node, it terminates with the payoffs $P(2k)=2k$, and $D(2k)=2k-1$ in the period $2k$.

The formulae are complicated, but everything becomes clear if a picture is made. For $k=4$ the game is reflected by the following tree (the name 'centipede' refers to its shape):



This is not a typical tree, because it grows horizontally (from left to right). The first decision is taken by the player 1, the second – by the player 2, the third – by the player 1 again, and so on. At any stage (node) each player receives a higher payoff if he (or she) stops the game than if the game is stopped by the other player in the next node (stage). By the sequential rationality principle, player number 2 is better off if he (or she) stops the game in its ultimate ($2k$ th) decision node. Knowing this, player number 1 has an incentive to stop the game in its penultimate ($(2k-1)$ th) decision node. Knowing this, player number 2 has an incentive to stop the game even earlier, i.e. in the $(2k-2)$ th node, and so on. By the sequential rationality principle (and backward induction), the game should be stopped in its first node.

Experiments do not confirm the outcome predicted by the sequential rationality principle in this case. Hence in real life people's behaviour is more complicated than some game theoretic models assume. Apparently people trust each other. If they expect that the game will last long, they will get more by its end than in its beginning. Nevertheless there is always a risk that the game will be stopped by the other person, and then the person who did not stop it earlier regrets, because the payoff could have been higher.

Nash equilibrium is the most important concept in game theory. It helps to explain and interpret certain economic phenomena, but there are two problems with it. The first problem is its non-uniqueness. There are games with multiple Nash equilibria. Economists – who expect that people make decisions consistent with the game theoretic framework – do not know which Nash equilibrium is then more likely than others. SPNE, that is a Nash equilibrium which survives subgames, is deemed to be better reflecting real life circumstances. The second problem is whether people behave like the expected utility theory assumes. It turns out that sometimes they do not; they trust each other, they do not confine to their own well-being, and they care for what the others gain from joint activities.

Questions and answers to lecture 9

9.1 How many subgames does the 'matching pennies' game have?

The 'matching pennies' game is its only subgame (it does not have any proper subgames). You can prove it using the theorems T.20 and T.21. The number of decision nodes is three, but two of them are 'hidden' in an information set, so they cannot start a subgame. Thus only the initial node can start a subgame. The same can be proved directly from the definition D.30. A subgame can start from the initial node, and it cannot start from any subsequent decision nodes since they belong to an information set consisting of more than one node. Intuitively the 'matching pennies' game does not have any smaller subgames, because what the second player does has to be confronted with the decision of the first player.

9.2 In the 'entry deterrence' game in its "normal" form the F has four strategies: D&C, D&A, E&C, and E&C. Do the first two of them make sense?

If the F does not enter then what it plans to do after entering is not relevant (the payoffs are zero for the entrant in both cases). Nevertheless listing D&C, D&A separately does not have a formal meaning only. Whether the entrant seems aggressive or not may influence the incumbent's choices in the future. From the point of view of vNM theory, only the expected payoffs count. Yet experiments we referred to demonstrate that actual behaviour of players can be more complex.

9.3 Why cannot subgames start from a two-node information set?

Even though each node in a two-node information set may have some successors, decisions made in this node take into account the fact that the player who moved did not know whether it was this node or the other one from the same information set. Thus the logic of the game requires that earlier decisions must be taken into account too.

9.4 Please provide examples of Nash equilibria that contain strategies which do not survive taking subgames.

The 'entry deterrence' game analysed in the class has two Nash equilibria that belong to this category. The A strategy of the incumbent (being aggressive) may have some deterring advantage. Perhaps some potential entrants may take the decision D (Do not enter) because they do not want to risk. Nevertheless this strategy is detrimental for the incumbent if the subgame is played (if the entrant took a decision E). Hence the A strategy does not survive in the second stage.

9.5 What will happen in the 'entry deterrence' game, if the oligopolistic competition in its second stage is *à la* Bertrand?

In the Bertrand model oligopolists compete by prices (not by quantities, like in the Cournot model). In the Bertrand model the only equilibrium price is the same as in the competitive markets: if firms have identical cost curves, the price is equal to the marginal cost. In the long run their profits must vanish. This will change payoffs in the second stage and perhaps will motivate the entrant not to enter.

9.6 Can sequential rationality principle be reconciled with pre-commitments?

No. Pre-commitments do not allow making choices in every node based on comparing the payoffs. Choices are made at earlier stages. They may influence choices to be made by other players. But the player who made a pre-commitment cannot choose any more. Therefore sequential rationality principle does not work (as stated in its definition explicitly). By making pre-commitments players think of course about potential payoffs, but they do not look at separate nodes necessarily.

9.7 How can the centipede paradox be explained?

The paradox is that experimental results contradict the sequential rationality principle. According to the principle, backward induction suggests that the first player should stop rather than continue, but experiments demonstrate that players can continue the game until the

end. To me the most convincing explanation is that the potential gain by stopping the game before the other player does it is 1, 2, or 3. In contrast, the potential gain by continuing the game for, say, 20 rounds is many times higher. Consequently players are ready to take the risk of being cheated in order to obtain a much higher reward. This mechanism is reinforced by trust that some people may reveal with respect to fellow players. Especially if you know the person who plays with you, you are likely to trust more than when you do not know whom you play with.

9.8 In the centipede game the joint payoffs increase by 2 in every round (except the first and the last one when they increase by 1). Can this be a reason that people continue rather than stop?

Probably yes. So-called side payments will be analysed in lectures 12 and 14 in a more detailed way. Here let us simply hypothesise that players can observe the fact of growing joint payoffs. Consequently they can make an agreement that whoever stops the game will share the extra gain with the other player (who did not stop the game). If they have a mechanism to enforce such an agreement, or if they trust each other, they can continue the game for a long time.

10 – Repeated games I

Games analysed so far were usually one-shot games. The concept of mixed strategies assumed that a game was played over and over again with the same payoff structure. Now we are going to look at the same game to be played repeatedly. However, instead of calculating expected payoffs, we will calculate present values of payoffs that will be received in the future.

To this end, we have to recall the idea of discounting. If I asked you to give me 100 euro today, and promised to give you 100 euro back one year from now, most of you would say "no" (even in the class today students confirmed that they would not accept such a deal). If I changed my question: "you give me 100 euro today, and I give you 130 euro next year (no inflation, no risk of losing the money); do you agree?" Some of you will probably say "yes" (confirmed in the class too). Let us then bargain. I ask: "you give me 100 euro today, and I give you 101 euro next year". You say "no". I say "you give me 100 euro today, and I give you 125 euro next year". You say "yes". I say "you give me 100 euro today, and I give you 102 euro next year". You say "no". And the negotiation process can go on. If there is an amount, say, 105 euro to be paid back after a year, and you say "I am indifferent – I can agree to such a deal, but I do not insist; I can do without such a deal either", then we say that 5% is the discount rate applied by the person who agreed.

The formula for finding equivalents of money amounts in different time moments can be iterated. In general, it is:

$$X_t = X_0(1+\rho)^t, \text{ or } X_0 = X_t/(1+\rho)^t,$$

where ρ is a discount rate. Please note that in my examples $t=1$ (the time span between "now" and the "future" is one year). The formula then reads:

$$X_1 = X_0(1+\rho)^1, \text{ or } X_0 = X_1/(1+\rho)^1,$$

For two years it will be

$$X_2 = X_0(1+\rho)^1(1+\rho)^1 = X_0(1+\rho)^2, \text{ or } X_0 = X_2/(1+\rho)^2,$$

and so on. The formula linking X_t with X_0 is called the "Present Value" (PV) of X_t in moment $t=0$. If it is applied to X_1 , X_2 , X_3 , and so on (cumulatively), the formula of PV reads:

$$PV = X_0/(1+\rho)^0 + X_1/(1+\rho)^1 + X_2/(1+\rho)^2 + \dots + X_T/(1+\rho)^T,$$

where T is the last year that the decision (project) is expected to imply a cost or a benefit (costs are entered with "minus" signs, and benefits are entered with "plus" signs). If benefits and costs belong to different time periods, then we need discounting to make them comparable. Namely – before adding sums from different time moments – we recalculate everything in present value terms, as if everything was to happen in the beginning (at the time of the analysis), in $t=0$.

The most common economic examples of repeated games refer to oligopolistic competition, namely to Bertrand and Cournot models. We will start with Bertrand, since calculations are simpler. Competition *à la* Bertrand means that firms compete by prices (quantities are less important; they can be easily adjusted to the demand). In this type of competition between firms characterised by identical cost functions the only equilibrium price – like in perfectly competitive markets – is equal to the marginal cost. If a firm charges a price higher than that, it can be undercut by another firm offering a somewhat lower one but higher than the marginal cost. Thus the only Nash equilibrium (no firm has a motivation to change the price unilaterally) is the marginal cost price.

In what follows we will assume for simplicity that there are only two firms 1 and 2 (it is a duopoly rather than a more complicated oligopoly), and they have the same linear cost functions with: $AC_1=AC_2=MC_1=MC_2=c=\text{const}$. If they compete *à la* Bertrand, their profits vanish, but they can collude to enjoy the monopolistic profit. In order to compare alternative competition strategies they need to calculate present values of profits obtained in various years, as stated in the following example.

Example I (Sequential Bertrand model)

A duopoly plays a Bertrand game repeatedly. After each period the rivals $j=1,2$ can (simultaneously) announce new prices. It is assumed that they maximise the Present Value (with the discount rate $\rho>0$) of their profits:

$$\pi_{j0} + \pi_{j1}/(1+\rho) + \pi_{j2}/(1+\rho)^2 + \pi_{j3}/(1+\rho)^3 + \dots$$

The game can be finite (if firms play it T times) or infinite.

Their strategies are characterised by the following theorem.

(T.23)

In the sequential Bertrand model, if the game is finite, then there is only one SPNE; it consists of Nash strategies from a single Bertrand model, i.e. (c,c) , where $c=AC_1=AC_2=MC_1=MC_2$.

Proof:

No price $p < c$ can be sustained since it implies losses. Yet no price $p > c$ can be sustained either, since each firm has a motivation to offer a price higher than c but lower than p in order to undercut the rival's price, get all the clients, and increase the profit.

T.23 explains what happens if firms do not collude. If they collude in order to enjoy the monopolistic price then – again by T.23 – each firm can undercut the rival's price, get all the clients, and increase the profit. So-called Nash reversion strategy is a mechanism providing incentives for keeping promises rather than cheating. Rivals can encourage themselves not to cheat by threatening to go back to Nash strategy and thus losing the profit. Its formal definition is given below.

(D.33) Nash-Bertrand reversion strategy

In the sequential Bertrand model we define $H_{t-1} = \{(p_{11}, p_{21}), (p_{12}, p_{22}), (p_{13}, p_{23}), \dots, (p_{1t-1}, p_{2t-1})\}$ for $t > 1$, i.e. a history of strategies applied. In the t th stage the rivals choose $p_{jt}(H_{t-1}) = p^M$ if $t=1$ or H_{t-1} consists of (p^M, p^M) only. Otherwise $p_{jt}(H_{t-1}) = c$. p^M is the monopolistic price, i.e. the price giving the two firms the maximum profit available jointly for the suppliers. In the standard model (two identical firms as in T.23, linear demand curve $p = a - by$) the monopolistic price is $p^M = (a+c)/2$. The joint supply is $y^M = (a-c)/(2b)$, and the joint profit is $\pi^M = (a-c)^2/(4b)$.

It is easy to see that the effectiveness of Nash-Bertrand reversion strategy depends on the discount rate. If a firm cheats then it loses all future profits. If the discount rate is sufficiently large then this loss is lower than a one-time extra profit achieved thanks to breaking the agreement to charge the monopolistic price, that is undercutting the rival's price, and getting all the clients. If the discount rate is low then cheating is not attractive, as summarised in T.24.

(T.24)

In the sequential Bertrand model, if the game is infinite, then Nash-Bertrand reversion strategy is SPNE if and only if $\rho \leq 1$.

Proof:

If firms collude to form a monopoly, their joint profit is $\pi^M = (a-c)^2/(4b)$. Each firm gets a half of this amount, that is $(a-c)^2/(8b)$. By cheating (undercutting the price by a small amount), one can almost double the profit by claiming the entire market. The cheating firm can increase its profit almost by $(a-c)^2/(8b)$. Let us call it π . At the same time it will lose all the future profits whose present value is $\pi/(1+\rho) + \pi/(1+\rho)^2 + \pi/(1+\rho)^3 + \dots$. This is an infinite geometric series with the first element $a_1 = \pi/(1+\rho)$, and the quotient $q = 1/(1+\rho)$. Students should recall from the high school that its sum is $a_1/(1-q)$, that is π/ρ . As long as $\rho \leq 1$, cheating is not attractive (the present value of lost profits is higher than the one-time extra profit).

Please note that $\rho \leq 1$ means that the discount rate is not higher than 100%. Typical discount rates are much lower than 100%. Thus Nash-Bertrand reversion strategy seems to be a good solution. It is also easy to prove that not only a monopolistic price can be sustained in this way. Any price not higher than the monopolistic one can be sustained as well (T.25).

(T.25)

In the sequential Bertrand model, if the game is infinite and $\rho \leq 1$, then each fixed and common (for both players) selection of the price $p \in [c, p^M]$ backed by the Nash-Bertrand reversion strategy is an SPNE.

By the same token, if the discount rate is very high cheating is attractive. As a result, sustaining any price higher than the competitive one (c) is impossible in this case even when backed by Nash-Bertrand reversion strategy (T.26).

(T.26)

In the sequential Bertrand model, if the game is infinite and $\rho > 1$, then the only SPNE consists of $p=c$ to be selected by both players in each stage.

Oligopolists can compete with quantities rather than prices. Competition *à la* Cournot means that firms compete by quantities. These cannot be easily adjusted to the demand; suppliers have to determine them in advance. In what follows we will assume for simplicity (as before) that there are only two firms 1 and 2 (it is a duopoly rather than a more complicated oligopoly), and they have the same linear cost functions: $AC_1=AC_2=MC_1=MC_2=c$. If they compete *à la* Cournot, their profits do not vanish, but they are less attractive than for a monopoly. If they collude they can enjoy the monopolistic profit. In order to compare alternative competition strategies they need to calculate present values of profits obtained in various years, as stated in the following example.

Example II (Sequential Cournot model)

A duopoly plays a Cournot game repeatedly. After each period the rivals $j=1,2$ can (simultaneously) announce new supplies. It is assumed that they maximize the Present Value (with the discount rate $\rho > 0$) of their profits: $\pi_{j0} + \pi_{j1}/(1+\rho) + \pi_{j2}/(1+\rho)^2 + \pi_{j3}/(1+\rho)^3 + \dots$. The game can be finite (if firms play it T times) or infinite.

The definition of the Cournot sequential game replicates the definition of the Bertrand sequential game (example I). T.27 is based on the basic microeconomics facts about how duopolists compete by quantities.

(T.27)

In the sequential Cournot model, if the game is finite, then there is only one SPNE; it consists of Nash strategies from a single Cournot model, i.e. (y^C, y^C) , where in the standard Cournot duopoly ($AC_1=AC_2=MC_1=MC_2=c=\text{const}$) with a linear demand curve $p=a-by$, $y^C=y_1=y_2=(a-c)/3b$. The total supply is thus $y=2(a-c)/(3b)$, the price is $p^C=(a-2c)/3$, and the joint profit is $\pi^C=2(a-c)^2/(9b)$; each firm makes a half of this, i.e. $(a-c)^2/(9b)$.

Proof (standard – single round – Cournot model):

In a single round duopolists solve the following problem: if the price is implied by the total supply that is $p=a-b(y_1+y_2)$, and the first rival makes a quantity decision that maximizes its profit expecting that the second rival does the same, they solve two simultaneous maximization problems:

- $\max_{y_1} \{ (a-b(y_1+y_2))y_1 - cy_1 \},$
- $\max_{y_2} \{ (a-b(y_1+y_2))y_2 - cy_2 \}.$

First Order Conditions (FOC) for these problems are:

- $a - 2by_1 - by_2 - c = 0$,
- $a - 2by_2 - by_1 - c = 0$.

Solving these equations yields

- $y_1 = y_2 = (a-c)/3b$,
- $y = 2(a-c)/3b$ (the total supply),
- $p = (a+2c)/3$.

The profit of each firm is $(a-c)^2/(9b)$. The derivation of it proves that neither of the firms has an incentive to unilaterally move away. Thus they are in Nash equilibrium if they supply exactly what the solution suggests: $y_1 = y_2 = (a-c)/3b$.

Now we can define the Nash-Cournot reversion strategy in the same way as in D.33.

(D.34) Nash-Cournot reversion strategy

In the sequential Cournot model we define $H_{t-1} = \{(y_{11}, y_{21}), (y_{12}, y_{22}), (y_{13}, y_{23}), \dots, (y_{1t-1}, y_{2t-1})\}$ for $t > 1$ i.e. a history of strategies applied. In the t th stage the rivals choose $y_{jt}(H_{t-1}) = y^M/2$ if $t=1$ or H_{t-1} consists of $(y^M/2, y^M/2)$ only. Otherwise $p_{jt}(H_{t-1}) = y^C$. $y^M/2$ is the half of the monopolistic supply, i.e. the (joint) supply giving the two firms the maximum profit available jointly for the suppliers. In the standard model (two identical firms as in T.27, linear demand curve $p = a - by$) the monopolistic joint supply is $y^M = (a-c)/(2b)$, and the resulting price $p^M = (a+c)/2$. The joint profit is then $\pi^M = (a-c)^2/(4b)$.

Cheating in the sequential Cournot model is possible as well, but it requires more sophisticated activities. At the same time, the penalty for cheating – that is Nash-Cournot reversion strategy – is less harmful (it lowers the profit rather than eliminates it completely). Consequently duopolists do not require the same very high discount rate (like in the Bertrand case) in order not to have an incentive to cheat. Nevertheless it can also be proved that if a discount rate is very high, the Nash reversion strategy does not work. We will quote the relevant theorems without proofs.

(T.28)

In the sequential Cournot model, if the game is infinite, then Nash-Cournot reversion strategy is SPNE if $\rho \leq 8/9$

(T.29)

In the sequential Cournot model, if the game is infinite and $\rho > 8/9$, then the only SPNE consists of $y = y^C$ to be selected by both players in each stage.

In some game theory textbooks (and in some other mathematical textbooks) so-called folk theorems are quoted. The name – 'folk theorem' – is applied in order to denote a useful theorem that everybody relies on, but so trivial that nobody bothers to prove and/or to authorise with one's name. Several game theory results are nick-named 'folk theorems'. In some translations into non-English languages these folk theorems are called Mr. Folk theorems. This is inadequate, because the word "folk" is not a noun, but rather an adjective meaning "widespread" or "popular".

Folk theorem (example)

In a finite game, by backward induction one can prove that the only SPNE consists of non-cooperative decisions (cheating, that is violating promises to cooperate).

This is an example of the fact that everybody knows of (for instance, think of the centipede game), but nobody bothers to demonstrate it formally (T.18 proves only a part of it). Folk theorems are used in game theory, but one needs to pay attention to whether they can be applied in a given context.

As the difference between an infinite game and a finite one with a large number of stages is difficult to appreciate in an experimental setting, while the consequences are drastically different (as evident by comparing the Folk Theorem above with previous theorems), economists introduce the concept of finite games of unknown duration (see exercise E-10 in the overheads to this lecture).

In examples I and II, 'Nash reversion strategies' involved 'punishing' the non-cooperating rival by switching to a Nash (non-cooperative) solution for ever. An alternative concept is to 'punish' the non-cooperating rival only once by choosing a Nash solution and then moving back to a cooperative one from the next stage on. There are empirical findings which indicate that this other 'Nash reversion strategy' provides (on average) higher payoffs for the players.

In the 1970s, Robert Axelrod, an American mathematician and political scientist, raised the question of the most effective scheme to encourage cooperation. Without cooperation, economic agents are likely to apply Nash strategies which can lead to very bad outcomes (see prisoner's dilemma). If they cooperate to improve the outcome, they are vulnerable to cheating, that is to an un-cooperative behaviour. As theorems T.24, T.25, and T.28 demonstrate, the threat of Nash reversion can be helpful. But there is an empirical question of what such threats imply. A Nash reversion strategy can be very expensive. It provides an incentive not to cheat, but if one player cheated then both players are doomed to a Nash equilibrium for ever. Perhaps other incentives can prove to be also helpful but less expensive.

Robert Axelrod decided to test it empirically. Anatol Rapoport (1911-2007), an American mathematical psychologist, demonstrated that a strategy called 'Tit for Tat' was found to be the most successful one; it provided an incentive not to cheat, and – at the same time – it was not so expensive as Nash reversion. 'Tit for tat' requires a punishment for the non-cooperative behaviour, but – unlike the original Nash reversion strategy – not for ever but only once. In other words, if one player broke the promise (it does not matter whether deliberately, or by a mistake), he (or she) is punished by the other player by applying the Nash strategy in the next round only. After this, that is after the one-time punishment, the other player resumes playing the cooperative strategy that was agreed upon.

In both examples (I and II above) there does not have to be a formal collusion between the rivals. Nevertheless their choices are similar and they may create an impression that the rivals colluded. That is why this type of behaviour is sometimes referred to as 'tacit collusion'.

Questions and answers to lecture 10

10.1 Does the speed of repeating the games have impact on strategies chosen?

Yes. As theorems T.24, T.25, and T.28 demonstrate, strategies may depend on the discount rate applied (because calculating present values requires discounting). The rate depends on the time interval. If a rate of, say, 4% refers to one year, then the rate applicable to 1 quarter is 1% approximately (strictly speaking it is somewhat less, namely $100(1.04^{1/4}-1)$). Therefore, when duopolists contemplate whether to cheat or not, if one round lasts 3 months, they will refer to 1% rather than 4%.

10.2 When do duopolists compete *à la* Bertrand, and when do they *à la* Cournot?

In the Bertrand model they compete by prices while in the Cournot model – by quantities. The textbook example of the former is the software market, while that of the latter: Boeing *versus* Airbus.

Let us start with software. Some time ago, computer programmes were sold on CDs. Burning a CD costs 1 eurocent or so, which was negligible. Now they are sold in the internet; you have to pay something and then you can download it to your computer. The programme had to be uploaded on the seller's server only once, and it does not matter, whether it was downloaded by users 1 thousand times or 1 million times. The cost does not change for the seller. One crucial decision to be taken by the seller is the price, while the quantity can be arbitrary; it will be determined by the market.

Boeing with its main production facility in Seattle is the most important American airplane supplier, and Airbus with its main production facility in Toulouse is the most important European airplane supplier. The two firms control the global market of wide body passenger jets. There are some other producers, but their market shares are much smaller. The global demand for large airplanes is very limited, maybe 500 per year. If Boeing and Airbus are to produce, say, 300 each per year, the market will not absorb such a supply at a satisfactory price. The quantity decisions are irreversible. If the global demand turns out to be higher than expected, say, 800 per year, neither Boeing nor Airbus can increase their production rapidly, because investing in appropriate capacity takes time. If they built facilities to produce less, they cannot produce more. Can they produce less than the capacity permits? Not really. The reason to decrease production was that the market did not accept the price the producer asked for. If the price is going to be higher (and it must be higher, because the cost of investing in capacity has been already spent), the demand will be even lower. In other words, the quantity cannot be lowered without incurring financial losses.

The Bertrand and Cournot models differ in the way the producers compete with each other. In the former they set prices. Whoever sets a lower price wins all the clients. In the latter they decide on quantities supplied, and the market determines the price that both of them have to face. They can be surprised by this price, but it is too late to change their quantity decisions. In order to use either of the models, economists have to determine whether the duopolists compete by prices or by quantities. If quantities can be adjusted easily, Bertrand model is perhaps more adequate. Otherwise, the Cournot model is more suitable. There are also mixed models that require more sophisticated mathematical treatment.

10.3 Can cheating be observed in sequential Bertrand models in real life situations?

Some students doubt whether Bertrand models reflect what they observe in grocery stores. Neighbouring stores offer similar prices, but there are no symptoms of price undercutting aimed at "stealing" clients. It seems that competitors colluded (at least informally), and they


keep prices high even when they could decrease them down to the marginal cost level. They seem to play a sequential Bertrand game without cheating.

Many grocery stores announce price rebates in daily newspapers, but these are not "regular" price decreases; they are announced as exceptions to the "regular" price. If a grocery is accused by competitors of lowering the price below what can be expected (or what was agreed to), it can explain that there was no price decrease announced; it was just a rebate offered in a marketing campaign which is perfectly justified.

Even though explicit price undercutting incidents are rare, systems of rebates may serve the purpose as well. One of the funniest rebates I saw was a rebate available to buyers whose grandmother had her birthday. If the grocery that announced it were accused by competitors for lowering the price, it could respond that it was available only to those whose grandmother had a birthday. It would be difficult to prove that clients whose grandmothers did not have their birthday took advantage of the lower price as well. Economic analyses of market behaviour should go beyond what the words suggest.

10.4 Christ (*Mt* 25:16) observes that in a long run it is possible to double one's wealth as a result of business operations. Please calculate the relationship between time and rate of return.

The parable of talents, i.e. bags of gold (*Mt* 25:14-30), alludes to the fact, that if you invest your money adequately, it is possible to get twice as much. But it does not refer to a specific rate of return. Neither does it refer to a specific time horizon. It only says that $2x = x(1+r)^T$, where x is the amount of money, r – rate of return, and T – time horizon. This equality reduces to $2 = (1+r)^T$, because x cancels out. It can be solved by taking logarithms of both sides: $\ln 2 = T \ln(1+r)$, and ultimately $T = (\ln 2) / \ln(1+r)$. Alternatively, one can calculate r as a function of T . To this end, exponents have to be taken of both sides of the equation $\ln(1+r) = (\ln 2) / T$. After rearrangements: $e^{\ln(1+r)} = e^{(\ln 2) / T}$, that is $1+r = e^{(\ln 2) / T}$, i.e. $r = e^{(\ln 2) / T} - 1$. This relationship looks like in the picture below (r is measured in percentage points):



By looking at the picture, you can easily see that when $T=1$, then $r=100\%$ ($r=1$), and when $T=10$, then $r=7\%$ ($r=0.07$). The first case is obvious: if the rate of return is 100% then you can double your wealth in just one year. The second case is less obvious: if the rate of return is 7% then doubling your wealth requires ten years ($1.07^{10} = 2$). The rate of return of 100% is very rare. The rate of return of 7% can be observed much more frequently. If such a rate is available, it takes 10 years to double one's wealth.

Please note that in this exercise r means 'rate of return' rather than 'discount rate'. The latter corresponds to the former in the case of a risk-free investment opportunity. The parable does not say anything about whether business operations were risky or not.

10.5 Why is a Cournot solution less attractive than a monopolistic collusion between two competitors?

In the Cournot equilibrium (under the assumptions adopted in this lecture) firms enjoy the profit of $(a-c)^2/(9b)$. The monopolistic profit is $(a-c)^2/(4b)$. If the firms share it equally, then each enjoys $(a-c)^2/(8b)$, that is more than in the Cournot equilibrium.

10.6 Statement "In a finite game, by backward induction one can prove that the only SPNE consists of non-cooperative decisions" is nicknamed Folk Theorem. Why?

Backward induction is an easy way to predict what decisions players are likely to take. We start from the last node. If there is an opportunity to increase the payoff at the expense of the rival expecting that he (or she) will reciprocate in the future, we will do it, because there is no future. If we are likely to take a non-cooperative decision in the ultimate node, the rival – who is a rational person – will not take a cooperative decision in the penultimate node. By iterating this argument many times, we reach the initial node always concluding that a cooperative decision is not likely. The proof is so trivial that no people would like to associate their names with such a theorem. That is why it is called a Folk Theorem.

10.7 'Tit For Tat' won the tournament announced by Axelrod. Is it possible that 'Tit For Two Tats' turns out even better?

Yes. The original 'Tit For Tat' proved to be effective in sequential games of the 'Prisoner's Dilemma' type. If players play such a game sequentially, and they promised to deny always, then if one player confesses, the other player confesses in the next round, but he (or she) starts denying after this. 'Tit For Two Tats' is a strategy that penalises two defections in a row. If one player defects only once, there is no retaliation. If he (or she) defects twice, then retaliation comes.

It is an empirical question whether this alternative 'Nash reversion' strategy is effective or not. It is "more forgiving" in a sense that one defection goes unpunished. One cannot preclude that a single defection is a mistake rather than a deliberate decision. Hence leaving it unnoticed may be a good solution. Nevertheless motivating to cooperate by alternative strategies depends on a particular type of a game. 'Tit For Tat' proved effective for the 'Prisoner's Dilemma' type. Other types may find other strategies to be better.

10.8 Should 'tacit collusions' be penalised like other collusions by anti-trust laws?

No, but it is very difficult to check if the 'tacit collusion' results from an illegal agreement, or from a natural logic of competition. If firms take similar (perhaps even identical) decisions, this may result not from an agreement, but rather from the common business methods applied. Of course, if there is evidence that the firms agreed to share the market (which is illegal) then penalties can be imposed. But if the firms collude then they try to hide this carefully. The best example of such a collusion that can go undetected for decades is the famous moon-phase scandal. Several American construction firms agreed on how to win government procurement contracts without offering realistic costs. They agreed that the winner is indicated by the moon-phase of the moment that the government announces a call. Moon-phases change regularly, but their period is not a rational number (it is approximately 29.5 days, but the exact number is different – the same phase may start on another week day and month day next year). The firm indicated to win declared cost that was higher than its marginal cost. But other

firms declared cost that was much higher, so the indicated firm always won the contract as the cheapest one. Authorities did not find anything suspicious, because there was no correlation between who wins and the day of the week or month. It would have never been detected unless one firm leaked the information.

The most popular way to detect a collusion is to prove that the price asked by the seller is higher than the marginal cost (suggesting that it could have been lower) without collecting evidence that the collusion resulted from an illegal agreement between competitors. But large firms have such complicated accounting systems that it is very difficult to demonstrate that the price is excessive. Besides – as the Cournot model shows – firms can determine similar quantities without consulting each other.

11 – Repeated games II

This lecture is devoted to refining the concept of Nash equilibrium in a dynamic setting. Its first part is around what in English is meant by 'bygone is bygone'. It means that the past should not determine the future. You cannot change the past, so it is better to check what the future brings.

The first ideas to be introduced are so-called Pareto domination and perfection. Vilfredo Pareto (1848-1923) was a talented Italian sociologist and mathematician, who made numerous contributions to economics. Both ideas are based on the Pareto improvement principle, that is checking whether it is possible to improve the situation of one individual without affecting somebody else adversely. If this is possible, we have not achieved a Pareto optimum yet.

Most students are perhaps familiar with the definition of a Pareto optimum. Let me recall it briefly just in case somebody forgot it. A Pareto optimum is such an outcome that nobody can be made better off unless somebody is made worse off. Assuming that there are 2 apples to be allocated to 2 people (both of them like apples), both people can be given 1 apple, respectively. This is a Pareto optimum, since neither can be given 2 apples unless the other person is deprived of an apple (but this would make him or her worse off). The allocation can be different. One privileged person can be given 2 apples while the other one is left without any. Surprisingly, this is a Pareto optimum too, since in order to make the deprived person better off, at least 1 apple should be taken away from the privileged person making him (or her) worse off.

Different allocations – some of them considered unfair – can satisfy the definition of a Pareto optimum. To see an allocation which violates it, let us consider 2 apples distributed in the following way. 1 apple goes to one person, the other person gets nothing, and 1 apple is left aside. This is not a Pareto optimum, because one person can be made better off by getting the apple that was left aside, and nobody will be made worse off. Intuitively, a Pareto optimum can be understood as an allocation which uses up everything what is available (not necessarily in a fair way).

(D.35) Pareto domination

A pair of strategies (i_1, j_1) is Pareto-dominated by (i_2, j_2) , if $P_{i_2 j_2} \geq P_{i_1 j_1}$ and $D_{i_2 j_2} \geq D_{i_1 j_1}$. In other words, if players are free to choose and rational they should avoid playing Pareto-dominated strategies (strategies that make both of them worse off).

(D.36) Pareto perfection (informal)

'Players are not likely to choose strategies that are Pareto-dominated, irrespective of previous choices'.

The reason this is an informal definition is that the words "not likely" are used. This is not a very precise definition. Pareto perfection is sometimes called 'renegotiation proofness', since – irrespective of previous (non-binding) agreements – players are unlikely to stick to their threats (such as "Nash reversion strategy") if a Pareto superior option is available at a given stage.

Also the term 'Markov perfection' is defined in an informal way. Andrey Markov (1856-1922) was a Russian mathematician who is best known for the theory of stochastic processes (now called 'Markov processes'). The concept of a state variable has to be explained first. This is a term used in several mathematical theories. In statistics it means that the conditional probability distribution of future states of the process in question (conditional on both past and present states) depends only upon the present state, not on the sequence of events that preceded it.

Perhaps an example from physics is more informative. When an object is moving in a given direction, we may want to know its position at any time. In the case it is at rest, it is clear that its position is known at any time (for ever). Its "initial position" is the only state needed to describe it. However if the object is moving, knowing the position at present time is not sufficient to describe its future positions. Hence another state is required; speed can serve this purpose. But if the speed is not constant, the two states "position and speed" are not sufficient to describe the motion of the object, and a third state is required: namely "the acceleration". And so on.

While it may not be clear what state variables are needed in order to describe a system analysed, one thing is unquestionable; its future does not depend on how it evolved in the past.

(D.37) Markov Perfection (informal)

In a repeated game, payoffs depend only on a state variable (not on the history of specific strategies).

The following example describes a game which is Markov perfect, that is driven by the state variable (in this case the remaining resource), not by other considerations (like the past performance).

Example I (Extraction of a common resource)

Two players, 1 and 2, exploit the stock of a common resource $K^t \geq 0$. In each period $t=1,2,3,\dots$ they plan how much of the stock to extract a^t_1, a^t_2 . If $a^t_1 + a^t_2 < K^t$, then their actual extraction quantities are $s^t_i = a^t_i$ for $i=1,2$. If $a^t_1 + a^t_2 \geq K^t$ (the plans cannot be accomplished for physical

reasons) then $s_i^t = K^t/2$ for $i=1,2$. Please note that the actual extraction does not have to be equal to the planned one, and the resource is exhausted. Their payoffs are $\pi_i(s_i^t)$ for $i=1,2$. The stock of the resource is $K^{t+1} = f(K^t - s_1^t - s_2^t)$. If the stock is depleted completely (like when the plans exceed what is available) then it will never be created again: $f(0)=0$.

Markov-perfect strategies $(s_1^1, s_1^2), (s_2^1, s_2^2), (s_3^1, s_3^2), \dots$ depend only on K^1, K^2, K^3, \dots and not on the history of earlier decisions (of either player). Critics say that it would be difficult to find systems which comply with this assumption. Perhaps it is more realistic to assume that players retain some memory of how they behaved in the past, as observed by some students.

Exercise (E-11) included in the overheads to this lecture makes the example more specific by assuming a peculiar function f defining the evolution of the resource.

Another informal definition characterises "trembling hand perfection" which is the topic of the second part of the lecture. This is yet another attempt to distinguish between various Nash equilibria from the point of view of their realism understood as the likelihood of being encountered in real life situations.

D.38 Trembling-hand-perfect Nash equilibrium (THPNE) (almost formal definition)

THPNE is a Nash equilibrium that is robust to small perturbations. A strict mathematical definition is much more complex and it will not be quoted here. Instead an example will be analysed to explain which Nash equilibria are robust enough (to be considered THPNE), and which are not.

Example II

In the following game there are two Nash equilibria: (U,L) and (D,R).

		Second	
		L	R
First	U	(1,1)	(2,0)
	D	(0,2)	(2,2)

However, only (U,L) is THPNE, as the following argument explains.

Since (U,L) is a Nash equilibrium, the players are likely to choose pure rather than mixed strategies. In other words, the First is likely to play U and D with probabilities 1 and 0, respectively. Likewise the Second is likely to play L and R with probabilities 1 and 0, respectively. Neither has an incentive to unilaterally change the decision. Now let us assume that they make mistakes occasionally (their hands 'tremble'). I.e. let us assume that the First chooses U and D with probabilities $1-\varepsilon$ and ε , respectively (please note that $1-\varepsilon+\varepsilon=1$; if the hand does not 'tremble', $\varepsilon=0$). What is the optimum mixed strategy of the Second player?

If the Second plays L then the expected payoff will be $1(1-\varepsilon)+2\varepsilon=1+\varepsilon$. If the Second plays R then the expected payoff will be $0(1-\varepsilon)+2\varepsilon=2\varepsilon$. Thus, for small ε , the Second will maximize the expected payoff by choosing L with high probability (almost surely) and R with low probability (almost never).

Now let us assume that the hands of the Second player tremble, but the First one never fails to choose an appropriate strategy. Thus, let us assume that the Second player chooses L and R with probabilities $1-\varepsilon$ and ε , respectively, then the First will maximize his/her payoff by preferring U rather than D. Hence (U,L) turns out to be a Nash equilibrium robust with respect to such perturbations, i.e. a THPNE.

At last let us analyse the other Nash equilibrium, that is (D,R). Assuming that the First player chooses U and D with probabilities ε and $1-\varepsilon$, respectively, the Second will get the expected payoff of $1\varepsilon+2(1-\varepsilon)=2-\varepsilon$ (if playing L) or $0\varepsilon+2(1-\varepsilon)=2-2\varepsilon$ (if playing R). Therefore the Second will choose L more frequently. By the symmetry of payoffs, the First will choose more frequently U when the Second plays L and R with probabilities ε and $1-\varepsilon$, respectively (that is when the hands of the Second player tremble, but the First one never fails). Either calculation explains that (D,R) is not THPNE. All arguments can be made mathematically more precise, but this is not required if the concept of THPNE is only intuitively defined (as it is in this course).

The difference between the two Nash equilibria (U,L) and (D,R) is that (U,L) makes worse off only the player who made the mistake. The other player is actually made better off by this mistake. In contrast, (D,R) makes worse the other player drastically. Thus intuitively, THPNE can be understood as a Nash equilibrium that is not a very risky one.

If the same or a similar game is played repeatedly, then several questions can be raised. The first one is about what the players remember. If they do not remember their past performance, then Nash equilibria are 'Markov perfect'; they depend only on the present payoffs rather than strategies played in the past. If they do not remember earlier threats, Nash equilibria are 'Pareto perfect'. 'Pareto perfection' is also called 'renegotiation proofness' understood as avoiding Pareto dominated strategies. If players agreed to use certain strategies (like, for instance, 'Nash reversion'), but it turns out that these strategies can be substituted by more attractive ones in a given round, they will do so. The second question is about stability of Nash equilibria. If players never fail to use their Nash strategies, they are likely to stick to their earlier choices. But if one of them has a trembling hand, that is, tries to use his (or her) Nash strategy, but makes occasional mistakes, then the question is how the other player is likely to react. If he (or she) expects to be better off by moving away from the earlier strategy, the Nash equilibrium was not trembling-hand-perfect. Otherwise it was.

Questions and answers to lecture 11

11.1 Can there be games with a Nash equilibrium achieved in a Pareto optimum?

Yes. Please recall the game (called coordination game)

	Second		
		L	R
	First		
	U	(1,1)	(-1,-1)
	D	(-1,-1)	(1,1)

It has two Nash equilibria (1,1) which happen to be Pareto optima at the same time.

11.2 Domination of strategies was defined in the fourth lecture (D.15 and D.16). Now Pareto-domination is defined (D.35). What is the difference between the two?

Domination of strategies referred to inequalities satisfied for one player only. Pareto domination refers to two inequalities: for payoffs of the first player, and the payoffs of the second player. If two strategies are dominated (for either of the players) then they are also Pareto dominated (this concept refers to both players).

11.3 Do players of the prisoner's dilemma game (GT-4, page 23) have Pareto dominated strategies?

Yes. D is a dominated (by C) strategy for either of the players. At the same time:

$$P(D,D)=-12 < -1 = P(C,C), \text{ and } D(D,D)=-12 < -1 = D(C,C).$$

This means that D is Pareto dominated by C. The same conclusion can be drawn from the answer to question 11.2.

11.4 Please argue that Nash reversion strategy is not necessarily Pareto-perfect.

Let us refer to the Bertrand-Nash reversion strategy as defined in the previous lecture (GT-10, D.33). If one of the players defected, the other one should offer the price $p=c$ (this is what he or she promised). But by doing so, the payoffs (of both players) will be wiped out for sure. However, if the price offered is higher than c , then there is a possibility of having a positive profit. Thus Nash reversion should be renegotiated if players behave rationally.

11.5 Is it realistic to expect that players of the common resource extraction game apply Markov perfection?

Not really. Perhaps under very peculiar conditions players can look at the remaining stock of the resource (state variable) and take decisions without any reference to how they behaved in the past. A more realistic assumption is that they choose strategies taking into account their reputation. If, for instance my rival has a reputation of a non-reliable person, I would prefer to gain as much as possible now rather than planning for the future. On the contrary, if my rival has a reputation of an honest and responsible person, I would not hesitate to sacrifice the short run for the sake of a more distant future.

11.6 Is the Nash equilibrium of the prisoner's dilemma game (C,C) trembling hand perfect?

The prisoner's dilemma analysed in the class has the following payoff matrix:

		Second	
		C	D
First	C	$(-12,-12)$	$(0,-18)$
	D	$(-18,0)$	$(-1,-1)$

If the First player has a trembling hand (plays C with the probability of $1-\varepsilon$ and D with the probability of ε), then the payoffs of the second player are $-12(1-\varepsilon)+0(1-\varepsilon)=-12+12\varepsilon$ if C is played or $-18(1-\varepsilon)-1\varepsilon=-18+17\varepsilon$ if D is played. For small ε C is better than D. From the

symmetry of the payoff matrix, the same conclusion is obtained if the Second player has a trembling hand. Hence (C,C) is THPNE.

11.7 Are the Nash equilibria (U,L) and (D,R) of the following game

	Second		
		L	R
	First		
	U	(5,1)	(0,0)
	D	(4,4)	(1,5)

trembling hand perfect?

Let us assume that the First player has a trembling hand. Thus he (or she) plays U with probability $1-\varepsilon$, and D with probability ε . The payoff enjoyed by the Second player is $1(1-\varepsilon)+4\varepsilon=1+3\varepsilon$ when playing L or $0(1-\varepsilon)+5\varepsilon=5\varepsilon$ when playing R. For small ε L is better than R. From the symmetry of the payoff matrix, the same conclusion is obtained if the Second player has a trembling hand. Hence (U,L) is THPNE. Once again by the symmetry of the payoff matrix (symmetry of positions of both players), (D,R) is also THPNE.

12 – Coalitions I

This lecture addresses the problem of coalitions – one of the most interesting, but also the most complicated topics in game theory. In the case of two-person games coalitions are trivial; they consist of a single player or of both of them. The two players cannot form a coalition against anybody else, because nobody else is there. In the case of three-person games, they start to be essential. Three different coalitions can be established against any single player; any two can collude in order to harm the third. But the player excluded from the coalition does not have to stay idle. He (or she) can try to break the existing coalition and to establish a new one against somebody. Examples developed in this lecture will explain how coalitions can be formed.

(D.39)

Let there be I players $\{1, \dots, I\}$ as in D.20. A coalition $K \subset \{1, \dots, I\}$ improves upon, or blocks, the outcome (s_1^*, \dots, s_I^*) if there is a set of strategies (s_1, \dots, s_I) such that for every $i \in K$ $u_i(s_1, \dots, s_I) > u_i(s_1^*, \dots, s_I^*)$.

Please note that this definition does not say anything about the feasibility of a coalition. In particular, there can be one-player "coalitions" that can be hardly effective – especially if the rest of the players make a powerful group. The concept of coalitions does not refer to any majority voting. This somewhat strange notion is used in order to define the concept of a "core" – another key idea in game theory.

(D.40)

The outcome (s_1^*, \dots, s_I^*) has the core property, if there is no coalition that can improve upon it. The set of outcomes that have the core property is called the core.

Example I

The 'Prisoner's Dilemma' game given by the following payoff matrix

		Second	
		C	D
First	C	(-12,-12)	(0,-18)
	D	(-18,0)	(-1,-1)

has an empty core. Namely, it can be seen that for every outcome (cell in the matrix) at least one player can make a coalition (with himself or herself) to improve his/her payoff

Once again, let me stress, that the concept of coalitions has nothing to do with how a theoretically available "coalition" can be implemented.

Definitions D.39 and D.40 can be easily extended to games with infinite number of players, and the next definitions do not depend on the finite number of players; wherever there is a summation over the set of players, integration can be applied (in the case of a continuum number of players). Nevertheless – in order to keep things simpler – in our definitions there is always summation rather than integration. Besides, all examples refer to games with a finite number of players.

(D.41)

A coalitional game with transferable payoffs is defined as $\langle N, v \rangle$, where N is the finite set of players, and v is a real function defined over the set of all possible coalitions $K \subset N$. $v(K)$ is called the worth of the coalition. $v(N)$ is the worth of the game. If x_1, \dots, x_n are payoffs then $x(K) = \sum_{i \in K} x_i$ is called a payoff profile. A payoff profile is feasible if $x(K) = v(K)$. By convention, $v(\emptyset) = 0$.

The worth of the game $v(N)$ is an abstract number. Intuitively it can be understood as how much can be obtained from the game, irrespective of how it is distributed among its participants ("transferability" of payoffs allows different distributions). The worth of the coalition $v(K)$ is a fraction of the worth of the game $v(N)$. A payoff profile $x(K)$ is the payoff enjoyed by the coalition participants. It is called feasible, if it is equal to the worth of the coalition. If it is higher than the worth of the coalition ($x(K) > v(K)$) then coalition participants pay themselves more than the coalition can achieve; this is not possible. If it is lower than the worth of the coalition ($x(K) < v(K)$) then coalition participants pay themselves less than the coalition can achieve; this is possible, but irrational – some revenues are left unused.

(D.42)

Shapley value of a coalitional game with transferable payoffs $\langle N, v \rangle$ for the player i is:

$$Sh_i(N, v) = (1/(card N)!) \sum m(K(\pi, i), i)$$

where the summation extends over all permutations of the set N . The number $m(K(\pi, i), i)$ is the increase of the sum of payoffs the player i 'brings' to a coalition $K(\pi, i)$ consisting of those members of the set N who in the permutation π precede i .

Note


Elements of Shapley value can be interpreted as the threat of leaving made by the player i to a coalition he/she is a member of (K) consisting of players whose place precedes i :

$$(v(K)-v(K\setminus\{i\})).$$

The sum of these numbers extends over all possible orderings (permutations) of the members of N . Thus $Sh_i(N,v)$ can be also interpreted as an expected value the player i is 'entitled to' if other players behave rationally.

Before we proceed with further examples, let me start with perhaps the simplest case.

Example II (left and right hand gloves)



Let players 1 and 2 own one left-hand glove each and player 3 own one right-hand glove (matching either of the left-hand gloves). The payoff from matching a pair is 12, while the payoff from an unmatched pair (or from a single glove) is 0. Calculate the Shapley values of the game where players try to get something useful out of what they have.

If the first player makes a coalition with the third one, or if the second makes it with the third, they will get a useful pair of gloves. However, if the first makes a coalition with the second, they will not get anything useful. You can see that the third cannot be left out from any reasonable coalition. Likewise, you can also feel that the first and the second can be substituted for each other; the third one can refer to it when he (or she) negotiates with them about whom to cooperate with. The first and the second are probably aware of this and they cannot expect too much if they are invited by the third to form a coalition. The definition D.42 takes this into account in a systematic way.

First of all, we have to list all the possible coalitions. To do this we have to consider all possible orderings of the numbers 1, 2, and 3 (members of the set N of players). There are six such orderings: $\{1,2,3\}$, $\{1,3,2\}$, $\{2,1,3\}$, $\{2,3,1\}$, $\{3,1,2\}$, and $\{3,2,1\}$.

Ordering	1	2	3
$\{1,2,3\}$	0	0	12
$\{1,3,2\}$	0	0	12
$\{2,1,3\}$	0	0	12
$\{2,3,1\}$	0	0	12
$\{3,1,2\}$	12	0	0
$\{3,2,1\}$	0	12	0
Σ	12	12	48

Let us take the first player (1) and think of the coalition formed by the players who precede the player number 1 in $\{1,2,3\}$. There is nobody to the left, therefore the coalition is empty. It produces nothing, and if the first player joins such an empty coalition they will not produce anything useful either; the increase in value of such a coalition is zero. We put 0 in this row.

Now let us take the next ordering, i.e. $\{1,3,2\}$ and think of the coalition formed by the players who precede the player number 1 there. Again, there is nobody to the left, therefore the coalition is empty, and if the first player joins such a coalition they will not produce anything useful either; the increase in the value of such a coalition is zero. We put 0 in this row again. In the ordering $\{2,1,3\}$ there is player number 2 who precedes him (or her) and who produces nothing. If another left glove is added, they will not get anything useful out of it either. Thus the increase in the value of such a coalition is zero, and we put 0 in this row again. In the ordering $\{2,3,1\}$ there are players number 2 and 3 who precede him (or her) and who can sell a useful pair of gloves. The value of it is 12. If the additional left hand glove is available, then nothing more can be sold. Hence the increase in the value of such a coalition is zero, and we put 0 in this row again. Things change when we take the next ordering, that is $\{3,1,2\}$. Now there is the player number 3 who precedes him (or her) and who has a right hand glove. The left hand glove brought by the player number 1 makes a useful addition, and the increase in the value of such a coalition is 12. Therefore we put 12 in this row. Finally, we look at the ordering $\{3,2,1\}$. There are players number 3 and 2 who precede him (or her), and the situation is like in the case of ordering $\{2,3,1\}$; the additional player (number 1) does not add anything, the increase in the value of such a coalition is zero, and we put 0 in this row again.

By looking at increases in the value of coalitions when the player number 1 joins, we filled out the second column of the table above. Likewise the next column can be filled out for the player number 2. The only difference is that this player increases the value of coalition in the row $\{3,2,1\}$ only (the player number 3 with a right hand glove cannot sell anything unless somebody with a left hand glove joins).

There will be more non-zero entries in the last column (for the player number 3). In the first ordering $\{1,2,3\}$ players 1 and 2 precede him (or her); but they have two left gloves only, and therefore they produce nothing. If the player number 3 joins, they can sell a pair of gloves, and the value of the coalition increases by 12. In the second ordering $\{1,3,2\}$ player 1 precedes him (or her), but with the left hand glove only he (or she) cannot sell anything. The value of coalition increases by 12 when the player 3 with a right hand glove joins. The ordering $\{2,1,3\}$ is like $\{1,2,3\}$, and the ordering $\{2,3,1\}$ is like $\{1,3,2\}$. In both cases, the value of coalitions increases by 12. In orderings $\{3,1,2\}$ and $\{3,2,1\}$ there is nobody to the left, so joining an empty coalition with a right hand glove does not produce anything useful. Both entries are therefore 0.

The sums in the lowest row are equal to what is denoted by $\sum m(K(\pi,i),i)$ in the Shapley formula (D.42). These numbers have to be divided into $(\text{card}N)!$ in order to calculate Shapley values. The number of players is 3, and $3!=6$. Thus the Shapley values for this game are $Sh_1=12/6=2$, $Sh_2=12/6=2$, and $Sh_3=48/6=8$. They sum up to 12, that is to the value of one pair of gloves, i.e. to what the game offers to its players.

The formula D.42 looks complicated, but the idea of the Shapley value is intuitively simple: it reflects what players can bring by joining a coalition. Something that one needs to be careful about is to check all the possible coalitions. In order not to lose anything, the formula is based on permutations (orderings of the players). But several orderings can correspond to the same coalition. For example 1,2 and 2,1 mean the same coalition of the same two players, but they appear separately. In order not to double count, there is this requirement to look at those who precede, that is are listed to the left.

From the very beginning it was clear that Shapley values of players 1 and 2 (bringing the same things) must be lower than the Shapley value of player 3 (bringing a unique thing). But the numbers we calculated look strange; the Shapley value of both players 1 and 2 combined (4) is lower than the Shapley value of the third one (8). Nevertheless in the next lecture we will see that it is the only one satisfying certain reasonable criteria (T.31). Apparently the game theory is not about being nice to each other, but about trying to get as much as possible from what you have, and others do not have. The next example provides another argument.

Example III (capitalist production)

A capitalist owns a factory and each of w workers owns his/her labour only. Without the capitalist workers cannot produce anything. Any group of m workers ($m > 0$) can produce $f(m)$, where $f: \mathcal{R} \rightarrow \mathcal{R}$ is a linear function $f(m) = mp$ if they are employed by the capitalist (note that $f(0) = 0$) with $p > 0$ interpreted as the marginal and average productivity of a worker (the same for any worker).

This can be modelled as the following coalitional game with transferable payoffs $\langle N, v \rangle$. $N = \{c\} \cup W$, where c denotes the capitalist and W – the set of workers. The game is worth

- $v(K) = 0$ if $c \notin K$, and
- $v(K) = f(\text{card}(K \cap W))$ if $c \in K$.

The second bullet states that the workers who produce – there are m such workers – produce $f(m) = mp$; and this is $v(K)$.

The core of this game is $\{x \in \mathcal{R}^{\text{card}N}: 0 \leq x_i \leq p \text{ for } i \in W \text{ and } \sum_{i \in N} x_i = wp\}$, where w is the number of workers participating (please note that in the summation $\sum_{i \in N} x_i = wp$ there is the capitalist also – not only the workers). Its Shapley values are

- $Sh_c(N, v) = wp/2$, and
- $Sh_i(N, v) = p/2$ for $i \in W$.

Note that the core includes $[wp/2, p/2, \dots, p/2]$, but it also includes the 'competitive' outcome $[0, p, \dots, p]$ (the capitalist gets nothing, while the workers get their marginal productivity).

To see that the core is $\{x \in \mathcal{R}^{\text{card}N}: 0 \leq x_i \leq p \text{ for } i \in W \text{ and } \sum_{i \in N} x_i = wp\}$, let us assume that one worker does not like the result and wants to establish another coalition to get a higher payoff. But since $\sum_{i \in N} x_i = wp$, somebody else must have a lower payoff, which contradicts the definition of the core.

Example IV (A three-player majority game)

Three players can obtain a payoff of 1; any two of them can obtain jointly $\alpha \in [0, 1]$ irrespective of the actions of the third; each player alone can obtain nothing irrespective of the actions of the remaining two. The parameter α is fixed for the game (it cannot be changed by the players). The only thing the players can manipulate with is how they split α between them.

This is a coalitional game with transferable payoffs $\langle N, v \rangle$ where $N = \{1, 2, 3\}$, $v(N) = 1$, $v(K) = \alpha$ whenever $\text{card } K = 2$, and $v(\{i\}) = 0$ for any $i = 1, 2, 3$. The core of this game is the set of all

nonnegative payoff vectors $[x_1, x_2, x_3]$ such that $x(N)=1$ and $x(K)=\alpha$ for every two-player coalition (without a third player). Therefore if $\alpha > 2/3$ the core is empty.

To see this, let us assume that all three players are in coalition ($K=N$). Then they can enjoy equal payoffs of $1/3$ each. Any two of them can make a coalition to get jointly α . In other words, if $\alpha = 2/3$, the allocation $[1/3, 1/3, 1/3]$ is equally good for them. Yet if $\alpha > 2/3$, then no allocation can be sustained, because if one player enjoys $x_i \geq 1/3$, then the other two get less than $2/3$ (or exactly $2/3$), and they can form a two-player coalition to get more than $2/3$ (when $\alpha > 2/3$).

Ordering	1	2	3
{1,2,3}	0	α	$1-\alpha$
{1,3,2}	0	$1-\alpha$	α
{2,1,3}	α	0	$1-\alpha$
{2,3,1}	$1-\alpha$	0	α
{3,1,2}	α	$1-\alpha$	0
{3,2,1}	$1-\alpha$	α	0
Σ	2	2	2

The Shapley values for all the players are equal to $1/3$ as the table above explains. By dividing 2 (the numbers in the bottom row) into $3!=6$ one gets the result. Without using D.42, by the mere symmetry, it can be argued that the Shapley values must be the same, and since they add up to 1, they must be $1/3$ each.

Example V (Miners game)

$I > 2$ miners discovered $[I/2]$ heavy identical pieces of a valuable mineral ($[x]$ – called *entier*, integer part – is the largest number not higher than x). Each piece can be carried out (and privately sold for the price of 1, e.g. $1/2$ for each carrying miner) by two miners. Will they agree on how to handle this discovery?

The problem can be modelled as a coalitional game with transferable payoffs $\langle N, v \rangle$, where

- $v(K) = (\text{card } K)/2$ if card K is even; and
- $v(K) = ((\text{card } K) - 1)/2$ if card K is odd.

If $I \geq 4$ and even then the core consists of a single payoff profile $[1/2, \dots, 1/2]$ which is fairly obvious. However, if $I \geq 3$ and odd then the core is empty. To see this, let us assume that there are 3 miners and one piece. Any two of them can carry out the piece, sell it, and share the revenue of 1 somehow. At least one of them enjoys the payoff of $1/2$ or less (if they split the revenue uniformly, then each enjoys the payoff of $1/2$). The third miner is left with 0. He (or she) can offer to make a coalition with the guy who has not more than $1/2$, and offer the share of, say, $3/4$ or something more attractive than the present payoff. In other words, no solution can be sustained.

By the way, one can calculate that the Shapley values in this game are the same for all the miners. They are equal to p/I , where p is the number of pieces, and I is the number of miners. As the number of miners is unknown, it would be difficult to apply the D.42 formula directly. By the symmetry, however, no miner should have a higher value than others.

The examples above illustrate the fact that analysing coalitions can be difficult, and it can yield strange results (like in the miners game, where the core depends on whether the number of miners was odd or even). Also the definition of the Shapley value is somewhat strange. It is intuitively obvious that different players may have different negotiation advantages. Nevertheless specific examples demonstrate that this asymmetry is sometimes unexpected.

Questions and answers to lecture 12

12.1 Please prove that for the following coordination game (example I in GT-5)

		Second	
		L	R
First	U	(1,1)	(-1,-1)
	D	(-1,-1)	(1,1)

the two Nash equilibria make its core.

Indeed, if the payoffs are 1 for each of the players, nobody can improve the outcome.

12.2 Please prove that the core of the following game (GT-5, page 27) is empty

		M	
		Y	N
F	Y	(4,2)	(-2,3)
	N	(2,1)	(-1,0)

Indeed, for any outcome (cell) there is at least one player who is not satisfied with the payoff.

12.3 Please prove that the core of the following game (GT-5, example II) is empty

		Second	
		L	R
First	U	(5,1)	(0,0)
	D	(4,4)	(1,5)

Indeed, for any outcome (cell) there is at least one player who is not satisfied with the payoff.

12.4 Please prove that even if the game has a Nash equilibrium which is a Pareto optimum, it may have an empty core.

For example the game

		Second	
		L	R
First	U	(3,6)	(0,0)
	D	(4,4)	(6,3)

has the unique Nash equilibrium (4,4). It happens to be a Pareto optimum. Yet the core of the game is empty, since for any outcome (cell) there is at least one player who could get a higher payoff elsewhere.

12.5 In the 'glove game' (example II in this lecture) the payoff profile can be [2,3,7] (the first player is paid 2, the second – 3, and the third – 7). The first player threatens to leave the coalition unless he/she is paid 3. Is his threat effective?

No. If he/she leaves, the second and the third player can establish a coalition where the second is paid 3 or more, and the third is paid 9 or less. This makes them better off, so they will not press the first player to stay. However, if the third player threatens to leave unless he (or she) receives more than 7, it is a different story. If the first and the second establish a coalition without him (or her), they cannot produce anything. Hence his/her threat can be effective. In lecture 14 (GT-14) we will analyse how players are likely to behave when they play coalitional games with transferable payoffs.

12.6 In the example III (capitalist production), if there are three workers and a capitalist, the workers produce $3p$ (p – average and marginal productivity of a worker), the capitalist gets $3p/2$, and workers get $p/2$ each. If the capitalist offers one of the workers to be paid better if he/she works with the second worker, but without the third one, is this credible?

No, it is not, because if the capitalist is left with two workers only, the production will be $f(2)=2p$. If the capitalist is to get not less than before, that is $3p/2$, then only $p/2$ is to be shared between the two workers left. But this leaves them with lower payoffs than before.

12.7 This is a slightly modified version of E-12, an exercise accompanying the overheads to this lecture. Let us assume that three-players majority game is played with $\alpha=3/4$ (example IV above). The payoff profile is not symmetric: $[1/4, 1/4, 1/2]$. One of the players who gets $1/4$ approaches the other one who gets $1/4$ suggesting that once they get rid of the third player (who gets $1/2$) they will get $\alpha=3/4$ to be split into halves ($3/8$ each), so that both of them can be better off. Does this offer make sense?

It may result in some changes, but players are likely to seek some other solutions. Faced with the risk of being excluded from the coalition, the third player is likely to offer to reduce his/her payoff, perhaps to the symmetric profile of $[1/3, 1/3, 1/3]$, but the first player may then suggest to get rid of the second one, and to get $\alpha=3/4$ to be split into halves ($3/8$ each). This will improve the hypothetical situation of the first and the third at the expense of the second one. No matter how they manipulate with payoffs, there will always be somebody who can be made better off by suggesting an alternative coalition. Hence as long as $\alpha=3/4$, the core of the game is empty.

12.8 In the example V (miners game) let us assume that there are 5 miners and three (not two) heavy pieces discovered. Why is the core of the game empty?

The number of miners is odd. Let us assume that miners 1 and 2 agreed to carry out the first piece, miners 3 and 4 agreed to carry out the second piece, and miners 1 and 5 agreed to carry out the third one. If they share the revenues from what they carry uniformly, then the miner 1 gets 2, and miners 2, 3, 4, and 5 get 1 each. Then any of the miners number 3, or 4 can

approach the miner number 2 or 5 to offer them to carry out the first piece or the third one for just 1/3 of the share (unlike 1/2 to be claimed by the miner number 1). By making this offer they are likely to motivate the miner number 1 to offer something alternative, but no matter what they agree to do, there will always be at least one miner willing to change the agreement. This corresponds to the emptiness of the core.


13 – Coalitions II

By the end of this lecture we will go back to the general question of coalitions. Yet most of the time will be devoted to the most important economic example of coalitional games with transferable payoffs, namely to the market economy. The latter can be considered a game between economic agents who analyse certain prices in order to build supply and demand schedules. The demand and supply are confronted with each other and market prices are formed.

In some analyses we make a simplifying assumption that economic agents do not produce anything; they merely exchange what is available. In more complicated analyses we relax the assumption on the lack of production, and we admit that what is available originates from inventories but it can be produced too.

A so-called pure exchange economy can be considered a game. In what follows we will confine to the case where two consumers hold two goods and contemplate whether to exchange some units of one good for some units of the other one. The players consider only allocations which do not make them worse off. The following terms will be used (the first subscript denotes the number of the consumer, and the second – the number of a good; this notation allows an arbitrary number of consumers):



- Endowment (initial allocation) of the i th consumer: ω_{i1}, ω_{i2}
- Total endowment of the j th commodity: $\omega_j = \omega_{1j} + \omega_{2j}$
- Gross demand (final allocation) of the i th consumer: x_{i1}, x_{i2}
- Total demand for the j th commodity: $x_j = x_{1j} + x_{2j}$
- Excess demand of the i th consumer: $x_{i1} - \omega_{i1}, x_{i2} - \omega_{i2}$



In order to get insights into relationships between the supply, the demand and the prices, so-called Edgeworth box will be applied. This name refers to an English economist Francis

Ysidro Edgeworth (1845-1926), who first applied this technique. This is a graphical analysis of feasible allocations in a pure exchange economy (superposition of two coordinate systems for the analysis of a consumer's choice: the width of the rectangle = $\omega_{11} + \omega_{21}$, the height of the rectangle = $\omega_{12} + \omega_{22}$; the second system is rotated by 180°).

Please imagine the standard consumer choice model recalled in the first lecture (GT-1) – the left part of the picture above – with its image rotated by 180° – something that you see in the right part of the picture above. The good number 1 is measured horizontally, and the good number 2 is measured vertically. An additional detail you need to decide is how close the coordinates should be. Edgeworth suggested that the "box" they make should be the $\omega_{11} + \omega_{21}$ wide and the $\omega_{12} + \omega_{22}$ high. If we follow these suggestions, we can arrive at pictures like below (the shaded area in the upper picture indicates combinations of goods preferred by both consumers to what they hold initially).



Pictures refer to abstract goods number 1 and number 2. But let us be specific. The good number 1 is apples, and the good number 2 is oranges. The first person (A) brought 8 apples and 2 oranges, and the second person (B, whose axes were rotated by 180°) brought 2 apples and 3 oranges. Thus they have 10 apples and 5 oranges jointly. They can stick to what they brought. However, they can also contemplate whether it would make sense to exchange some of them. The decisions may depend on relative prices of the goods they have. For example, person A can think: "if one apple goes for two oranges, I could exchange 2 of my apples for 4 oranges, but if two apples go for one orange I could also improve the utility from what I have". The person B can think: "if one apple goes for two oranges, I need to sacrifice two oranges to get an additional apple; this is not attractive, but – for instance – if one apple goes for one orange, this would change my calculation". Having considered alternative relative prices, A and B may agree to exchange some of the goods they have in order to improve their utilities.

Edgeworth's trick allows us to have a correspondence of points in a plane with four numbers rather than two (as we usually have). Usually a point with coordinates (8,2) (8 apples and 2 oranges) corresponds to two numbers: 8 and 2. But in the Edgeworth box, like above, it corresponds to four numbers: 8 apples and 2 oranges owned by person A, and 2 apples and 3 oranges owned by the person B. We will see later, that in fact, it corresponds to five numbers. In addition to the numbers of apples and oranges owned by A and B, it shows the equilibrium price ratio, equal to the slope coefficient of the tangent line to both indifference curves in X^* .

Now let us go back to the abstract terminology. Pictures illustrate the following situation. Indifference curves (i.e. the sets of points yielding the same utility) of the agent A, $I_A(\alpha)$ are given by the formula $x_{2A}=\alpha/x_{1A}$, while indifference curves of the agent B, $I_B(\beta)$ are given by $x_{2B}=\beta/x_{1B}$ ($\alpha, \beta > 0$ – parameters); additionally, we assume that the total quantity of the first good is 10, while that of the second – 5. Moreover, the diagram corresponds to the initial allocation of the first good 8:2, while of the second one – 2:3 (point X_0). There are two indifference curves containing this point: $x_{2A}=16/x_{1A}$ ($\alpha=16$) and $x_{2B}=6/x_{1B}$ ($\beta=6$). A would prefer to be on a higher indifference curve, say, in (9,3), i.e. on the curve $x_{2A}=27/x_{1A}$ ($\alpha=27$). At the same time, B would like to have more of everything, i.e. to be in, say, (3,4), i.e. on the curve $x_{2B}=12/x_{1B}$ ($\beta=12$). It is impossible to satisfy these expectations at the same time. A solution which can place both agents in a jointly preferred point (one should solve a system of simultaneous equations) is: $x_{1A}^*=6$, $x_{2A}^*=3$, $x_{1B}^*=4$, $x_{2B}^*=2$, $\alpha=18$, $\beta=8$, $p=p_1/p_2=0.5$ (see the second picture). Agents A and B are on $I_A(18)$ and $I_B(8)$, respectively, and they are better off than in X_0 . One can see from the figure that they cannot improve their situations further simultaneously. In other words, (6,3) is a Pareto optimum (a combination such that it is impossible to improve the situation of one person without deteriorating the situation of the other person). Equilibrium prices which satisfy this solution are multiple, e.g. $p_1=1$, $p_2=2$, or $p_1=7$, $p_2=14$, or $p_1=0.5$, $p_2=1$ etc, as long as $p_1/p_2=p=0.5$.

Let us calculate the answer ($\mathbf{x}=[6,3,4,2]$ and $p_1=p_2/2$) algebraically. In other words, we need to find x_{1A} , x_{2A} , x_{1B} , x_{2B} , α , β , and $p=p_1/p_2$. At the first glance it seems as if we had 7 unknowns. Nevertheless this number can easily be reduced to 5, since $x_{1B}=10-x_{1A}$ and $x_{2B}=5-x_{2A}$. Thus let $p=p_1/p_2$, $x_1=x_{1A}$ and $x_2=x_{2A}$. By the definition of the indifference curves:

$$\alpha=x_1x_2, \text{ and } \beta=(10-x_1)(5-x_2).$$

In the market equilibrium both indifference curves are tangent to the price ratio, i.e.

$$-\alpha/(x_1)^2=-p, \text{ and } -\beta/(10-x_1)^2=-p$$

(one needed to calculate the derivatives with respect to x_1 of the following two functions:

$$x_2 = \alpha/x_1, \text{ and } x_2 = 5 - \beta/(10 - x_1))$$

which allows to calculate α and β : $\alpha = p(x_1)^2$ and $\beta = p(10 - x_1)^2$. Thus from the definitions of the indifference curves: $p(x_1)^2 = x_1 x_2$ and $p(10 - x_1)^2 = (10 - x_1)(5 - x_2)$; and after cancellations we get: $p x_1 = x_2$ and $p(10 - x_1) = (5 - x_2)$. These are 2 equations with 3 unknowns. We need an additional equation. By the necessity of equating expenditures with revenues (please note that the same equation is derived irrespective of whether the balance reflects the consumer A or B) we get $p_1(4 - x_1) = p_2(x_2 - 4)$, i.e. $p_1/p_2 = (4 - x_2)/(x_1 - 4)$ (with $x_1 \neq 4$). Hence the third equation reads

$$p = (4 - x_2)/(x_1 - 4).$$

By substituting p into the previous equations we get:


$$x_1(4 - x_2) = x_2(x_1 - 4) \text{ and } (10 - x_1)(4 - x_2) = (5 - x_2)(x_1 - 4).$$

These are two equations with two unknowns. After cancellations we obtain: $4x_1 - 2x_1x_2 = -4x_2$ and $60 - 9x_1 + 2x_1x_2 = 14x_2$. By adding the two we get $x_2 = 6 - x_1/2$ which can be substituted into the first one, yielding

$$(x_1)^2 - 10x_1 + 24 = 0.$$

The last quadratic equation has two roots: $x_1 = 6$ or $x_1 = 4$ implying $x_2 = 3$ or $x_2 = 4$, respectively. The second solution has to be rejected (in order to calculate p properly; i.e. – as noted above – to avoid dividing into zero).

Hence the answer that could have been seen in the picture, can be also calculated algebraically. Both persons are better off if they agree to the price ratio $p = p_1/p_2 = 1/2$ (one apple goes for two oranges), A sells two apples, and buys one orange, and B sells one orange and buys two apples. As a result of this exchange, A has 6 apples, and 3 oranges, while B has 4 apples, and 2 oranges. Both of them are now on higher indifference curves.



An additional definition has to be introduced. A contract curve is the set of all allocations which are Pareto optima and which are preferred by both players over their initial allocation. It is depicted as a thick line in the following picture. A contract curve in an Edgeworth box is a core in a pure exchange economy game. To see that it has the core property, we have to demonstrate that if players are on the contract curve, they cannot improve their situations. In fact they can if they start from X_0 (the first can move to the right to be at a higher indifference

curve, while the second is still in the shaded area, i.e. at a higher indifference curve than in X_0). But if they start from a point which is on the contract curve already (where indifference curves are tangent to each other), they will not agree to make a transaction, because moving one to a better indifference curve requires the second to be on a worse indifference curve.

What you see on the picture above can be defined formally as a Walrasian (competitive) equilibrium. The concept originates from Marie-Esprit-Léon Walras (1834-1910) a French economist, who worked in the University of Lausanne. The adjective 'competitive' refers to the fact that the market equilibrium we talk about is competitive in a sense that buyers and sellers cannot manipulate prices (they take prices as given). Its formal definition follows.

(D.43)

X^* is a Walrasian (competitive) equilibrium in an Edgeworth box pure exchange economy if:

- $u_A(x_{1A}^*, x_{2A}^*) \geq u_A(x_{1A}, x_{2A})$ for all $(x_{1A}, x_{2A}) \in B_p(X_0)$,
- $u_B(x_{1B}^*, x_{2B}^*) \geq u_B(x_{1B}, x_{2B})$ for all $(x_{1B}, x_{2B}) \in B_p(X_0)$,
- $B_p(X_0) = \{(x_{1A}, x_{2A}, x_{1B}, x_{2B}) \in \mathbb{R}^4 : p_1 x_{1A} + p_2 x_{2A} \leq p_1 \omega_{1A} + p_2 \omega_{2A}, \text{ and } p_1 x_{1B} + p_2 x_{2B} \leq p_1 \omega_{1B} + p_2 \omega_{2B}\}$

$B_p(X_0)$ is the combined budget set of consumers A and B, that is the set of bundles they can buy given the prices $p=[p_1, p_2]$. Please note that $p_1 \omega_{1A} + p_2 \omega_{2A}$ is the amount of money that A has if he/she decides to sell all his/her initial allocation. Likewise $p_1 \omega_{1B} + p_2 \omega_{2B}$ is the amount of money that B has if he/she decides to sell all his/her initial allocation. Example below explains what in the Walras model is understood by money implicit in the initial allocation of goods.

An old lady without any revenues lives alone in a large villa worth 1 million € and starves. If she sold the villa (or rented it to somebody), she could meet all her needs like food, medication and a satisfactory housing. When asked about her preferences, she says that the villa is not for sale. Is the lady rich or poor? The fact is that she has no cash (no money income, as observed by one of the students). But on the other hand, her wealth is quite impressive; it is 1 million € not in cash, but in real estate. From the Walras model point of view this is irrelevant. Her wealth allows her to buy whatever we think that she needs for decent living. It is her preference that she wants to use it for living in the villa (and nothing else). If she sold it, she could have enough money to buy other things. But she does not want.

The Walras model assumes that whatever we have has some implicit value. The value depends on market prices. We may prefer to keep it rather than sell and use the money to buy other goods, but this is our sovereign decision. In other words, "everything is for sale", even though sometimes we prefer to stick to what we have instead of using it to exchange for something else.

There are two fundamental theorems in so-called welfare economics: the first and the second. They establish a relationship between Pareto optima (PO) and Walras equilibria (WE). The first corresponds to the implication $WE \Rightarrow PO$, and the second: $PO \Rightarrow WE$. What they establish is almost an equivalence. Strictly speaking this is not an equivalence, because these theorems require different assumptions. As a rule, the first is easy, and it does not require almost any assumptions. The second is more difficult, and it requires some mathematical assumptions, like the convexity of certain functions. Welfare economics theorems are stated and proved in

many areas of economics. Here we will confine to the first one phrased in the language of a pure exchange economy.

(T.30) The first fundamental welfare economics theorem

In an Edgeworth box pure exchange economy, if X^* is a Walrasian equilibrium then X^* is in the core of a corresponding coalitional game with transferable payments (it is a Pareto optimum)

Proof:

It is sufficient to demonstrate that no coalition can improve upon X^* . If exchanges were to be carried out at given prices this would have been obvious given both inequalities (no improvements are possible) and the feasibility condition (each consumer can spend no more than he/she obtains from selling his/her endowments). But the core property is more general. It refers to potential coalitions that do not necessarily apply price mechanisms.

Thus let us assume that there is another feasible allocation X' which is at least as good as X^* for one consumer and strictly better for the other one. Then

$$p_1(x'_{1A} + x'_{1B}) + p_2(x'_{2A} + x'_{2B}) > p_1(x^*_{1A} + x^*_{1B}) + p_2(x^*_{2A} + x^*_{2B}),$$

which means that X' must have been too expensive to be chosen in equilibrium (in equilibrium X^* was chosen instead; X' must have been unaffordable). However, in a pure exchange economy: $x'_{1A} + x'_{1B} = x^*_{1A} + x^*_{1B} = \omega_1$, and $x'_{2A} + x'_{2B} = x^*_{2A} + x^*_{2B} = \omega_2$. Hence $p_1\omega_1 + p_2\omega_2 > p_1\omega_1 + p_2\omega_2$ which is a contradiction.

Now we go back to coalitional games with transferable payoffs – something the pure exchange economy was the most important economic example of. We define four axioms that "value functions" (values assigned to players, like the Shapley value) should satisfy. They are obvious. D.44 states that what is assigned to players should correspond to the value of the game. If it was higher, then what is gained from the game would not be sufficient to have all the players paid. If it was lower, then some gains from the game would be "wasted" (not allocated to the players). D.45 states that if players are renumbered (or renamed) their values do not change. D.46 – additivity – states that if two valuation systems are combined (perhaps one looking at some payoffs from the game, and another looking at payoffs of some other sort), then the value of any coalition is simply the sum of the two. Finally the dummy axiom (D.47) states that if there is a player who does not contribute to any coalition, then his (or her) value is zero. Formally these axioms are defined below.

(D.44)

A value function φ_i defined for each player in a coalitional game with transferable payoffs $\langle N, v \rangle$ is called efficient if $\sum \varphi_i(N, v) = v(N)$; i.e. 'no value is wasted'. (Summation extends over all the players.)

(D.45)

A value function φ_i defined for each player in a coalitional game with transferable payoffs $\langle N, v \rangle$ is called symmetric if $\varphi_i(N, v) = \varphi_h(N, v')$, where v and v' are identical except that the roles of players i and h are permuted.

(D.46)

A value function φ_i defined for each player in coalitional games with transferable payoffs $\langle N, v_1 \rangle$ and $\langle N, v_2 \rangle$ is called additive if $\varphi_i(N, v_1 + v_2) = \varphi_i(N, v_1) + \varphi_i(N, v_2)$, where $\langle N, v_1 + v_2 \rangle$ is a coalitional game with transfers defined through $(v_1 + v_2)(K) = v_1(K) + v_2(K)$

(D.47)

A value function φ_i defined for each player in a coalitional game with transferable payoffs $\langle N, v \rangle$ satisfies the dummy axiom if $\varphi_i(N, v) = 0$ for a player who does not contribute to any coalition, i.e. for every coalition K , $v(K \cup \{i\}) = v(K)$.

It turns out that the Shapley value (D.42) is the only value function satisfying all four definitions (D.44-D.47). In other words – despite earlier doubts caused by strange results referred to in the previous lecture – the values attached to various players cannot be defined in a different way, if certain obvious conditions are to be satisfied. The theorem below which states this is difficult to prove, so its proof will not be quoted here.

(T.31)

A value function φ_i defined for each player in a coalitional game with transferable payoffs $\langle N, v \rangle$ satisfying definitions D.44-D.47 is a Shapley value.

Questions and answers to lecture 13

13.1 What is the difference between total demand and excess demand?

The total demand is what a consumer demands at given prices. The excess demand (sometimes called net demand) is the difference between what the consumer demands and what he (or she) holds as the initial allocation (endowment). In the class example, the first consumer brought 8 apples and 2 oranges ($\omega_{11}=8$, and $\omega_{12}=2$). His (or her) total demand (final allocation) was 6 apples and 3 oranges ($x_{11}=6$, and $x_{12}=3$). Thus his (or her) excess demand for apples was -2, and the excess demand for oranges was 1. These numbers indicate how many units of a good the consumer wants to buy. The negative value of excess demand means that the consumer wants to sell it rather than buy.

13.2 Calculate the Walrasian equilibrium (including the price ratio) in a pure exchange economy game where the two players have the allocations of the two goods (4,4) and (6,1), respectively, and their indifference curves are given by the following hyperboles: $x_{2A} = \alpha/x_{1A}$, and $x_{2B} = \beta/x_{1B}$, respectively ($\alpha, \beta > 0$ – parameters).

This question is included in my overheads as E-13. Some students may recognise that the Edgeworth box is identical with what we see on page 89, except that the initial allocation corresponds to the upper intersection of indifference curves $x_{2A} = 16/x_{1A}$ ($\alpha=16$) and $x_{2B} = 6/x_{1B}$ ($\beta=6$). We will solve the problem algebraically.

Obviously the answer is the same as calculated in the class ($\mathbf{x}=[6,3,4,2]$ and $p_1=p_2/2$). Let us find x_{1A} , x_{2A} , x_{1B} , x_{2B} , α , β , and $p = p_1/p_2$. As before, it seems as if we had 7 unknowns. This number can easily be reduced to 5, since $x_{1B} = 10 - x_{1A}$ and $x_{2B} = 5 - x_{2A}$. As before, let $p = p_1/p_2$, $x_1 = x_{1A}$ and $x_2 = x_{2A}$. By the definition of the indifference curves:

$$\alpha = x_1 x_2, \text{ and } \beta = (10 - x_1)(5 - x_2).$$

Having calculated the derivatives with respect to x_1 of functions

$$x_2 = \alpha/x_1, \text{ and } x_2 = 5 - \beta/(10 - x_1),$$

and knowing that in the market equilibrium both indifference curves are tangent to the price ratio (the slope coefficients of their derivatives are equal to the price ratio), i.e.

$$-\alpha/(x_1)^2 = -p, \text{ and } -\beta/(10 - x_1)^2 = -p$$

allows us to calculate α and β : $\alpha = p(x_1)^2$ and $\beta = p(10 - x_1)^2$. This allows us further to calculate $\alpha = p(x_1)^2$, and $\beta = p(10 - x_1)^2$. When substituted to the definitions of indifference curves, it gives $p(x_1)^2 = x_1 x_2$ and $p(10 - x_1)^2 = (10 - x_1)(5 - x_2)$. After cancellations we get: $px_1 = x_2$ and $p(10 - x_1) = (5 - x_2)$. The rest is as in the class. By the necessity of equating expenditures with revenues, we get

$$p_1(4 - x_1) = p_2(x_2 - 4), \text{ i.e. } p_1/p_2 = (4 - x_2)/(x_1 - 4) \text{ (please remember that } x_1 \neq 4).$$

Hence the additional equation (needed in order to find the three unknowns) reads

$$p = (4 - x_2)/(x_1 - 4).$$

By substituting p into the previous equations we get:

$$x_1(4 - x_2) = x_2(x_1 - 4) \text{ and } (10 - x_1)(4 - x_2) = (5 - x_2)(x_1 - 4).$$

These are two equations with two unknowns: $4x_1 - 2x_1x_2 = -4x_2$ and $60 - 9x_1 + 2x_1x_2 = 14x_2$. By adding the two we get $x_2 = 6 - x_1/2$ which can be substituted into the first one, yielding

$$(x_1)^2 - 10x_1 + 24 = 0.$$

The last quadratic equation has two roots: $x_1 = 6$ or $x_1 = 4$ implying $x_2 = 3$ or $x_2 = 4$, respectively.

The second solution has to be rejected (in order to calculate p properly; i.e. – as noted above – to avoid dividing into zero).

13.3 Demonstrate that in a pure exchange economy where the two players have the allocations of the two goods (8,2) and (2,3), respectively, and their indifference curves are given by the following hyperboles – $(x_{2A}) = \alpha/(x_{1A})$ and $(x_{2B})^{1/2} = \beta/(x_{1B})^{1/2}$, respectively ($\alpha, \beta > 0$ – parameters) – the Walrasian equilibrium consists of the allocation (6,3,4,2) and the price ratio 1:2.

Please observe that by raising to the second power the definition of indifference curves for B you obtain the class example with constant β^2 instead of β . The calculations are the same as in the class, except that the constant β^2 should be put instead of β . Moreover, the calculations do not have to be carried out, since the Walras equilibrium has been calculated already: the allocation is (6,3,4,2), and the price ratio is 1:2. What needs to be done is to check that the slope coefficient of both indifference curves for the allocation (6,3,4,2) is the same and it is equal to -1:2.

The slope coefficient of the first indifference curve is $-\alpha/(x_1)^2 = -18/36 = -1/2$. The slope coefficient of the second indifference curve is $-\beta/(10 - x_1)^2 = -8/16 = -1/2$. Indeed they are the same and equal to -1/2 (that is to the price ratio).

13.4 Please prove (graphically or algebraically) that in a pure exchange economy game where two players have the allocations of two goods (8,2) and (2,3), respectively, and their indifference curves are given by the following hyperboles: $x_{2A} = \alpha/x_{1A}$, and $x_{2B} = \beta/x_{1B}$, respectively ($\alpha, \beta > 0$ – parameters), the allocation (7,3) and (3,2) is preferred by the first and offers the same well-being for the second, but it cannot be attained as a Walras equilibrium.

Graphical proof is perhaps easier, but it requires drawing very accurate pictures. Thus it is safer to solve the problem algebraically. First we will demonstrate that the allocation (7,3) is preferred by the first player. According to the formula $\alpha = x_{1A}x_{2A}$, the first (initial) allocation

gave the utility of 16, and the second one – 21. Hence the second is preferred by the first player. Now we will demonstrate that (3,2) gives the second the same utility as (2,3). But this is obvious, since $\beta = x_{1B}x_{2B}$ (in either case the product is 6). Hence the first part of the question is answered.

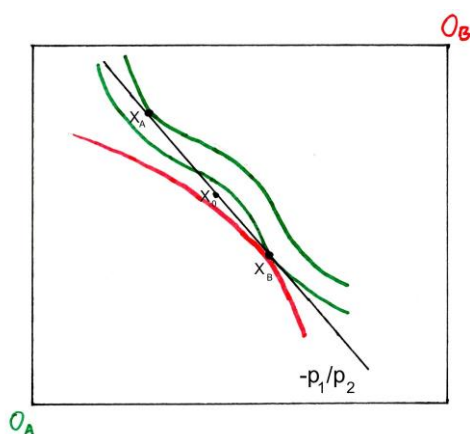
This allocation cannot be attained as a Walras equilibrium. In order to prove this, one has to recall that in Walras equilibria (D.43) offers of both consumers bring them to the highest indifference curve they can afford given the prices. In other words, their indifference curves must be tangent to each other (if they are not tangent – if they intersect like in the upper picture on page 89 – both consumers are better off when they do some additional exchange). Hence it is sufficient to check if the indifference curves are tangent or not. In (7,3) the slope coefficient of the tangent line to the first indifference curve is $-\alpha/(x_1)^2 = -21/49 = -3/7$. The slope coefficient of the tangent line to the second indifference curve is $-\beta/(10-x_1)^2 = -6/9 = -2/3$. This ends the proof.

13.5 Do lower left and upper right corners of an Edgeworth box (0_A and 0_B) belong to the contract curve?

Yes they do. The first corresponds to the allocation such that B does not have anything (everything is owned by A), and the second to the allocation such that A does not have anything (everything is owned by B). These are Pareto optima, since in either case there is no transaction that makes both of them better off simultaneously.

13.6 Is it possible that a Pareto optimum cannot be achieved as a Walras equilibrium?

Yes. The second welfare economics theorem lists mathematical assumptions that need to be satisfied to let Pareto optima be attained as Walras equilibria. One condition is that indifference curves are convex. Picture below demonstrates what may happen if the convexity assumption is violated.



This is an Edgeworth box for two consumers: the green one (A), and the red one (B). You may recall it from my overheads. The red one has convex indifference curves, while the green one does not. X_B is a tangency point of two indifference curves and therefore it is a Pareto optimum. The black line is a tangent to both indifference curves. Its slope coefficient ($-p_1/p_2$) gives the price ratio that can bring the consumers to this Pareto optimum. The two consumers have endowments X_0 . The point is located on this tangent line. But when the consumers face price ratio p_1/p_2 (the only one that can bring them to this Pareto optimum), then the red one

would like to sell some units of the first good, buy some units of the second good, and enjoy the point X_B . At the same time, the green consumer would like to sell (not to buy!) some units of the first good, buy some units of the second good, and enjoy the point X_A . The consumers have offers that are not consistent (they would like to buy what the other one does not want to sell), and a Walras equilibrium cannot be attained.

13.7 Please give an example of inefficient value functions φ_i (violating D.44).

Let us consider the capitalist production game (example III in the previous lecture GT-12) where the capitalist is assigned the value of $\varphi_c = mp/3$, and each of the $m > 1$ workers participating in the coalition $\varphi_w = p/2$. The sum of their values is $mp/3 + mp/2 = (5/6)mp < mp$. The sum of what the coalition members receive is less than what they produce ((1/6) of the game value is 'wasted').

13.8 Please consider the 'glove game' like in example II in the previous lecture (GT-12). Let us assume that the price of the pair of gloves consists of two components: 50% is for the bottom, and 50% for the top. Thus the price of the pair of gloves is 6 € for the bottom and 6 € for the top. Everything else is like in the class example. Assuming that players play two games: one to produce the bottom, and another one to produce the top, are their Shapley values additive?

Yes, they are. The Shapley values for these 'partial' games are 1 for the owners of left hand gloves, and 4 for the owner of the right hand glove. The Shapley values calculated in the class (2 and 8, respectively) are the sums from the two (hypothetical partial) games.

14 – Axiomatic bargaining

The problem analysed in two previous lectures (especially in GT-12) was how to guess what coalitions players are likely to make. The problem analysed now is different. We contemplate a coalition consisting of all (I) players, and look for players' desired behaviour which is needed to have such a coalition happen. And if such a coalition functions, how players can be motivated not to leave it. In other words, we analyse bargaining criteria that can be applied in order to have everybody participating in the game. Players bargain by manipulating payoffs available to them. Their sum is determined by the value of the game. But their allocation can be changed. If somebody is more important than somebody else (as – for instance – indicated by Shapley values), then he/she can expect to have a higher share in potential gains.

An important concept is the so-called 'starting point'. This is what happens if the game is not played. All the players refer to the same 'starting point'. If no agreement is reached – the 'starting point' continues; thus it can be considered a 'threat'. We do not care whether it is credible or not – this is a separate question.

The following definition sets the scene for the bargaining problem. We look for rules of choosing payoff combinations that turn out attractive for potential players.

(D.48)

- U – the set of all possible payoff combinations,

- $U \subset \mathbb{R}^I$
- U is closed and convex
- U satisfies 'free disposal' property (i.e. if $\mathbf{u} \in U$, and $\mathbf{u}' \leq \mathbf{u}$ then $\mathbf{u}' \in U$)

Non-zero two-person games as defined in D.14 do not satisfy D.48, because their payoff sets are finite (and a finite set consisting of more than one point cannot be convex). However, later on, D.14 was extended to cover infinite payoff sets. As before, elements of U can be interpreted as (expected) utilities.

The free disposal property needs to be explained. The condition $\mathbf{u}' \leq \mathbf{u}$ means that \mathbf{u}' is less attractive for everybody than \mathbf{u} (every coordinate of \mathbf{u}' is smaller or equal to the corresponding coordinate of \mathbf{u}). In other words, a less attractive outcome is feasible if a more attractive is. If some (negative) payoff is interpreted as a waste disposal, then the axiom says that more waste is feasible. This is not the case if the disposal is costly. The 'free disposal' name refers to the fact that players can get rid of the waste at no cost, which sometimes is not feasible. This property seems obvious, but for some applications it cannot be assumed.

The next definition looks complicated, but in fact it is not. It simply says that a bargaining solution $f(U)$ is to be selected from U . It is just one possible combination of payoffs; a combination that players agree to.

(D.49)

A function $f: 2^{\mathbb{R}^I} \rightarrow \mathbb{R}^I$, a rule that assigns a solution vector $f(U) \in U$ to every bargaining problem, is called a bargaining solution.

Now we will define several common sense axioms that bargaining solutions should satisfy.

(D.50)

The bargaining solution f is independent of utility origins (IUO) if $f(U) = f(U')$, where $U' = \{\mathbf{u}' \in \mathbb{R}^I: \exists \mathbf{u} \in U [\mathbf{u}' = (\mathbf{u}_1 + \alpha_1, \dots, \mathbf{u}_I + \alpha_I)]\}$ for some $\alpha = (\alpha_1, \dots, \alpha_I) \in \mathbb{R}^I$. It is understood that in U' the 'starting point' is $\mathbf{u}^* = (\mathbf{u}_1^* + \alpha_1, \dots, \mathbf{u}_I^* + \alpha_I)$ where \mathbf{u}^* was the 'starting point' in U .

Parameters $\alpha_1, \dots, \alpha_I$ can be understood as utilities of having nothing. Often it is assumed that $\alpha_1 = 0, \dots, \alpha_I = 0$, but this does not have to be the case. Some people declare negative utilities when they have nothing, and some – positive utilities.

(D.51)

The bargaining solution f is independent of utility units (IUU), or invariant to independent changes of units, if for any β ($0 < \beta \in \mathbb{R}^I$), $f_i(U') = \beta_i f_i(U)$ for all $i = 1, \dots, I$, whenever $U' = \{(\beta_1 u_1, \dots, \beta_I u_I): \mathbf{u} \in U\}$.

Coefficients β_1, \dots, β_I can be understood as 'exchange rates' between units applied by different people. For instance, if utilities are measured in money, but the first person works with British pounds, the second – with Norwegian crowns, the third – with American dollars, etc., then bargaining outcomes should be the same if we calculate payoffs in euros.

(D.52)

The bargaining solution f satisfies the Pareto property (P), or is Paretian, if for every U $f(U)$ is a weak Pareto optimum, i.e. there is no $\mathbf{u} \in U$ such that $u_i > f_i(U)$ for every i .

Please note the difference between a weak Pareto optimum, and a Pareto optimum defined earlier. A weak Pareto optimum is when it is not possible to make everybody better off; the concept of Pareto optimality defined earlier (see GT-11) was different: a Pareto optimum was when it was not possible to make one person better off. To see the difference, let us refer to the apples example, assuming that apples cannot be divided (only whole apples can be allocated). The allocation of four apples between two persons where the first person was given two apples, the second person one apple, and one apple was set aside did not satisfy the Pareto optimality, since it was possible to give the last apple to one person, without depriving the other one of anything. Hence this was not a Pareto optimum. But it was a weak Pareto optimum; this single apple can make one person better off, but it cannot make better off both of them (because it cannot be divided). This also explains why the concept is called 'weak': it is easier to achieve a weak Pareto optimum than a 'standard' Pareto optimum.

(D.53)

The bargaining solution f is symmetric (S) if whenever U is a symmetric set (i.e. U remains unaltered under permutations of axes) all the coordinates of $f(U)$ are equal (i.e. $f_1(U)=\dots=f_l(U)$).

In simple language D.53 says that if players are renumbered, then bargaining rules do not change.

(D.54)

The bargaining solution f satisfies individual rationality (IR) whenever $f(U) \geq \mathbf{u}^*$ (f is not worse than the 'starting point').

(D.55)

The bargaining solution f is independent of irrelevant alternatives (IIA) if whenever $U' \subset U$ and $f(U) \in U'$ it follows that $f(U') = f(U)$.

The last property is fairly obvious but its name should be explained. The set $f(U)$ contains an allocation such that all the players agree to, if they are to choose from U ; $f(U) \in U$. If $f(U) \in U'$, that is $f(U)$ is an element of a smaller set ($U' \subset U$), then it does not depend on any payoff combinations coming from $U \setminus U'$. These other combinations are called 'irrelevant alternatives' – something that does not have any role for finding the solution $f(U)$.

For notational simplicity, it will be assumed from now on that $\mathbf{u}^* = \mathbf{0}$ (justified by D.50). In other words, it will be assumed that in the starting point nobody enjoys any utility (if $\mathbf{u}_i^* \neq \mathbf{0}$, for the player i , then $\alpha_i = \mathbf{u}_i^*$ makes the required transformation).

Now we will look at the four most frequently used examples of bargaining solutions.

Example I (Egalitarian solution)

It is also called the Rawlsian solution: $f^e(U) = (u^e, \dots, u^e)$ where $u^e = \max\{\min\{u_1, \dots, u_l\} : \mathbf{u} \in U\}$.

Example II (Utilitarian solution)

It is sometimes called Benthamian solution. $f^u(U) = (u^u_1, \dots, u^u_l)$ where $u^u_1 + \dots + u^u_l \geq u_1 + \dots + u_l$ for all $\mathbf{u} \in U \cap \mathcal{R}_+^l$ (f^u maximizes the sum of utilities).

Example III (Nash solution)

$f^n(U) = (u_1^n, \dots, u_I^n)$ where $u_1^n \times \dots \times u_I^n \geq u_1 \times \dots \times u_I$ for all $\mathbf{u} \in U \cap \mathcal{R}_+^I$ (f^n maximizes the product of utilities). Alternatively it can be defined as a maximiser of the sum of logarithms of utilities $\ln u_1 + \dots + \ln u_I$.

Example IV (Kalai-Smorodinsky solution)

$f^k(U) = (u_1^k, \dots, u_I^k)$ where $u_i^k = \delta u_i^{\max}$ with

- $u_i^{\max} = \max \{u_i : \mathbf{u} \in U \cap \mathcal{R}_+^I\}$ and
- $\delta = \max \{\lambda : \lambda(u_1^{\max}, \dots, u_I^{\max}) \in U \cap \mathcal{R}_+^I\}$

All these solutions refer to some names. In the fourth lecture (GT-4) the name of John Nash was introduced. Please note however, that Nash solution defined in example III above has nothing to do with Nash equilibrium (these are two completely different concepts). John Rawls (1921-2002) is one of the most influential philosophers of the 20th century. He is best known for his concept of 'justice as fairness', that is justice understood as not taking the side of anybody, and acting 'behind the veil of ignorance'. If one does not know anything about who is going to benefit from a particular solution, then it is best to identify the least privileged one ($\{\min\{u_1, \dots, u_I\}\}$), and to maximise his/her position ($\max\{\min\{u_1, \dots, u_I\}\}$), like in the example I above.

Jeremy Bentham (1748-1832), an English philosopher, is considered the founder of the utilitarianism. His words "it is the greatest happiness of the greatest number that is the measure of right and wrong" can be translated into the criterion quoted in example II above. Ehud Kalai and Meir Smorodinsky are excellent game theorists. The solution they defined in example IV above is different from the Nash solution. The Nash solution finds payoff allocations such that players maximise their product. The maximum is achieved when gains from a hypothetical reallocation are equal for all the players (if these gains are not equal then the product can be increased by allocating more to where it is lower). The Kalai-Smorodinsky solution finds allocations such that relative gains from a hypothetical reallocation are equal. A relative gain is understood as a fraction of what a player can get from the game at maximum. The following example explains the difference.

Example V (Nash *versus* Kalai-Smorodinsky)

Let us assume that there are two players with the starting point 0. They can improve their position if they play the following game:

		Player 2	
		L	R
Player 1	U	(0,3)	(0,0)
	D	(0,0)	(6,0)

This game has two Nash equilibria: (U,L), and (D,R). The former gives higher utility to the second player, while the latter – to the first one. The latter offers the highest sum of payoffs. If these are transferable then (D,R) is the best the players can enjoy (jointly). If they make a coalition to play (D,R) then they need to agree on how the total payoff of 6 is to be allocated. If they apply the Nash bargaining solution, they will split it equally: 3+3. However, if they apply the Kalai-Smorodinsky solution they will split it 4+2 (since the maximum enjoyed by the first is twice as high as the maximum enjoyed by the second; δ in the example IV above is equal to 2/3).

Definitions of various bargaining solutions can be confronted with the six axioms (D.50-D.55). It turns out that the Nash solution satisfies all of them. Moreover, it is the only solution that satisfies the first five of them, as the following theorem (too difficult to prove here) states.

(T.32)

The Nash solution is the only bargaining solution which satisfies IUO, IUU, P, S, and IIA.

Proving its compliance with the sixth axiom (T.33) is easy.

(T.33)

The Nash solution satisfies IR.

Proof:

By definition, $f^n(U) \geq 0$ as a product of non-negative numbers. By IUO and IUU 0 can be chosen as the 'starting point'.

Questions and answers to lecture 14

14.1 Any finite set with at least two different points cannot be convex. Do you see why?

The definition of a convex set X states that for every $x, y \in X$ and for every $\lambda \in [0, 1]$

$$\lambda x + (1 - \lambda)y \in X$$

(the entire segment with the ends x and y lies in X). Thus, if X contains two different points x and y , it has an infinite number of points.

14.2 Electricity production does not satisfy the 'free disposal' property (D.48). Can you see why?

Let us assume that electricity technology is characterised by two numbers: u_1 – the quantity of kWh produced, and u_2 – the quality of the ambient air. The latter is a monotonically decreasing function of pollutants emitted. The inequality $\mathbf{u}' \leq \mathbf{u}$ is satisfied when $u'_1 = u_1$ and $u'_2 < u_2$ (the same electricity production, and deteriorated environment caused by the emission). This is not necessarily possible. For instance in a centrally planned economy in Poland – especially in the 1950s – power plants emitted so much, that their buildings suffered from the excessive pollution; their corrosion was catastrophic. As a result, the number of kWh produced could not be sustained. Apparently waste disposal was not 'free'.

14.3 Consumer preferences \leq are reflected by a utility function u if and only if $x \leq y$ whenever $u(x) \leq u(y)$ (preferred bundles are characterised by higher utility). There are no other requirements. In particular, if u' is a strictly increasing transformation of u , it reflects the same preference relationship. This property of utility functions led to distinguishing between 'ordinal' and 'cardinal' utilities. The former reflect preferences and nothing more; the latter inform about the 'intensity' of preferences. For instance, when $u(x) = 2u(y)$, we say that the

consumer enjoys x twice as much as y . Do definitions D.50 and D.51 refer to 'ordinal' or 'cardinal' utilities?

They refer to 'ordinal' ones. They do not make use of specific numbers. Translating utility functions (D.50) and scaling them (D.51) can change their values dramatically. In contrast, vNM theory refers to 'cardinal' utilities. Calculating expected values requires that preferences with respect to bundles can be compared not only 'ordinally' (better – worse), but also 'cardinally' (how much better).

14.4 Is IIA axiom (D.55) always reflected in negotiation processes?

Not necessarily. Some game theorists say that the entire set of alternatives may affect the agreement reached. Let us suppose that players look at three options: A, B, and C. Let us suppose also that the preference relation of player 1 is: $C \gg B \succ A$ (C is much better than B, which is somewhat better than A) while the preference relation of 2 is reversed: $A \gg B \gg C$ (as before, \gg reads "much better"). The fact that option C becomes available allows player 1 to say: "if I give up my best option (that is C), I have a right to demand that at least my second-best option (B) will be chosen". If only options A and B are available, probably A will be chosen (because by 2 it is preferred more strongly over B than B over A by 1). However, if the set of options contains C as well, 1 can argue for B convincingly.

14.5 Is the Rawlsian concept of acting 'behind the veil of ignorance' realistic?

When I first learned about the concept, I thought that it was a theoretical idea only. Acting 'behind the veil of ignorance' seemed to be practically impossible. When taking decisions about how to organise health care, people are supposed not to know whether they are healthy or not; when taking decisions about how to finance old age pensions, people are supposed not to know whether they are young or not; when taking decisions about how to use exhaustible resources, people are supposed not to know whether they live in the 20th century or earlier (or later).

I started looking at the Rawlsian concept of acting 'behind the veil of ignorance' as something realistic when a drug dealer was found in the high school attended by my children. The principal convened a meeting of all the parents to discuss a decision to be taken with respect to the drug dealer. The parents asked who was found to be the drug dealer. But the principal said that she would not tell. They asked which class the person attended, but she did not tell. They asked about whether this was a girl or a boy, but she did not tell either. The decision was to be taken without any knowledge of that sort. This was exactly acting 'behind the veil of ignorance' like in the Rawls' concept.

At the first glance the Rawls' concept of acting 'behind the veil of ignorance' seems to be a purely theoretical idea. Nevertheless it can be implemented sometimes in order to stimulate fair decisions (when decision takers do not know whether what they decree affects themselves or not).

14.6 Please discuss the statement made by Jeremy Bentham: "it is the greatest happiness of the greatest number that is the measure of right and wrong".

The statement – the foundation of utilitarianism – may cause objections. If I am to be worse off by a little bit, but – as a result – somebody else is to be much better off, then the

reallocation is considered by Jeremy Bentham as an improvement (the total happiness increases). There are two problems with this approach. First, utilities (especially 'ordinal' utilities) of various consumers are difficult to compare: it is difficult to judge whether what one loses is less than what someone else gains. Second, it is difficult to implement a reallocation which is not a Pareto improvement (the losers may vote against even if what they lose is small when compared to what the winners gain).

Yet utilitarianism is appealing. Implementation problems can be solved by taking decisions by appropriate bodies, not by the people affected directly (we pay taxes even though we would like to pay less if we were asked to quote the "right" amount). Besides, the society considers it appropriate if (moderate) losses of some of its members are compensated by gains of others.

14.7 In order to identify the best promising students, some universities add the score in mathematics to the score in history, and some multiply the scores. Please discuss pros and cons of both approaches.

Let us assume that there are two high school graduates. One (A) scored 20 in mathematics and 90 in history, and another one (B) scored 50 in mathematics and 50 in history. Who is better? The total score of A is 110 (55 per field on average), and the total score of B is 100 (50 per field on average). Thus from this point of view, A is better than B. But if products instead of sums are taken into account, then A is characterised by 1800 (geometric average is 42), and B is characterised by 2500 (geometric average is 50). Thus from this point of view B is better than A. How applications of these two candidates should be assessed by the university they apply to?

A is better than B from the point of view of the total score. He/she did poorly in mathematics, but this deficit was more than compensated by history; at least in one field the candidate performed very well. But from the point of view of covering the entire spectrum of competencies developed by the high school, B is better than A (there is no field where he/she did very poorly).

My class questions revealed that some of you would prefer to have colleagues resembling A or B. Also in the case of university preferences, it would be up to broader criteria of how to measure the adequacy of a solution. If we look for a combination that is characterised by a high arithmetic average, then the first approach is appropriate. If we do not like extremes, then the second approach (based on the geometric average) should be favoured. The Nash solution aims at eliminating allocations that are very asymmetric (e.g. good in mathematics and bad in history or *vice versa*).

14.8 An alternative definition of the Nash solution is to maximise the sum of logarithms of payoffs (rather than the product of payoffs). Please prove the equivalence.

In the first definition in example III (page 3)

$$f^n(U) = (u_1^n, \dots, u_l^n) = f(u_1^n) \times \dots \times f(u_l^n) = \max \{ f(u_1) \times \dots \times f(u_l) \text{ for all } \mathbf{u} \in U \cap \mathbb{R}_+^l \}.$$

In the alternative definition, the maximisation of $f(u_1) \times \dots \times f(u_l)$ is substituted with the maximisation of $\log f(u_1) + \dots + \log f(u_l)$. According to the formula known from the high school: $\log(a \times b) = \log a + \log b$. Logarithm with a base higher than 1 (usually its base is e

and the logarithm is denoted as \ln , or its base is 10 and the logarithm is denoted as \log_{10}) is a one-to-one strictly increasing function. Thus the maximum of $a \times b$ can be found by maximising $\log(a \times b)$.

14.9 The so-called Jensen inequality for concave functions states:

$$g(x_1) + \dots + g(x_n) \leq n \times g(c/n), \text{ if } x_1 + \dots + x_n = c.$$

What does this imply for bargaining solutions?

Let us take a logarithm as an example of the concave function g , $x_1 = u_1, \dots, x_n = u_n$, and c as the value of a game (the total utility to be allocated among the players). The Jensen formula reads:

$$\log u_1 + \dots + \log u_n \leq n \times \log(c/n).$$

The left hand side is the alternative definition of the Nash solution (example III above). The inequality says that the sum of logarithms is maximised, if the payoff is distributed uniformly.

Another example of a concave function is the utility observed by risk averse people (see T.5 in GT-2). If the gains from a game are to be distributed to such people, the uniform distribution maximises their sum.

14.10 Can the Kalai-Smorodinsky solution be identical to the Nash one?

Yes. The Nash solution equates potential gains from the coalition enjoyed by all the players. Kalai-Smorodinsky solution equates potential relative gains from the coalition enjoyed by all the players. If relative gains (that is gains understood as proportions of gains enjoyed when claiming the maximum payoff) are the same as absolute gains, both solutions identify the same allocation. This happens, for example, when the game is symmetric (all the players have identical maximum payoffs). The game from example V was not symmetric.

15 – Evolutionary game theory

The last question addressed in this class is about the stability of Nash equilibria. Nash equilibrium has been the most important concept in game theory. An interesting problem is what happens if some terms of a game change slightly. This new game will have a new Nash equilibrium (at least in mixed strategies). Will it be much different from the original one? Or perhaps the old one will maintain its characteristics.

This problem was looked at in the eleventh lecture (GT-11) when we introduced the concept of "trembling hand perfection". The assumption was that one of the players applied a 'wrong' strategy (that is a strategy that was not envisaged in a Nash equilibrium), and the question was what would be the best response of the other one. Here we will analyse the problem in a more systematic way. It turns out that stability depends both on the game structure, and on the environment structure.

Let us start with the oligopolistic competition. In order to keep things simple, it will be assumed that there are only two producers who compete with each other *à la* Cournot, that is by setting quantities (T.27; see also the answer to question 10.2). Let us sum up the model.

Example I (Cournot duopoly)

- Two producers supply the same market
- The price is implied by the total supply: $p=a-b(y_1+y_2)$
- Both firms have identical linear cost functions and $MC_1=MC_2=AC_1=AC_2=c=\text{const}$
- The first rival makes a quantity decision y_1 that maximizes its profit expecting that the second rival does the same (with respect to y_2)
- Thus they solve two maximization problems:
 - $\max_{y_1}\{(a-b(y_1+y_2))y_1-cy_1\}$
 - $\max_{y_2}\{(a-b(y_1+y_2))y_2-cy_2\}$

FOC (*First Order Conditions*) for these problems are:

- $a-2by_1-by_2-c=0$
- $a-2by_2-by_1-c=0$

Solving these equations yields


- $y_1 = y_2 = (a-c)/3b$
- $y = 2(a-c)/3b$
- $p = (a+2c)/3$

This is a 'standard' Cournot-Nash solution.

Rewriting the FOC one can yield

- $y_1=(a-c)/(2b)-y_2/2$, and
- $y_2=(a-c)/(2b)-y_1/2$

These are so-called reaction curves, i.e. the best response (profit maximizing output) of one player given the decision (output) of the other one. To simplify calculations from now on, we put specific numbers: $a=9$, $b=1$, and $c=3$. Reaction curves are thus $y_1=3-y_2/2$, and $y_2=3-y_1/2$, and the Nash equilibrium is (2,2).



Let us assume now that players missed the equilibrium. For instance, let $y_1^0=3$ (instead of 2). Assuming that this choice will be repeated, in the second round the rival may wish to choose

$y_2^1 = 3 - 3/2$ (according to the reaction curve), i.e. $y_2^1 = 3/2$. If the first player expects this choice to be repeated then his/her best reaction in the next round will be $y_1^2 = 3 - (3/2)/2 = 9/4$. In yet another round the likely supply offered by the second one may be $y_2^3 = 3 - (9/4)/2 = 15/8$. Then $y_1^4 = 3 - (15/8)/2 = 33/16$, $y_2^5 = 3 - (33/16)/2 = 63/32$, and so on.

It is easy to see that both players are likely to approach (2,2), i.e. to the Nash equilibrium. The example suggests that the Nash equilibrium in the Cournot game is stable. Even if the players miss it, they are likely to move towards it if they play the game repeatedly (as picture below illustrates).


Let us now look at a situation called in economic textbooks the 'cobweb model'. This strange name becomes clear when one sees the graphs. It can be considered a game between producers (who control the supply curve) and customers (who control the demand curve).

Example II ('Cobweb model')

The so-called 'cobweb model' refers to a competitive market characterized by a demand curve $p = a - by$, and a supply curve $p = c + dy$ ($a, b, c, d > 0$, and $a > c$). Equilibrium quantity is $y = (a - c)/(b + d)$, and equilibrium price is $p = (ad + bc)/(b + d)$. An additional assumption is that adjustments are done in discrete time intervals (e.g. years), suppliers react to market prices, but the quantity they offer in the next period is not flexible (i.e. it corresponds to the previous period price rather than to the present one).


This can be modelled as a game with one player determining the demand and the other one – the supply. Assuming that disequilibrium is always considered a bad outcome (if the price is higher than expected, suppliers regret that they did not offer more; if the price is lower than expected, suppliers regret that they did not offer less), its Nash equilibrium is when the actual price is equal to the expected one.

If the equilibrium quantity and price were achieved, then $y^t = (a - c)/(b + d)$ and $p^t = (ad + bc)/(b + d)$ for $t = 0, 1, 2, \dots$. Let us see what happens if the price faced by the suppliers was not their expected price. To simplify calculations from now on, we put specific numbers ($a = 12$, $b = 2$, $c = 0$, $d = 1$). Thus the demand curve is $p = 12 - 2y$, and the supply curve is $p = y$. The equilibrium quantity is $y = 12/3 = 4$, and the equilibrium price is $p = 12/3 = 4$.



Let us assume that for whatever reason $p^0=5$ (higher than the equilibrium price). Reacting to this price, the supply will be $y^1=5$. But with this quantity offered (higher than the equilibrium quantity), the demand will bring the price to $p^1=2$. The next period supply will be thus $y^2=2$, and the price $p^2=8$. It is easy to see that the quantities and prices will diverge from the equilibrium. Hence the equilibrium outcome is not stable: if players miss the equilibrium they will never return to it.

Picture above illustrates the mechanism of a divergence process rather than exact arithmetic solutions (corresponding to numbers adopted). But the trajectory can be different, like in the following graph.



The red trajectory resembles a cobweb – like in the previous picture – but this time prices and quantities converge to the equilibrium (y,p) . It is easy to see that the trajectory is diverging when the slope of the demand curve is steeper than that of the supply curve (one of the students observed correctly, that it is linked to elasticities). It is converging if the relation of slopes is reverted. If the slopes are exactly the same, then the trajectory will be a closed loop (a rectangle). In the first case (in the example II above; see previous page) if the market started from a non-equilibrium, it will never be there. In the second case, if the market started from a non-equilibrium, it will converge to it gradually.

The next example explains why what we analyse today is called "evolutionary game theory". It refers to what may happen if a peaceful individual meets an aggressive one.

Example III (Mutation)

Let us consider the following 'Dove-Hawk' game:

		Second	
		D	H
First	D	(4,4)	(0, 8)
	H	(8,0)	(3,3)

If two animals find a prey they can behave like doves, i.e. share it peacefully. If one behaves like a dove and the other like a hawk, the militant one gets everything and the other one is left

with nothing. If both of them behave like hawks, they lose a quarter of the prey (because of fighting) and share the remaining part equally. The game has a unique Nash equilibrium (H,H). In other words, in equilibrium all animals are likely to behave as hawks.

Now let us assume that in principle all animals behave like hawks, but there are few 'mutants' who behave like doves. Their behaviour is hereditary, and an individual never changes it. Thus the 'game' they play has a probabilistic structure. The payoffs depend on probabilities of encounters (D-D), (D-H or H-D; the matrix is symmetric), and (H-H).

Let us assume that ε is the proportion of 'mutants' (i.e. those who turn out to be doves). Thus the proportion of 'non-mutants' (behaving like hawks) is $1-\varepsilon$. The expected outcome of the encounter of two individuals (in terms of 'payoffs' of the Dove-Hawk game) is 4ε for a 'mutant' and $5\varepsilon+3$ for a 'non-mutant'. Actually the expected payoff for 'mutants' is $4\varepsilon^2$, and for 'non-mutants' it is $(1-\varepsilon)(5\varepsilon+3)$, but taking into account their shares (ε and $1-\varepsilon$, respectively), it is 4ε and $5\varepsilon+3$ per individual. For a small ε the second number is larger than the first one, so the 'mutant' population is likely to have smaller offspring (because of hunger) and consequently its share will go down to zero.

The game above explains what may happen (the dovish mutation will disappear), but it does not give any clue as to how much time this may take. The following concepts aim at answering this question.

(D.56) Replicator game

Let us assume that a population consists of two types of individuals who play a symmetric D.14 (two-person, non-zero sum) game. However, they do not choose strategies, but each is determined to apply a certain pure strategy and they 'pass it over' to their offspring. Let ε_i be the statistical frequency of individuals who play strategy i ($\varepsilon_1+\varepsilon_2=1$). It is further assumed that the population reproduces asexually (i.e. an individual has offspring without matching) and 'reproduction success' is proportional to the payoff from the game.

The assumption of asexual reproduction is just for mathematical simplicity. Solving models with sexual reproduction would require to adopt an assumption on what happens if a mutant breeds with a non-mutant (whose behaviour is inherited by their offspring?).

(D.57) Evolutionary Stable Strategy

A population plays a Replicator game (D.56). Let there be a pure-strategy Nash equilibrium (i^*, i^*) in the game. It is called an Evolutionary Stable Strategy (ESS) if

$$\varepsilon_i^* P_{i^*, i^*} + \varepsilon_{-i^*} P_{i^*, -i^*} > \varepsilon_{-i^*} P_{-i^*, i^*} + \varepsilon_i^* P_{i^*, -i^*} \text{ (likewise for the second player's payoffs).}$$

It is easy to see that being 'hawkish' in the 'Dove-Hawk' game is an ESS ('dovish' population will shrink). The definition of ESS can be extended to more complicated games, e.g. consisting of more than two strategies.

There are Nash equilibria that do not determine an ESS. To see this, consider the game below. It has two Nash equilibria – (I,I) and (II,II) – but only the first one determines an ESS. To see this, please note that when I serves as i^* in the inequality above, we have ε_i at the left hand side, and 0 at the right hand side.

		Second	
		I	II
First	I	(1,1)	(0,0)
	II	(0,0)	(0,0)

(D.58) Replicator equation

A population plays a Replicator game (D.56). Let $\varepsilon_i(t)$ be the frequency of individuals of type i in the population at time t . A formula determining $\varepsilon_i(t)$ as a function (perhaps an indirect one) of time and the game characteristics is called a Replicator equation.

Example IV (Replicator equation)

The formula from D.58 may take the following (differential) form:

$$d\varepsilon_i(t)/dt = \varepsilon_i(t)(f_i(\varepsilon_1(t), \varepsilon_2(t)) - \varphi(\varepsilon_1(t), \varepsilon_2(t))),$$

where $\varphi(\varepsilon_1(t), \varepsilon_2(t)) = \varepsilon_1 f_1(\varepsilon_1(t), \varepsilon_2(t)) + \varepsilon_2 f_2(\varepsilon_1(t), \varepsilon_2(t))$ for $i=1,2$.

Function $f_i: [0,1]^2 \rightarrow \mathbb{R}$ is a 'fitness' of type i to the environment characterized by the game and distribution of types in the population, and φ is the average 'fitness' in the population.

The Replicator equation predicts that the share of the type that is less fit than the average will decrease (the right hand side of the differential equation is negative), and the share of the type that is more fit than the average will increase (the right hand side of the differential equation is positive). It also lets estimate how fast the process will be.

Please note that the 'Dove-Hawk' game (Example III) predicts that the share of 'dovish' individuals will shrink, but – without a Replicator equation – it cannot be more specific about the pace of this process.

The definition of Replicator game (D.56) can be modified by assuming that the population is homogeneous (one type only), but there are two strategies to choose. As before, ε_i is the statistical frequency of individuals who play strategy i ($\varepsilon_1 + \varepsilon_2 = 1$). Unlike in the original definition, playing a strategy is not only inherited, but it can also be changed by an individual. The change of ε_i is determined by a modified 'Replicator equation':

$$d\varepsilon_i(t)/dt = \alpha\beta\varepsilon_i(t)(f_i(\varepsilon_1(t), \varepsilon_2(t)) - \varphi(\varepsilon_1(t), \varepsilon_2(t))),$$

where α is the learning rate (how fast individuals learn that changing strategy can be beneficial), and β is the willingness to change (if an individual knows that the other strategy is more attractive, is it willing to change the behaviour?); other symbols are as in D.58

D.57 (ESS) can be extended to account for Replicator games with more complicated characteristics (such as learning/adoption).

There are two interpretations of the 'Dove-Hawk' game: a statistical one and a strategic one. In the statistical version, individuals are born to behave like doves or hawks, and they cannot change it. In the strategic version they can choose to behave like doves or hawks; they select

the behaviour which gives them a better chance to reproduce. The idea is fairly simple, but the theorem which summarises it is difficult to prove. It will be quoted here without a proof.


(T.34)

Let (i^*, i^*) be an ESS. Then the system of Replicator equations $d\varepsilon_i(t)/dt = \varepsilon_i(t)(f_i(\varepsilon_i(t)) - \phi(\varepsilon_i(t)))$ for $i=1,2$ has a solution $\varepsilon_i^*(t)=1=\text{const.}$ It is called a stationary solution.

Questions and answers to lecture 15

15.1 Please draw the trajectory of market supply offered by Cournot duopolists from example I.

The trajectory looks more or less like the red line in the picture below (replicated from GT-15). No matter what quantities the duopolists start with, they will approach the Nash equilibrium (i.e. where their response curves intersect).



15.2 Can cobweb models explain how real markets function?

Some people doubt whether market agents are so stupid that they cannot guess future prices. Indeed, many firms construct sophisticated models to predict future prices. They try to plan their supply based on predicted rather than past prices. Nevertheless in some sectors – e.g. in agriculture – the supply is difficult to change. Moreover, projections are uncertain, and it is very difficult to plan the supply which maximises the profit. Cobweb models turned out to be quite adequate for explaining market behaviour in some sectors.

15.3 Let us assume that in the 'Dove-Hawk' game (example III) most individuals are 'dovish' and only mutants are 'hawkish'. How does this change game-theoretic predictions?

Let us assume that ε is the proportion of 'mutants' (i.e. those who turn out to be hawks). Thus the proportion of 'non-mutants' is $1-\varepsilon$. The expected outcome of the encounter of two individuals (in terms of 'payoffs' of the Dove-Hawk game) is $8-5\varepsilon$ for a 'mutant' (a hawk) and $4-4\varepsilon$ for a 'non-mutant' a dove). For a small ε the first number is larger than the second one, so the 'mutant' population is likely to have larger offspring (because of food availability) and consequently its share will increase. The 'non-mutants' (i.e. 'dovish' individuals) are likely to

disappear gradually. The conclusion is the same as when 'mutation' is understood as switching from 'hawkish' to 'dovish' (the 'dovish' individuals are doomed to extinction). The only thing that cannot be determined without a replicator equation is the speed of this process. See question 15.5 below.

15.4 Let a Replicator equation (see D.58) read:

$$d\varepsilon_i(t)/dt = \varepsilon_i(t)(f_i(\varepsilon_1(t), \varepsilon_2(t)) - \varphi(\varepsilon_1(t), \varepsilon_2(t))),$$

where $\varphi(\varepsilon_1(t), \varepsilon_2(t)) = \varepsilon_1 f_1(\varepsilon_1(t), \varepsilon_2(t)) + \varepsilon_2 f_2(\varepsilon_1(t), \varepsilon_2(t))$, $f_1(\varepsilon_1(t), \varepsilon_2(t)) = \Delta/\varepsilon_1$, and $f_2(\varepsilon_1(t), \varepsilon_2(t)) = -\Delta/\varepsilon_2$. For every t , Δ (a small positive number) is calculated in such a way that $\varepsilon_1(t) \leq 1$, and $\varepsilon_2(t) \geq 0$. Then $\varphi(\varepsilon_1(t), \varepsilon_2(t)) = \varepsilon_1 \Delta/\varepsilon_1 - \varepsilon_2 \Delta/\varepsilon_1 = 0$. Which population is expected to grow, and which is expected to shrink?

The first one is expected to grow, and the second one is expected to shrink. Please note that the differential equation is reduced (after cancellations) to $d\varepsilon_i(t)/dt = \Delta$. Thus ε_1 increases and ε_2 decreases.

15.5 Please consider a modified exercise E-15 from the overheads to this lecture. Let it be phrased as follows. The 'Dove-Hawk' game from Example III is supplemented with the following 'Replicator equation': if the expected payoff of the mutant in the game at time t is lower than 1 then its share in the population is reduced to $1/3$, that is $\varepsilon(t+1) = \varepsilon(t)/3$; if the expected payoff of the mutant in the game at time t is lower than $1/10$ then $\varepsilon(t+1) = 0$. Discuss how much time it will take to eliminate mutants from the population.

The answer depends on what the original mutant population was. Please note that the expected payoff for a mutant is 4ε (calculated in the class as the relevant average payoff). If $4\varepsilon(0) < 1/10$ then $\varepsilon(1) = 0$, i.e. mutants are eliminated in one period. Otherwise – if $1/10 \leq 4\varepsilon(0) < 1$, i.e. $0.025 \leq \varepsilon(0) < 0.25$ – it will take 2 periods for $\varepsilon(t)$ to fall below 0.025 . For instance, if $\varepsilon(0) = 18\%$, then $\varepsilon(1) = 6\%$, and $\varepsilon(2) = 0$ (because 6% is less than $1/10$). The wording of the exercise – the lack of a more precise 'Replicator equation' – does not allow predicting what will happen if $\varepsilon(0) \geq 0.25$, that is when in the beginning the expected payoff is 1 or larger.

15.6 The Replicator equation on page 109 was modified (compared to its original formulation in example IV) by adding the term " $\alpha\beta$ " to account for learning and willingness to change. Discuss what this additional term implies.

The Replicator equation determines the share of the type that is less or more fit than the average (it predicts that the share of the less fit type will decrease, and the share of the type that is more fit than the average will increase). Assuming that both coefficients are between 0 and 1, the added term does not change this conclusion, but it may change the pace of the process.

In the 'Dove-Hawk' game in its strategic version (individuals do not inherit behaviour, but they choose it), we look at mechanisms leading to adopting a 'hawkish' behaviour by individuals who behaved 'dovish' before. They learned that in a cruel world it is better to behave like a hawk. The coefficient α reflects how fast an individual learns that changing behaviour leads (in terms of expected payoffs) to a better outcome. If it is close to 0, they learn very slowly. If it is close to 1, they learn faster.

Yet the awareness of what behaviour implies a better fit to the environment does not imply that an individual changes it instantly. Some of them will stick to the past behaviour even though they know, that changing it would be beneficial. The coefficient β reflects the fact whether an individual is willing to change the behaviour that was found to be unfit to the environment. In biology this can be caused by some natural circumstances. In economics this can be caused by some social or psychological circumstances. For instance, consumers may hesitate to change their habits, and firms may hesitate to apply a new technology even when they know that it would be beneficial.