

# Experimental economics

## Seminar II - Methodological foundations

Matej Lorko

matej.lorko@euba.sk

Student resources: [www.lorko.sk](http://www.lorko.sk)

### References:

- Weimann, J., & Brosig-Koch, J. (2019). Methods in experimental economics. Springer International Publishing.

# Methodological foundations

- The times are long gone when economists “just did an experiment” to see what happens when you let subjects make specific decisions. In the meantime, methodological standards and procedures have evolved. Following these procedures is an important prerequisite for obtaining experimental results that can claim to meet the scientific standards of the economics profession.
- What is actually a good experiment? A good experiment succeeds in testing the most important effect a theory describes, while at the same time controlling for all plausible alternative explanations. This is based on the fact that there can often (if not always) be more than one potential explanation for an empirically observable regularity. Good experiments are undoubtedly those that, from the many alternative causal relationships, can find the one that actually explains the observed phenomenon.
- It should be borne in mind, however, that this is a very ambitious goal that can only relatively rarely be achieved. This is simply because not all the alternatives are always known. Therefore, the design of an experiment should eliminate the possibility that the results are influenced by factors that are not connected to the assumed causal relationships. If a faulty design or the inept execution of an experiment is responsible for the observed regularity, then one has certainly not carried out a good experiment.
- For this reason, it is worth considering the way in which various elements of an experiment’s design can influence the behavior of the subjects. This is the only way to be certain that the results obtained are an empirically significant regularity and not just an artifact of a bad experiment.

# Actions and rewards

- At the heart of all experimental investigations is the ability to make observations under controlled conditions. For example, experiments with animals attempt to study their learning ability. For this purpose, certain signals are combined with rewards (in the form of food) and the researchers observe whether the laboratory animal is able to make a connection between the signal and the reward associated with it.
- By controlling the variation of signals and rewards, conclusions can be drawn about the animals' ability to learn. It is implicitly assumed in such experiments that the animals like the food offered as a reward and that they prefer to have more of it. Only then is it justifiable to assume that the animals are making an effort (learning is strenuous) to obtain the food. Now the assumption that apes, for example, like to eat sweet fruit and are particularly keen on certain "treats" is not too daring. It is easy to see that this is the case. The behavioral hypothesis for an ape is therefore "I prefer more bananas to fewer bananas".
- But what about experiments with humans? Ultimately, economic theory describes above all else how people make decisions. Of central importance are the preferences that are attributed to the actors. Rational choices are always related to the goal being pursued. Thus, it is not possible to make a prediction about what a rational decision-maker will do if it is not known which goal this person is pursuing.
- Experiments that test theories and also experiments that are not based on theory would therefore be worthless without assumptions about the underlying preferences. The problem is that people's preferences are doubtlessly more differentiated than those of apes. In other words, it makes little sense to assume that people prefer more bananas to fewer bananas. Fortunately, however, the difference between us (humans) and our closest genetic relatives is not so great, thanks to the fact that there is a banana equivalent of sorts for us, and that is money.

# Induced value method

- Vernon Smith introduced this equivalent systematically into experimental methodology in 1976 and gave it a name: the induced value method. The idea is very simple. It is assumed that the consumption of every good generates utility, for which there is a monetary equivalent – the willingness to pay for the good in question. If each utility value can be expressed in terms of money, then the utility function can also be replaced by a “money function”, and by introducing this money function into the experiment as a “payoff function”, one has induced from the outside the utility function that is used for the evaluation of options for action. The following example best illustrates this process.
- The induced value method requires that people react to money in the same way as apes react to bananas – more of which is always better than less. This is also one of a total of three requirements that Smith specifies need to be met for the induced value method to be applied. First, the utility function must grow monotonically in terms of money. In slightly more technical terms, if a decision-maker can choose from two alternatives and one of them has a higher payoff than the other, then the decision-maker will always choose the alternative with the higher payoff.
- Second, the payoffs have to be salient. The so-called salience requirement is understood as meaning that the decision to be taken by a subject in an experiment must also be payoff-relevant. It is worth considering a little more deeply what this implies. It is important that the payoff function is not too flat. If it is, taking different decisions has little impact on the resulting payoffs, which can result in subjects not putting much effort into actually making what for them is the best decision, because a mistake has little financial impact.
- The third requirement that Smith lists is the dominance of the payoffs. As an experimenter, one has to be aware of the fact that experimental subjects could also have other things on their minds than the money they can earn in an experiment. For example, people dislike getting bored or thinking that their precious time is being wasted. Boredom may lead to subjects having an incentive to make things more interesting by trying things out without paying too much attention to their payoffs. There are many other factors that might discourage subjects from focusing exclusively on maximizing their payoffs.
- They could form expectations about what the experimenter wants from them and behave accordingly (the experimenter demand effect mentioned earlier). They might also make social comparisons and try to outperform the other subjects by doing things that adversely affect them. They may even develop altruistic feelings or think about fairness. All this and much more is possible. What Smith means by “dominance” is that, despite all these distractions, the pursuit of the highest possible payoff still comes first, and in case of doubt, the alternative that ensures the highest payoff is chosen.

# Induced value method

- Any assessment of whether the induced value method works or not depends on the strength of the various motives that drive people in the laboratory. Economists tend to assume that the desire to earn more income overshadows everything else and is much stronger than the above-mentioned “distractions”. Psychologists take a slightly different view. Behavioral economics deals with phenomena that lead people to make distorted decisions precisely because they allow themselves to be distracted from their actual goal. And psychologists know that people are probably much more susceptible to such distortions than economists generally believe.
- Even if it is the case that the motive to earn money dominates the experiment, it may be that more money is not always better. Two examples are worth mentioning. First, it is possible that within the experiment there are non-monetary subjective costs which are not taken into account in the payoff function but which nevertheless influence the decision of the subjects.
- A simple example is the cost associated with reading and understanding instructions. How high these costs are depends on the complexity of the experimental design, on the one hand, and on the ability of the subjects to grasp the facts, on the other. If a subject finds it difficult to understand the task presented in the experiment, he will probably deviate from a payoff-maximizing decision due to the much greater effort required to determine it.
- A second example is the existence of altruistic preferences. If a decision-maker has such preferences, it is obvious that more own money need not always be better than less own money.

# Achieving the best possible control of the subjects' preferences

- The payoffs should be noticeable. This means that the subjects can only be expected to pay attention to the payoffs if they are structured in such a way that it is worth paying attention to them.
- Subjective costs should be minimized. This means that it should be made as easy as possible for the subjects to understand the task presented in the experiment and to make the best decision for them. This is one of the reasons why experiments should be simple.
- Use neutral language. The main aim is to avoid experimenter demand effects. This means that the subjects should not be given the impression that the experiment serves a specific purpose.
- Provide an opportunity to learn. Even simple games should, however, be practiced by the subjects before the actual experiment takes place. It is quite possible that learning processes take place in the first few rounds of an experiment. If the experiment is not designed to observe these learning processes, then learning the game has no place in the actual experiment. The aim is to test whether the subjects who know and understand the game behave as predicted by the experimental hypothesis or not. Therefore, the learning process must take place before the experiment.
- A very controversial question is whether to conduct an experimental session only once or whether to provide the subjects with a repeated opportunity to experience the experimental situation – for instance, by carrying out the experiment again at intervals of 1 week.
- Repetition tends to increase external validity, since most of the decision problems that are explored in experiments are in fact not one-off occurrences in the real world, but recur at irregular intervals. The disadvantage of repeating the session is that the experimenter has no control over what happens between the repetitions. It may well be that the subjects in the meantime gain experience which has a strong influence on their behavior. The problem is that there is no way of knowing what experience this is. It is presumably this methodological problem that has hitherto prevented experimenters from repeating sessions to any large extent.

# The Size of Payoffs

- It's about money. But how much money we are actually talking about is a question we have sidestepped a little so far. Do monetary incentives have an effect and if so, to what extent does it depend on the size of the payoffs? This is a question that is frequently directed at experimental economic research. Two extreme positions are conceivable. One of them assumes that it makes no difference at all whether monetary incentives are used or not. The other extreme position is that all the deviations from the model of rational choice that can be observed in experiments disappear if the payoffs are set at a sufficiently high level.
- Camerer and Hogarth (1999) analysed 74 papers in which the effect of different payoff levels was investigated. The most important message from their study is that the two extreme positions described above are wrong. Monetary incentives are not ineffective (i.e. decisions should not be elicited hypothetically, but provided with appropriate monetary consequences) and deviations from the rational choice model do not disappear at higher payoffs. The last point can be expressed a little more precisely. The authors have not found a single study where a deviation from the rational choice model observed at low payoffs disappears when the payoff is increased.
- Nevertheless, the work of Camerer and Hogarth also shows that the effect of incentives is not always the same. It may well depend on the special circumstances of the experiment. For instance, an increase in payoffs has an impact if the payoff a subject receives at the end of the experiment depends on the effort involved. A good example is provided by experiments testing memory ability. Here it is profitable for the subjects to be more attentive and the more they can earn, the more attentive they actually are.
- The size of the payoff, on the other hand, has no influence in experiments in which the subjects already have a sufficiently high level of intrinsic motivation or in which any additional effort is not worthwhile because the payoff function is flat. Such experiments show, however, that the variance decreases, i.e. the average amount given remains the same, but there are fewer deviations upwards or downwards. The size of the payoff also has little influence on the behavior under risk. At best, there is a slight tendency towards more risk-averse behavior.
- The following rule of thumb should suffice: it is necessary to set noticeable but not exorbitantly high incentives. As a rule, the size of the payoff should be based on the opportunity costs of the subjects in the experiment.

# Is It Okay to Take Money from Subjects of Experiments?

- In economic contexts, but also in other important situations in society, it is possible that decisions taken by people result in losses. Sometimes it is even the case that people may only be able to exert an influence on how high a loss is and are no longer able to avoid it altogether.
- An important and interesting question is whether decision-making behavior in the event of losses mirrors that of gains or whether there are systematic differences. The only way forward experimentally is to conduct experiments in which subjects actually face the risk of loss or even have to accept a loss with certainty. In such a case, the experimenter takes money away from the subject. Is he allowed to do this? Should he do this?
- It is sometimes claimed that it is unethical to take money from subjects of experiments if they make losses in the laboratory. This poses a dilemma for the experimenters. On the one hand, it is important to find out how people react to possible losses. On the other hand, experiments must be designed in such a way that the subjects end up receiving money and not having to pay anything. A popular method of overcoming this dilemma is to design experiments in such a way that, although there is a possibility that losses may occur in individual parts of the experiment, on average there will be no loss at the end of the experiment.
- Incorporating losses without any actual losses being incurred is one way of avoiding the dilemma. However, such approaches are compromises since they only reflect real losses in a limited sense. Not paying for all the losses that result in an experiment can bias decisions. A frequently implemented alternative to this is to pay the subjects a sufficiently high “show-up fee”, from which the potential losses can be paid.



# The House Money Effect

- Monetary incentives are usually created in the experiment by, in a sense, pressing money into the hands of the subjects, who then can use it in the experiment. The basic idea here is that the value of money does not depend on where it comes from. Whether you work hard for 10 euros, find it on the street or win it in a lottery, it makes no difference to the quantity of goods you can buy for that money.
- So why should 10 euros received as a gift be worth less than the 10 euros earned? This view stems from rational choice theory. The notion that money always has the same value cannot be shaken, and under the requirements of the neoclassical rational choice model it would simply not be reasonable to value endowed money any differently from earned money.
- However, it obviously makes a difference whether the money used in an experiment is “your own money” or money provided by the experimenter. This endowed money is something like a windfall profit, i.e. income that simply lands in your pocket without you having to do anything about it.
- It is quite obvious that the unexpected gain changes your behavior – even if this is difficult to reconcile with rational behavior. It can be assumed that after a windfall profit, the propensity to consume increases just as much as the willingness to take risks. If that is the case, then of course this is highly relevant for the design of the payoffs in an experiment. The only question is how to experimentally test the effect of monetary gifts. How do we get people to use their own money in an experiment? Such an experiment is likely to make recruiting subjects quite difficult. For this reason, a different approach is taken.
- Instead of the subjects having to spend their own money, the experimenter has them perform a task for the money they receive. The type of task can be freely chosen. The crucial point is that the subjects no longer have the feeling that they have been given the money to use in the experiment. This is not quite the same as using their own, self-earned money, but if it turns out that in this sense money that is not endowed is treated differently from money that is, then it is safe to assume that it is the “house money effect” that describes this.

# Is It Permissible to Lie to Subjects of Experiments?

- Experiments in which the subjects can suffer losses are very rare and even rarer are experiments in which they actually have to pay something. Most seldom, however, are experiments in which subjects are lied to – at least in experimental economic research. At first sight, this seems to be self-evident, since lies seem to be at least as unethical as asking for money. In light of this, it should be clear that such a thing is simply not done.
- On closer inspection, however, it can be seen that honesty in the laboratory is a specialty of economists and that there are other disciplines which are far from being as strict about this as economics. It is therefore worth taking a look at why economists insist on honesty and why, for example, experimental psychologists often fail to do so.
- There is a very broad consensus within the scientific community of experimental economists that deception cannot be tolerated. According to the argument, lying would lead to the experimenters gaining a reputation of not being honest. This, in turn, would have disastrous consequences because, if the subjects were to suspect that they were being lied to in the laboratory, how would it possible to monitor their preferences?
- If the experimenter did not know which game the subjects thought they were actually playing, he could, strictly speaking, no longer draw any conclusions from their behavior. Such a scenario must be prevented and can only be achieved by the experimenters defending their reputation of being honest.
- An important question does remain to be resolved in this connection, however. When does dishonesty begin? The rule could be formulated as follows. Everything that is said to the subjects must be true. However, the whole truth does not always have to be told all at once.

# Are Students the Right Subjects?

- Well over 90% of all laboratory experiments are carried out with student subjects. And this not only applies to laboratory experiments in economics. Is this the right choice? This is a question which is constantly being raised with concern. Is it really possible to learn something about the behavior of people in general from the behavior of students? Or are students actually too “special”, i.e. not sufficiently representative?
- Students simply have many advantages. First of all, they are readily available, being represented in large numbers at universities and blessed with a relatively large amount of freely available time. That is why they can take part in an experiment, for example, at 2 in the afternoon or at 10 in the morning. Another advantage is that it can be assumed that students generally understand relatively easily and quickly what is expected of them in the experiment. From the point of view of the experimenters, it is also an advantage that students are often short of money and therefore gladly take the opportunity to earn something by participating in an experiment. The relatively low opportunity costs in terms of time mean that the monetary incentives set in the experiment do indeed carry a high weighting.
- These advantages are to a degree mirror images of the disadvantages of conducting experiments with non-student subjects. Recruiting the latter is much more difficult and time-consuming. If they are working people, only laboratory hours after work can be considered. In other words, it is necessary to get people to spend their scarce free time in the evening in the laboratory instead of at home with their families. In addition, it is difficult to establish initial contact. Students can be recruited relatively easily in lectures and are therefore usually represented in large numbers in a database of experimental subjects. Recruitment can be done at the push of a button or with a few clicks.
- This is not the case for non-students. Moreover, the opportunity costs are significantly higher for employed people than for students, with experiments with non-students, therefore, always being more expensive than those with students.

# Are Students the Right Subjects?

- However, inviting students into your laboratory may create a double selection bias. First, students differ systematically from the average population. On the one hand, they are younger and better educated; on the other hand, they do not have the experiences of an average adult. For instance, they have no professional experience, do not know what it is like to pay income tax or to negotiate for their salary. These systematic differences make it difficult to transfer the decisions observed by students to the average population.
- The second selection bias comes into play when students participate voluntarily in experiments. It cannot be ruled out that only certain types of students participate in experiments. What is particularly worrying is that this self-selection process affects the preferences of the subjects, both their risk preferences and their social, or other regarding, preferences.
- Are students different from non-students? Exadaktylos et al. (2013) conclude that there is no significant difference in social behavior between student volunteers and non-student volunteers. These findings admittedly contrast with a whole series of observations showing significant differences between students and non-students.
- For example, Falk et al. (2013) found that in a trust game, non-students paid back significantly higher amounts to the first movers than students did. In the experiment already mentioned by Anderson et al. (2013), it can also be seen that students behaved much more selfishly than the adult volunteers and the truck drivers. Cappelen et al. (2015) examine social behavior in a dictator game experiment and in a trust game experiment. They also find significantly more pronounced social preferences in a group of subjects consisting of representative persons (of Norwegian society) compared to a group of students. Belot et al. (2015) come to the same conclusion. They also find that students are more likely to be able to think strategically than “normal citizens”.
- However, all in all, students are not the worst possible choice. The differences to the rest of the population tend to be moderate, while the use of experts in experiments is not unproblematic. Therefore, using students as subjects may well represent a good alternative in the vast majority of cases. This does not rule out the possibility that there may be specific questions in which it seems advisable to conduct experiments with a more representative population.

# Cultural Differences

- The experimental method has long been used worldwide. The papers that result are published internationally and are recognized by the international scientific community, for whom national borders are not really important. But all this does not change the fact that the subjects of the experiments are almost always recruited within the region in which the laboratory is located. This raises the question of whether experimental results obtained in different countries can be easily compared. In other words, does it matter which country the subjects come from and what cultural background they have?
- Studies in which cultural differences are investigated usually have to contend with methodological problems resulting from the need to ensure that experiments in different countries differ only in terms of the cultural background of the subjects involved. This is not trivial because two experiments taking place in two countries can also differ in other dimensions, for example frequently in the language spoken and possibly also in the currency in which the money is paid out. The language difference may necessitate the use of different experimenters who speak the national language, thus leading to a further difference.
- Intercultural studies have made it clear that social norms can be assumed to develop very differently on a local level. It should be noted at this juncture that this is an important insight in two respects. First, it once again demonstrates that human decisions are also massively influenced by social norms – and not just by a rational, materialistic process of evaluation. Second, it again renders invalid the rational choice theory's claim to provide an explanation for all behavior.
- It appears that certain behavioral differences can often only be explained by resorting to local social norms. It is an important research question of the future for which type of decisions this is valid and for which the universal claim of the rational choice model can still be made. However, given all the cultural differences revealed by the experimental research conducted so far, it should be emphasized that these are mostly limited to differences in degree and only very rarely include qualitative differences. In other words, the patterns of human behavior demonstrated in laboratory experiments seem to be very similar across cultural borders.