Introduction to Behavioral economics

Lecture VIII - **Behavioral game theory**

Matej Lorko
matej.lorko@euba.sk
Student resources: www.lorko.sk

References: Camerer, C. F. (2011). *Behavioral game theory: Experiments in strategic interaction*. Princeton University Press.

How do people actually play?

- ROCK - PAPER - SCISSORS
- Each of two players simultaneously announces either Rock, or Paper, or Scissors.
    - Paper beats (wraps) rock
    - Rock beats (blunts) scissors
    - Scissors beats (cuts) paper
- The player who names the winning object receives $1 from her opponent
- If both players name the same choice then no payment is made
- Draw the payoff matrix and determine whether there is a pure strategy Nash equilibrium.
    - What are the best responses to each strategy?
    - How would you play this game?

- https://www.youtube.com/watch?v=rudzYPHuewc

Behavioral game theory

- In standard models it is typically assumed that players are highly rational beings who completely understand the strategic situation and who always maximize their consistent preferences given their rationally formed beliefs about the behavior of their opponents.

- At the opposite extreme, in evolutionary models, players have no cognition and therefore “no choice” but are “programmed strategies” that survive or go extinct in an evolutionary contest.•

- By contrast, the approach of behavioral game theory (BGT) is to seek empirical information about how human beings – as opposed to highly rational beings or programmed strategies – behave in strategic situations.

- Thus, BGT takes the middle ground between these two extremes but builds on the great advances of formal game theory, without which BGT would not exist. BGT aims to answer the following research questions:

    - To what extent is standard game theory a useful approximation to the strategic behavior of real people?

    - If we observe deviations from what standard theory predicts, can we disentangle the reasons for the discrepancies?

    - What are the players’ preferences and their strategic reasoning processes?

    - How do people learn in games?

Behavioral game theory

- Coordination games: Experiments show that, after some initial miscoordination, play converges to an equilibrium. Yet, unless the players can communicate, they almost invariably end up playing the risk-dominant instead of the payoff-dominant equilibrium.

- Cooperation games: The striking result is that people cooperate much more than is compatible with a simple dominance argument that underlies the prediction in the prisoners' dilemma (if we assume that players maximize only their monetary payoffs).

- An important contribution of BGT in recent years concerns our understanding of human players' actual social preferences, i.e., to what extent people take the well-being of other players into account in their preferences.

- The results from the cooperation games suggest that many players are not purely selfish.

- Yet, simultaneous-move games of cooperation are rather blunt tools to measure social preferences because it is very hard to distinguish altruism, reciprocity, and selfishness.

- Therefore, games in so-called "extensive form," where players move sequentially, are more apt to measure social preferences than simultaneous-move games.

Ultimatum game

- Bargaining process - “This is my best offer, take it or leave it …”. If that ultimatum offer is accepted then it leads to a resolution, but if not, then it sometimes means substantial financial losses for both parties involved.

- 2 players (Proposer and Responder)

- 2 stages

- Stage 1: Player 1 proposes a division of a fixed pie (say 100 dollars) between the two players.
- Stage 2: Player 2 either:
    - Accepts the division (money are divided accordingly and the game ends).
    - Rejects the division (the game ends with 0 payoff for both players).

- You are Player 1. What would you do?

Ultimatum game

- Typical results: Player 1 offers between 30-50% (almost never more than 50% or less than 10%). Offers of 40-50% are rarely rejected, offers below 20% are rejected half of the time. What drives the behaviour of Player 1 and Player 2?

- Most researchers today agree that rejecting a positive offer in the ultimatum game indicates negative reciprocity (eye for eye, tooth for tooth). A person has negatively reciprocal preferences, if she is willing to pay some price to punish an opponent for behavior that is deemed unfair or inappropriate. The observation of negative reciprocity is not confined to ultimatum games. It has also been observed in social dilemma and public goods games where players had the opportunity to punish their opponents.

- Many cooperators were willing to incur costs to punish the defectors, even in one-shot games without any future interaction. Rejecting a positive offer in a one-shot ultimatum game or punishing defectors means to forgo money without any material benefit. Many people have a willingness to punish even in the absence of any present or future rewards.

- The friendly version of reciprocity is called positive reciprocity (nice to me, nice to you). Positive reciprocity means that people are prepared to pay a price to reward a friendly or a generous action by an opponent player. They are willing to pay this price even in the absence of any present and future material benefits.Thus, a purely self-interested individual would never exhibit positive reciprocity. And yet, positive reciprocity is quite common.

- Casual evidence and daily experience suggest that not only outcomes but also the “intentions” (the attribution of motivations) behind a decision matter for our evaluation of outcomes. People display apparent willingness to pay to achieve fairness or to punish unfair behavior.

Dictator game

- Possible explanations of the observed behavior in the ultimatum game: Player 1 may propose a positive amount for player 2 because of (1) altruistic other-regarding preferences or (2) fear that player 2 might reject a "selfish" proposal. To test for quantitative effects of altruistic other-regarding preferences and fear of rejection of proposals one can use a dictator control treatment.

- Again 2 players, first one divides a pie. However, this time, there is no Stage 2. Player 2 has no move, the game ends after the division by Player 1. You are Player 1. What would you do?

- Typical results: Player 1 in DG usually offers less than Player 1 in UG. However, many of then still offer a substantial amount (10-30%).

- Why? Social/other-regarding preferences.

    - Altruism – sacrifices in order to improve another's situation (prefers say 70-30)

    - Envious people, egoists – prefer 100-0

    - People that like fairness, or are inequality averse – 50/50

Altruism

- There is a lot of evidence that many people are prepared to make anonymous donations to charities or to spontaneously help others who are in need.
- A person has altruistic preferences if her utility increases with the well-being of others.
- The experimental tool to study this is the “dictator game”.
- A player (the “dictator”) is endowed with some money, say $10, and can then decide how much to pocket, and how much to pass on to a passive recipient, who cannot veto the offer. Of course, under standard assumptions, the dictator will keep everything.
- Under double-blind conditions, roughly two-thirds of the people give nothing and one third gives amounts between 10 and 50 percent of the pie.
- Offers are significantly lower than in the ultimatum game, because the dictators do not have to fear rejections.
- The significance of the results from the dictator games is that many people, even under complete anonymity, are willing to share their wealth with others.

Fairness

- Much of the behavior in Ultimatum/Dictator games can be explained in terms of people having a preference for what is fair or what is just. Yet the empirical nature of this preference differs in important ways from philosophical descriptions of fairness and justice.

- Fair choice is largely about tradeoffs. The choices we observe in both the ultimatum bargaining game and the dictator game make this quite evident. Models of this behavior assume that an equal split is what people think of as fair, yet many people demonstrate a willingness to strike a compromise between fairness and self-interest whether they are asked to take a smaller than 50–50 share or are deciding how much another should get. The heterogeneity and behavior combined with a tendency to believe others see things as you do might help explain why there are often arguments about what is fair in the first place.

- Fair choice is asymmetrically self-centered. The tendency to resist what is deemed unfair to one's self is, on aggregate, stronger than the willingness to sacrifice self-interest to treat others fairly. An important implication is that those most likely to sacrifice their own self-interest to punish an unfair distribution are those whose own relative standing would be most diminished by it.

- Fair choice is strategic choice. The influence of fair choice seemingly vanishes in competitive markets. In competitive markets, strategic considerations compel the fair minded to behave as if they are self-interested. Similarly, the influence of fair choice tends to be mitigated, in strategic ways, when information about payoffs and strategic options becomes incomplete or less transparent.

- Fair choice is predictable choice.This robustness, in turn, allows social choice research to contribute to the engineering of better incentives and institutions.

- Fair choice is a trigger of reciprocity and trustworthiness. Reciprocity can be thought of as gift exchange; one gives a gift in the hope a gift will be given in return. Trusting means making your own well-being vulnerable to the action of another. In both cases, we intuitively understand that what is expected in return is commensurate with the size of the gift or level of trust that has been invested in us.

- This is not to say that fairness is the only element important to human sentiment toward others. But it does seem to be a robust factor behind a lot of other-regarding behavior of concern to economics and business behavior.

- Check this! https://www.youtube.com/watch?v=meiU6TxysCg

Hold-up problem

- The hold-up problem (or commitment problem) is central to the theory of incomplete contracts, and shows the difficulty in writing complete contracts. A hold-up problem arises when two factors are present:
    - Parties to a future transaction must make noncontractible relationship-specific investments before the transaction takes place.
    - The specific form of the optimal transaction (such as quality-level specifications, time of delivery, what quantity of units) cannot be determined with certainty beforehand.
- The hold-up problem is a situation where two parties may be able to work most efficiently by cooperating but refrain from doing so because of concerns that they may give the other party increased bargaining power and thus reduce their own profits. When party A has made a prior commitment to a relationship with party B, the latter can 'hold up' the former for the value of that commitment. The hold-up problem leads to severe economic cost and might also lead to underinvestment.
- But in reality, we observe much less of a hold-up problem than we would expect from theory. Why?

Trust

- 2 players - Sender and Receiver. Both are initially given 10 dollars.Sender makes an investment – sends x of his 10 dollars to the Receiver. The investment is multiplied by 3 The Receiver can send some amount back.

- Receiver is in effect playing a dictator game – however, with a previous stage in which he could be given some trust. Sender usually sends around a half of his money. Receiver usually sends approximately the same amount back.

- Why senders actually make investments? Why receivers send something back?

- People like to be trusting and like to be trustworthy. Many day-to-day transactions in life require us to trust strangers, e.g., lawyer or accountant or car mechanic, buyer seler amazon. Trust is often important in reducing the costs of transacting deals (e.g., no formal contracts, enforcements needed).

- Neither – trust or reciprocity – could support cooperation without the other. Those who trust naïvely, without any calculation of expected reciprocity, would be easily exploited. On the other hand, those who engage in calculated and strategic trust without any tendency to reciprocate others' trust would be too opportunistic and it is unlikely that they will be trusted too often.

- Individuals in higher-trust societies spend less to protect themselves from being exploited in economic transactions. Written contracts are less likely to be needed and litigation may be less frequent. Individuals in high-trust societies are also likely to divert fewer resources to protecting themselves – through arbitrary tax payments, bribes, or private security services and equipment – from unlawful or criminal violations of their property.

- Low trust can also discourage investments and innovation. If entrepreneurs must devote more time to monitoring possible malfeasance, they have less time to devote to innovations in new products and processes. Societies characterized by high levels of trust are also less dependent on formal institutions to enforce agreements.

Social preferences

- People seem to live in two worlds. The market world and the social world. While the market world is often driven by rational utility maximization, the social world is often driven by social/other-regarding preferences. These preferences are often evolutionary. People care about equality, they behave reciprocally (both positively and negatively) and are often display trust in others.

- Social norms are learned behaviors, and we usually learn them by observing the behavior of others. Our social behavior seems to be governed by considerations of fairness, cooperation and the ‘warm glow’ we feel when we help other people.

- Social norms are useful because:

    - They help to reduce uncertainty about how to behave appropriately – just follow the example set by other people!

    - They help to coordinate the behavior of individuals, which reduces ‘cooperation losses’ for other group members

    - They facilitate group cohesion.

    - In other words, social norms promote social efficiency.

Public goods game

- N players
- Each player given 10 dollars
- One stage, players move simultaneously
- Each player has the option to transfer any fraction of his endowment to a public account
- Money in the public account is multiplied by 2 and split evenly between the players
-

Public goods games

- In the public goods game subjects secretly choose how many of their private tokens to put into a public pot. The tokens in this pot are multiplied by a factor (greater than one and less than the number of players, N) and this "public good" payoff is evenly divided among players. Each subject also keeps the tokens they do not contribute.

- The nature of the experiment is incentives and the problem of free riding. Public goods games investigate the incentives of individuals who free-ride off of individuals who are contributing to the common pool.

- The group's total payoff is maximized when everyone contributes all of their tokens to the public pool. However, the Nash equilibrium in this game is simply zero contributions by all; if the experiment were a purely analytical exercise in game theory it would resolve to zero contributions because any rational agent does best contributing zero, regardless of whatever anyone else does.

- While in experiments, we usually don't see all players playing Nash equillibrium (zero contribution), usually the average contribution falls with more rounds. It is because people who begin as coordinators usually tend to give up after some time and start free-riding too.

- What seems to help is the availability of altruistic punishment, which can turn free riders into conditional cooperators. However, if there are not enough contributors willing to punish (and incurr own costs), the coordination is basically impossible.

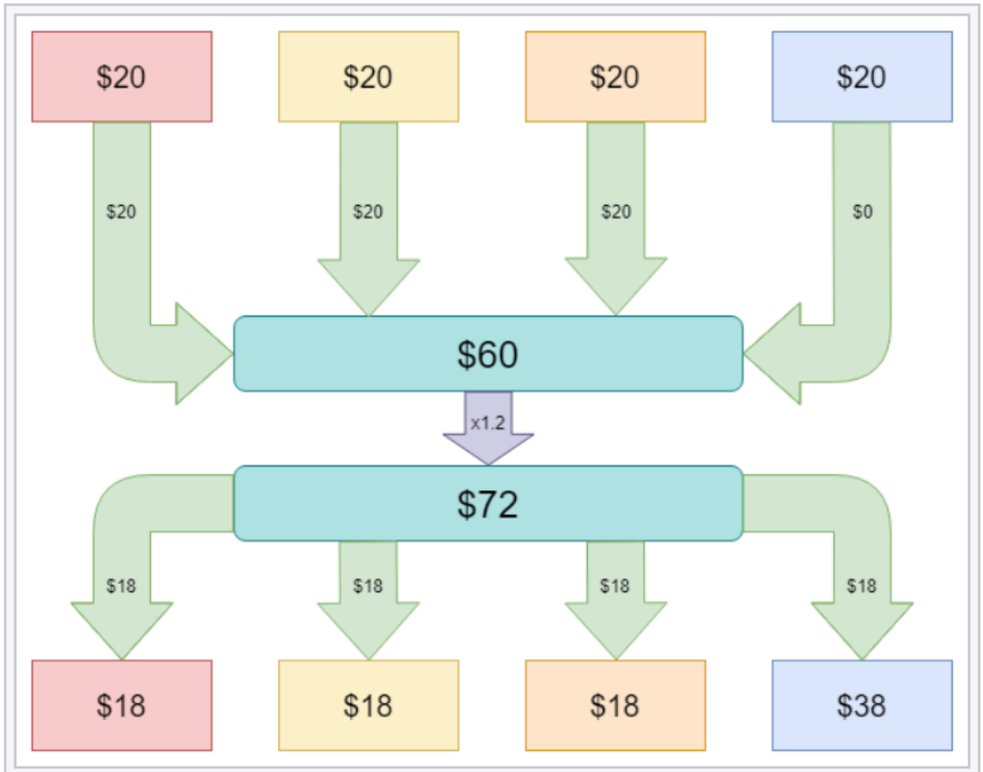

In this diagram of a public goods game, three players choose to contribute their full $20 while the fourth player chooses to contribute $0. The $60 is multiplied by a factor of 1.2 and the resulting $72 is distributed equally among the four players.

Public goods games

- An idea of a threshold public good game is that a good will be provided for the benefit of everyone if and only if people contribute enough to exceed some threshold. For example, if a church roof needs to be replaced, and doing so costs $100,000, then members of the church somehow need to raise the $100,000 threshold.

- Or, if flatmates have to clean a flat to the standard their landlord requires, between them they need to put in time cleaning the flat to the threshold standard. In this game there is no trade-off between risk dominance and Pareto dominance. The best outcome is that they contribute enough to exceed the threshold. The problem now is a conflict of interest over how much each should contribute. For instance, one flatmate may do little to clean the flat, in the hope that another flatmate will do a lot.

- The main question is whether people can coordinate by contributing enough, despite the conflict of interest. Yet again, we see that people are not great at coordinating. Yet again, we need to ask what might help people coordinate better. One thing that might matter is the institution in place to collect contributions.

- For example, if contributions fall short of the threshold then we might be able to give people a refund on their contribution. Sometimes this is not possible; for instance, flatmates cannot get back the time they have spent cleaning. But sometimes it is possible; for instance, the church could give back donations if insufficient funds are raised. Another possibility is to give a rebate if contributions exceed the threshold.

Cooperation in the real world

- The crux of the prisoners' dilemma and public good games problem is that all players can be better off if they cooperate, but individual rationality and the desire to maximize one's own pay-off dictates free-riding on the cooperation of others, which is the dominant strategy. When they both rely on their dominant strategies, they are collectively worse off. There is, thus, a tension between cooperating and maximizing the joint benefit, or free-riding and trying to maximize one's own pay-off at the expense of others.

- Collectively, we are better off if we cooperate, but the cooperative outcome is often hard to sustain, since, if everyone is cooperating, then one person can be better off by reneging and free-riding. But if it makes sense for one person to free-ride, then it does so for others as well, so, the equilibrium is that we all free-ride and we end up with global warming, fast depleting oceans and forests, and dirty streets.

- And once we arrive at that bad outcome, we might regret it, but we are often unable or unwilling to change the situation, because we would need everyone to change at the same time. One person choosing to cooperate while everyone else free-rides does not change things and makes the one co-operator worse off. But getting everyone to change their minds at the same time poses similar problems of collective action which led to the Nash equilibrium in the first place.

- However, if players know that they will interact over and over again, or that they can make binding commitments that can be enforced by a third party, then the outcome might be different.

Cooperation in the real world

- People are neither purely self-interested nor purely altruistic, but, rather, they are conditional co-operators whose behaviour is determined to a large extent by what they think their peers will do. So, if there is any doubt in the mind of one person that the other person might not cooperate and might go off on his/her own, then the secure option might be to go on your own in the first place.

- What seems absolutely crucial to successful cooperation is the creation of optimistic beliefs about the actions of our peers. More importantly, a majority of people are willing to cooperate as long as enough others do; they just need to be made aware of the fact that there are others like them. This seems to be the key to generating the requisite optimistic beliefs that can lead to successful collective action.

- The same holds for business environment. Financial incentives are important but simply raising incentives is poor managerial strategy; it is essential to reinforce the financial incentives with messages providing the insight that everyone is better off when everyone works harder.

- What also works well is mutual monitoring among the workers. But to sustain long-term coordination it is helpful to remind yourself of an availability bias. It is good to remember, that you will occasionally do more than your share, but it is useful to know that you are likely to have that feeling even when each member of the team feels the same way.

What we learned from experiments?

- Many people are willing to sacrifice their own monetary payoff to increase that of others (dictator game, trust game, public good game).
- Many people prefer fair outcomes and want to be trustworthy. They also want to trust.
- Many people reciprocate the kind action of another to them by kindness of their own (trust game, gift exchange game, moonlighting game). This is positive reciprocity.
- Many people reciprocate the unkind action of another to them with punishment (linear public good game, moonlighting game). This is negative reciprocity.
- Many people show both positive and negative reciprocity (moonlighting game, linear public good game). This is strong reciprocity.
- Many people reciprocate the kind or unkind action of another to someone other than them (dictator game with third-party punishment). This is indirect reciprocity.
- There is considerable heterogeneity in desires for giving and reciprocity (all experiments), with many giving zero if they have the chance