Behavioral economics

Lecture 8 - **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?

Coordination games

- Coordination games

  - Examples - Stag hunt game, Minimum effort game

  - Interesting features: multiple equillibria

    - one of which is typically a payoff-dominant equllibrium (Pareto efficient)

    - another is a risk-dominant (risk-free) equillibrium with lower payoffs

    - classical game theory can't predict which equllibrium will be played

  - Playing a strategy potentially leading to payoff-dominant equllibrium requires trusting the other players to also play that strategy

  - Insights from experiments featuring human subjects

    - after some initial miscoordination, play converges to an equilibrium

    - If the players can't communicate, they almost invariably end up playing the risk-dominant equllibritum

    - If the players can communicat before the game, they usually end up in a payoff-dominant equilibrium.

Cooperation games

- Cooperation games

  - Examples: prisoner's dilemma, public goods game

  - Interesting features

    - socially effective (Pareto efficient) outcome comes from cooperation

    - however, every player has individual incentives (dominant strategy) not to cooperate

    - Thus, classical game theory predicts no cooperation (Nash equllibrium)

  - Insights from experiments featuring human subjects

    - 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).
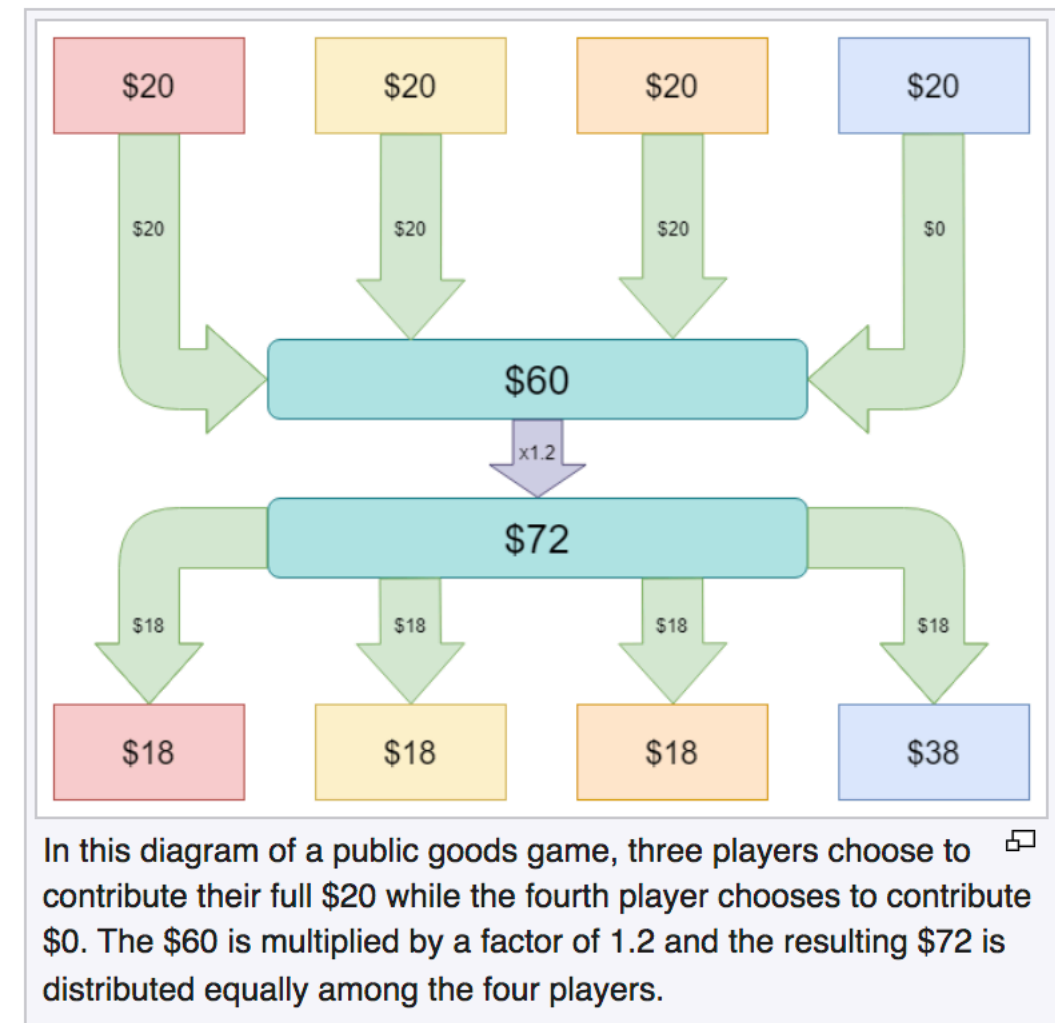
Axelrod's prisoner's dilemma tournament

- Robert Axelrod reports on a tournament that he organized of the N-step prisoner's dilemma (with N fixed) in which participants have to choose their strategy repeatedly and remember their previous encounters. Axelrod invited academic colleagues from around the world to devise computer strategies to compete in an iterated prisoner's dilemma tournament. The programs that were entered varied widely in algorithmic complexity, initial hostility, capacity for forgiveness, and so forth.

- Axelrod discovered that when these encounters were repeated over a long period of time with many players, each with different strategies, greedy strategies tended to do very poorly in the long run while more altruistic strategies did better, as judged purely by self-interest. He used this to show a possible mechanism for the evolution of altruistic behavior from mechanisms that are initially purely selfish, by natural selection.

- The winning deterministic strategy was tit for tat. The strategy is simply to cooperate on the first iteration of the game; after that, the player does what his or her opponent did on the previous move. Depending on the situation, a slightly better strategy can be "tit for tat with forgiveness": when the opponent defects, on the next move, the player sometimes cooperates anyway, with a small probability (around 1–5%, depending on the lineup of opponents). This allows for occasional recovery from getting trapped in a cycle of defections.

- After analyzing the top-scoring strategies, Axelrod stated several conditions necessary for a strategy to succeed:

  - Nice: The strategy will not be the first to defect, i.e., it will not "cheat" on its opponent for purely self-interested reasons first. Almost all the top-scoring strategies were nice.

  - Retaliating: The strategy must sometimes retaliate. An example of a non-retaliating strategy is Always Cooperate, a very bad choice that will frequently be exploited by "nasty" strategies.

  - Forgiving: Successful strategies must be forgiving. Though players will retaliate, they will cooperate again if the opponent does not continue to defect. This can stop long runs of revenge and counter-revenge, maximizing points.

  - Non-envious: The strategy must not strive to score more than the opponent.

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.

Behavioral game theory

- 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.

Labor markets

- Why do we observe involuntary unemployment in the world?

- Economic theory says that if the supply exceed the demand, then the price should fall and new equllibrium should emerge. However, in labor markets, we do not observe price (wage) falling too often, which is why oversupply of work (involuntary unemployment) arises.

- The main reason of employers to not decrease wages is that they fear that most productive workers would leave (an adverse selection argument), but a close second is that worker effort would decrease (a reciprocity argument). Along the lines of loss aversion, a pay rise is predicted to have some benefit, but it is clear that a wage cut is expected to have a significant negative effect. If employers are reluctant to cut wages because it will cause negative reciprocity then we have a plausible reason for wage stickiness.

- In laboratory, these predictions are usually tested using a gift-exchange game. The gift exchange model is used to explain workers' effort and wages provided by firms in the real world. George A. Akerlof described labor contracts as "partial gift exchange". Employees may exceed the minimum work required and firms may pay more than the market-clearing wage.

- According to Akerlof's model, this is because the worker's effort to some degree depends on the norm for effort. Thus, to affect these norms, firms may pay more. A worker may be willing to work hard if he believes the employer is being kind and offering a higher-than-equitable wage.

Gift exchange
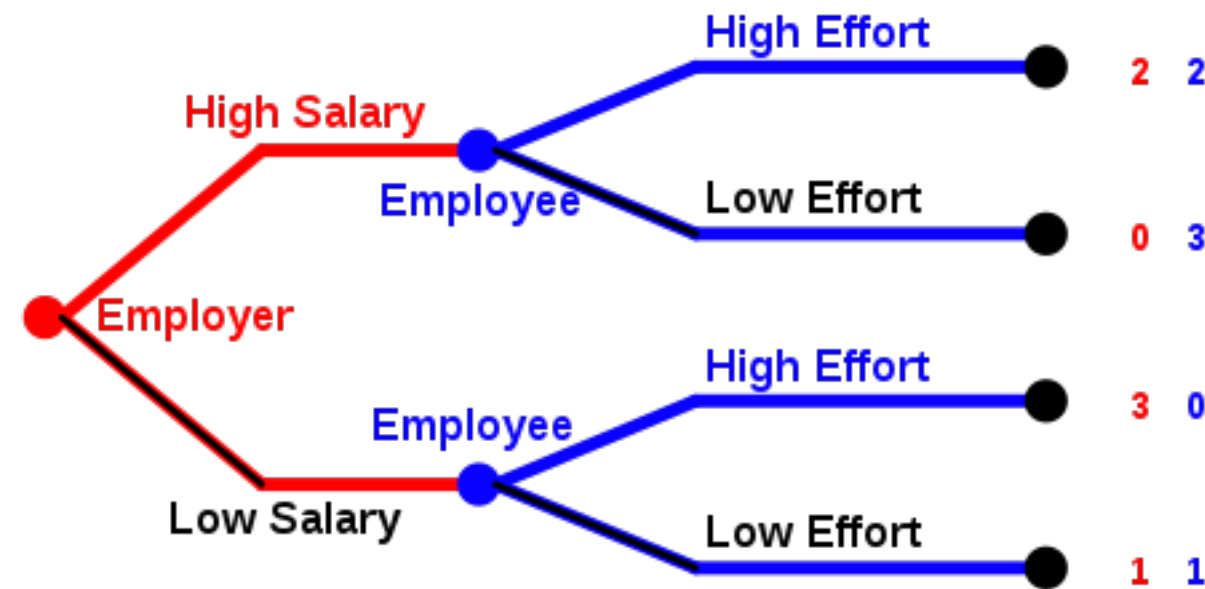
- Two players are at least involved in such game – an employee and an employer. The employer has to decide first, whether to award a low or higher salary. Having observed her wage, the employee chooses how much effort to put into working. Higher effort costs the worker, but also means the employer receives higher revenues.

- Cost and revenue are usually designed in such a way that higher worker effort is mutually beneficial – i.e. the extra effort would earn enough extra revenue that the employer could pay a high enough wage to offset the worker's effort cost. A worker motivated solely by her own monetary payoff would not put in effort, however, because it is too late to do anything about her wage, and so all it would do is lower her payoff. Given this, why should an employer pay a high wage? He should not.

- Like in trust games, game-theoretic solution for rational players predicts that employees' effort will be minimum for one-shot and finitely repeated interactions. Therefore, there is no incentive for the employer to pay a higher salary. If the employer pays a higher salary, it is irrational for the employee to put extra effort, since effort will reduce his or her payoff. It is also irrational for the employee to put extra effort while receiving a lower salary. Therefore, the minimum salary and the minimum effort is the equilibrium of this game.

- The payoff matrix of the gift-exchange game has the same structure as the payoff matrix of Prisoner's dilemma. The difference constitutes by the sequentiality of gift-exchange game.

- Many experiments observe that workers in the gift exchange game provide substantially more efforts than the minimum required. The workers' choices appear to reciprocate the firm, in the sense that higher wages tend to lead to higher effort, which are against the standard assumption of strictly self-interested behavior.

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.

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 (such as in the gift-exchange game). 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.

Positive reciprocity

- Positive reciprocity refers to the tendency for people to respond to kind and generous actions from others with similar kindness and generosity. It involves the exchange of positive actions or favors, where one person initiates an action that benefits another person, and that person responds with a similar action that benefits the first person in return.

- Positive reciprocity is generally considered to be good for several reasons:

  - Strengthens relationships: By engaging in positive reciprocity, people can build stronger relationships with others, as they feel a sense of trust and mutual benefit from the exchange of kind and generous actions.

  - Increases cooperation: Positive reciprocity can also promote cooperation among individuals or groups, as each person is motivated to continue behaving positively towards others in the hopes of receiving similar treatment in return.

  - Encourages prosocial behavior: When people experience positive reciprocity, they may feel more motivated to engage in prosocial behavior, such as helping others or volunteering, because they see the positive impact it can have on themselves and others.

  - Enhances well-being: Acts of kindness and generosity can promote positive emotions, such as gratitude and happiness, which can enhance overall well-being and life satisfaction.

Negative reciprocity

- Negative reciprocity refers to the tendency for people to respond to negative actions from others with similar negative actions.

- Negative reciprocity can occur for various reasons, including:

  - Retaliation: People may engage in negative reciprocity as a way to retaliate against others who have harmed them or violated their interests.

  - Self-defense: Negative reciprocity can also be a form of self-defense, where individuals respond to harmful actions from others in order to protect themselves or their resources.

  - Justice: Negative reciprocity can be seen as a way to restore justice or fairness, particularly in situations where individuals feel that they have been wronged or treated unfairly.

  - Competition: In competitive environments where resources or power are limited, negative reciprocity may be used as a way to gain an advantage over others or maintain social status.

  - Cultural norms: In some cultures, negative reciprocity may be seen as an appropriate response to certain situations, such as insults or challenges to one's honor or reputation.

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%).

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.

Altruism vs. warm glow

- Altruism and warm glow are both related to prosocial behavior, but they are different concepts.

- Altruism is the act of helping or doing something positive for others without expecting any personal benefit or reward in return. Altruistic behavior is motivated by concern for the welfare of others, rather than personal gain. Altruism often involves selfless acts, where the individual puts the needs of others before their own.

- On the other hand, warm glow is a positive feeling that individuals experience when they engage in prosocial behavior. Warm glow refers to the internal satisfaction or sense of happiness that individuals feel when they engage in actions that benefit others, even if there is no external reward or recognition for their actions. This feeling of warmth and happiness can serve as a motivation for individuals to engage in further prosocial behavior.

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 game

- 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.

Trust

- Trust is a belief or confidence that one person or entity will act in a reliable, competent, and ethical manner. It involves a willingness to be vulnerable and rely on another person or entity to act in one's best interests, based on the belief that they will do so.

- Trust is a fundamental aspect of human relationships and social interactions, as it enables individuals to cooperate, work together, and build relationships based on mutual respect and understanding. Trust can be built through various means, such as consistent and reliable behavior, honesty and transparency, and shared values and beliefs. Conversely, trust can be broken through actions that violate these expectations, such as dishonesty, betrayal, or a failure to fulfill commitments.

- Trust is good because it is essential for building strong, positive relationships and facilitating effective communication and cooperation between individuals or groups. When people trust each other, they are more likely to share information, work together towards common goals, and support each other in times of need.

- Trust can also help to create a sense of safety and security, as individuals can rely on each other to act in their best interests and avoid actions that could cause harm or damage. Trust is also important in business and other professional settings, as it can help to establish credibility and foster positive reputations for individuals or organizations. Overall, trust is a fundamental aspect of human social interactions and can lead to many positive outcomes, such as improved well-being, greater productivity, and stronger relationships.

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