Introduction to Behavioral economics

Lecture IV - **Strategic interactions**

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.

Strategic interactions

- A large number of decisions in our day-to-day lives require us to engage in “strategic decision making”.

- It means that what I decide to do in a particular situation will affect the well-being of another person (or group of people) – and, in turn, what someone else does will crucially impact upon my own well-being.

- Same holds true for firms and their partners or competitors.

- Theses situations can be thought of as “games” with us as “players”, and they can be analysed using the tools of game theory. This interdependence causes each player to consider the other player’s possible decisions, or strategies, in formulating his own strategy. Players may have similar, opposed, or mixed interests. A solution to a game describes the optimal decisions of the players.

Game theory

- Developed by John von Neumann and Oskar Morgenstern (1944)
- Game theory is a set of tools used to help analyze situations where an individual's best course of action depends on what others do or are expected to do. Game theory allows us to understand how people act in situations where they are interconnected.
- Connections between people arise in all sorts of situations. Sometimes through cooperation with others, we can achieve more than we can on our own. Other times, conflict arises where an individual benefits at the expense of others. And in many situations, there are benefits to cooperation but elements of conflict also exist.
- In stressing the strategic aspects of decision making, or aspects controlled by the players rather than by pure chance, the theory both supplements and goes beyond the classical theory of probability.
- Because game therory can help analyze any environment where the person's best action depends on others' behavior, it has proven useful to analyze strategic interactions in wide variety of fields.
- Examples:
    - In economics, the decisions of firms are affected by their expectations of a competitor's choice of product, price and advetising.
    - In political science, a cadidate's policy reform is influences by policy announcement of their rival.
    - In biology, animals must compete for scarce resources, but can be hurt if they are too aggressive with thewrong rival.
    - In computer science, networked computers compete for bandwidth.
    - In sociology, public displays of non-conformist attitudes are influeneced by other's behavior, which is shaped by social culture.
    - In sports, it is often necessary to predict the behavior of others (e.g., footbal penalty kicks, tennis serves…)

Types of games

- Number of players (player need not be an individual)
    - one-person - games against nature, no opponents, the player only needs to list available options and then choose the optimal outcome.
    - two-person
    - n-person (with n greater than two)
- Information
    - perfect - each player knows everything about the game at all times (chess)
    - imperfect - players do not know all of their opponents' possibilities (poker)
- The extent to which the goals of the players coincide or conflict
    - Constant-sum games are games of total conflict (pure competition), players have completely opposed interests
    - Variable-sum games - players may all be winners or losers (labour-management dispute)
        - cooperative (players can communicate and, most important, make binding agreements)
        - noncooperative (players may communicate, but they cannot make binding agreements)
- Number of options
    - Finite - each player has a finite number of options, the number of players is finite, and the game cannot go on indefinitely
    - Infinite

Constant vs. Variable sum games

- Constant-sum games are win-lose games. Whatever you win in those situations, the other party has to lose. The players in constant-sum games games have diametrically opposed interests, and there is a consensus about what constitutes a solution (=everybody plays their dominant strategy). Theoretically, the outcome of the game (the Nash equillibrium) is predictable.

- Most games that arise in practice, however, are variable-sum games; the players have both common and opposed interests. For example, a buyer and a seller are engaged in a variable-sum game (the buyer wants a low price and the seller a high one, but both want to make a deal).

- The effect of communication is particularly revealing of the difference between constant-sum and variable-sum games.

- Communication is pointless in constant-sum games because there is no possibility of mutual gain from cooperating.

- In variable-sum games, on the other hand, the ability to communicate, the degree of communication, and even the order in which players communicate can have a profound influence on the outcome.

- Generally, the more two players' interests coincide, the more important and advantageous communication becomes.

Variables

- Players: Who is interacting?
- Strategies: What are the options of each player? In what order do players act?
- Payoffs: How do strategies translate into outcomes? What are players' preferences over possible outcomes?
- Information/Beliefs: What do players know/believe about the situation and about one another? What actions do they observe before making decisions?
- Rationality: How do players think?
- A solution concept - "Nash equilibrium" = a profile of strategies where each player's strategy is a "best response" to the strategies of others (i.e. gives him the highest payoff among his strategies, given the others' strategies). Note that the Nash equilibrium requires players to have correct beliefs about the strategies of others. Nash equilibria can be in "pure strategies" (each player chooses one strategy with certainty) or in "mixed strategies" (players choose randomly among a set of strategies).

Saddlepoints

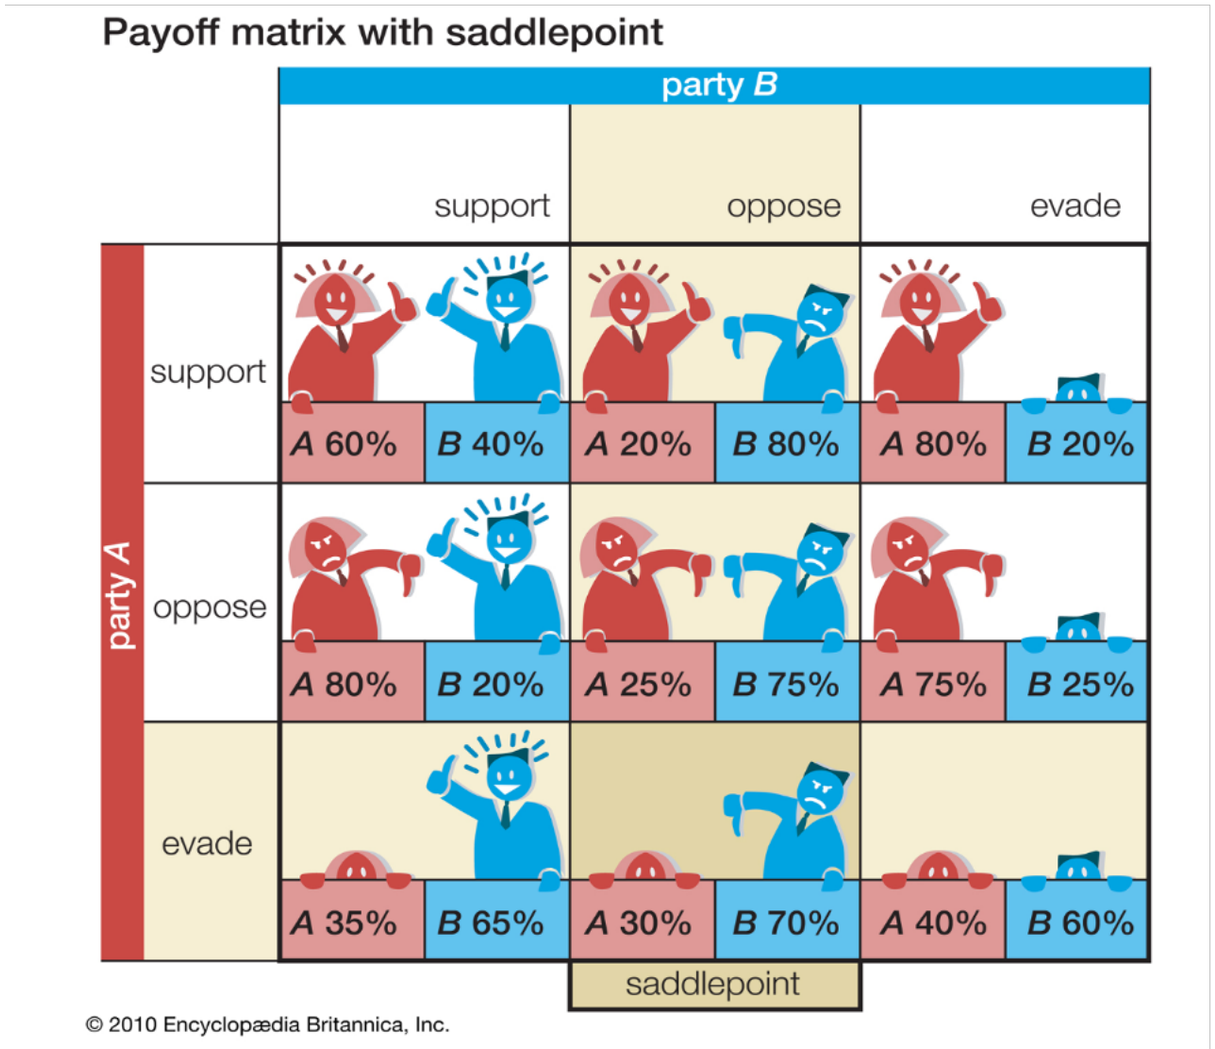

- A “saddlepoint” is the outcome that rational players would choose.
- A saddlepoint always exists in games of perfect information but may or may not exist in games of imperfect information.
- By choosing a strategy associated with this outcome, each player obtains an amount at least equal to his payoff at that outcome, no matter what the other player does.
- This payoff is called the value of the game; as in perfect-information games, it is determined by the players’ choices of strategies associated with the saddlepoint, making such games strictly determined.

Stag Hunt (coordination)

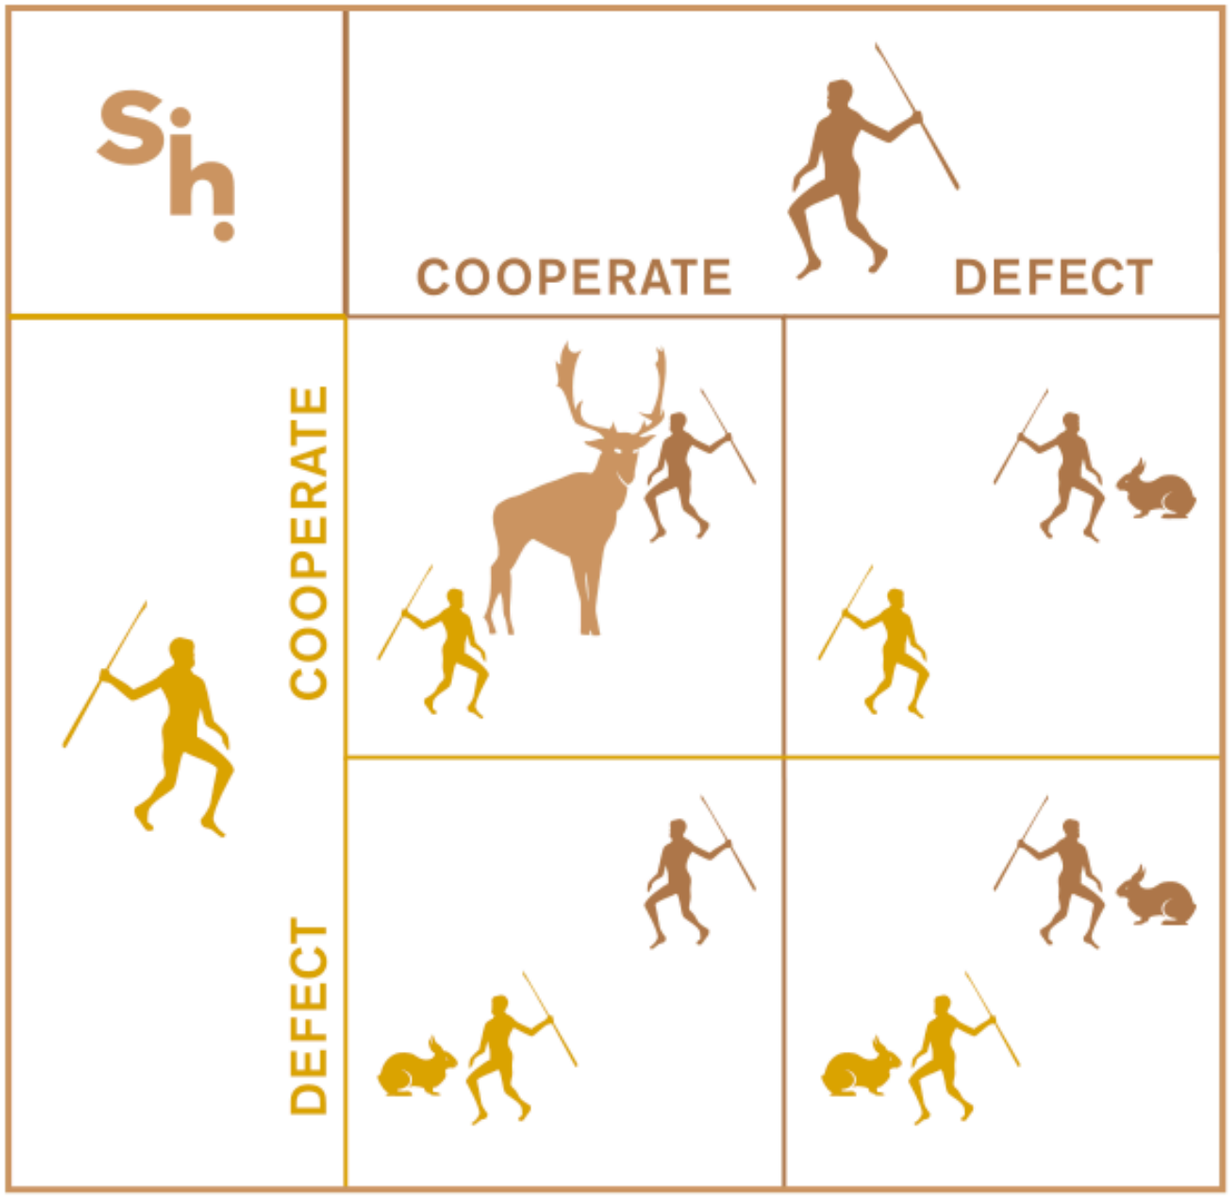

- The stag hunt, sometimes referred to as the assurance game, trust dilemma or common interest game, describes a conflict between safety and social cooperation.

- Two hunters must decide separately, and without the other knowing, whether to hunt a stag or a rabbit. However, both hunters know the only way to successfully hunt a stag is with the other's help. One hunter can catch a rabbit alone with less effort and less time, but it is worth far less than a stag and has much less meat.

- Therefore, it would be much better for each hunter, acting individually, to give up total autonomy which brings only the small reward of the rabbit. Instead, each hunter should separately choose the more ambitious and far more rewarding goal of getting the stag, thereby giving up some autonomy in exchange for the other hunter's cooperation and added might.

- Commentators have seen the situation as a useful analogy for many kinds of social cooperation, such as international agreements on climate change.

Stag hunt game

	Stag	Rabbit
Stag	(2, 2)	(0, 1)
Rabbit	(1, 0)	(1, 1)

Minimum effort game

Own action	Minimum action in group						
	7	6	5	4	3	2	1
7	130	110	900	70	50	30	10
6		120	100	80	60	40	20
5			110	90	70	50	30
4				100	80	60	40
3					90	70	50
2						80	60
1							70

Figure 7.2 Example stag hunt and minimum effort coordination games.

- From the past experiments, it seems that in a basic stag hunt game, the risk dominant rabbit equilibrium is chosen overwhelmingly. However, subjects are attracted to the payoff dominant stag equilibrium and will coordinate on it if they are able to communicate before the game.

- Coordination of the form captured by the stag hunt and minimum effort games (also known as weakest link games, which is basically a stag hunt with many participants) is also important in organizational settings, where the disparate parts of a production or service system all need to complete their tasks in a timely fashion. Similarly, a project with multiple critical paths can be delayed by any one path going over schedule. For example, an airplane can only take off if all the preparations are finished: passengers boarded, flight crew prepped, luggage loaded, etc. Any work group falling behind delays the whole flight. A natural question is then what kind of institutional features best support coordination, and what kind of managerial interventions can lead to coordination improvements.

- Large coordinating groups can emerge by taking small groups that are coordinating well and slowing growing them. Other managerial and institutional interventions can also improve coordination. Requiring players to pay an up-front fee increases coordination on higher actions. They argue this behavior is consistent with forward induction and loss aversion: Subjects are reluctant to enter the game intending to play an equilibrium that would yield negative payoffs. Brandts and Cooper (2007) add a subject in a managerial role that can both communicate with the coordinating employees and give them financial incentives. While financial incentives are helpful, communication is even more effective. Indeed, the most effective communication strategy is simple: Ask for high effort, emphasize the mutual benefits of high effort, and suggest employees are well paid.

Prisoners’ dilemma (cooperation)

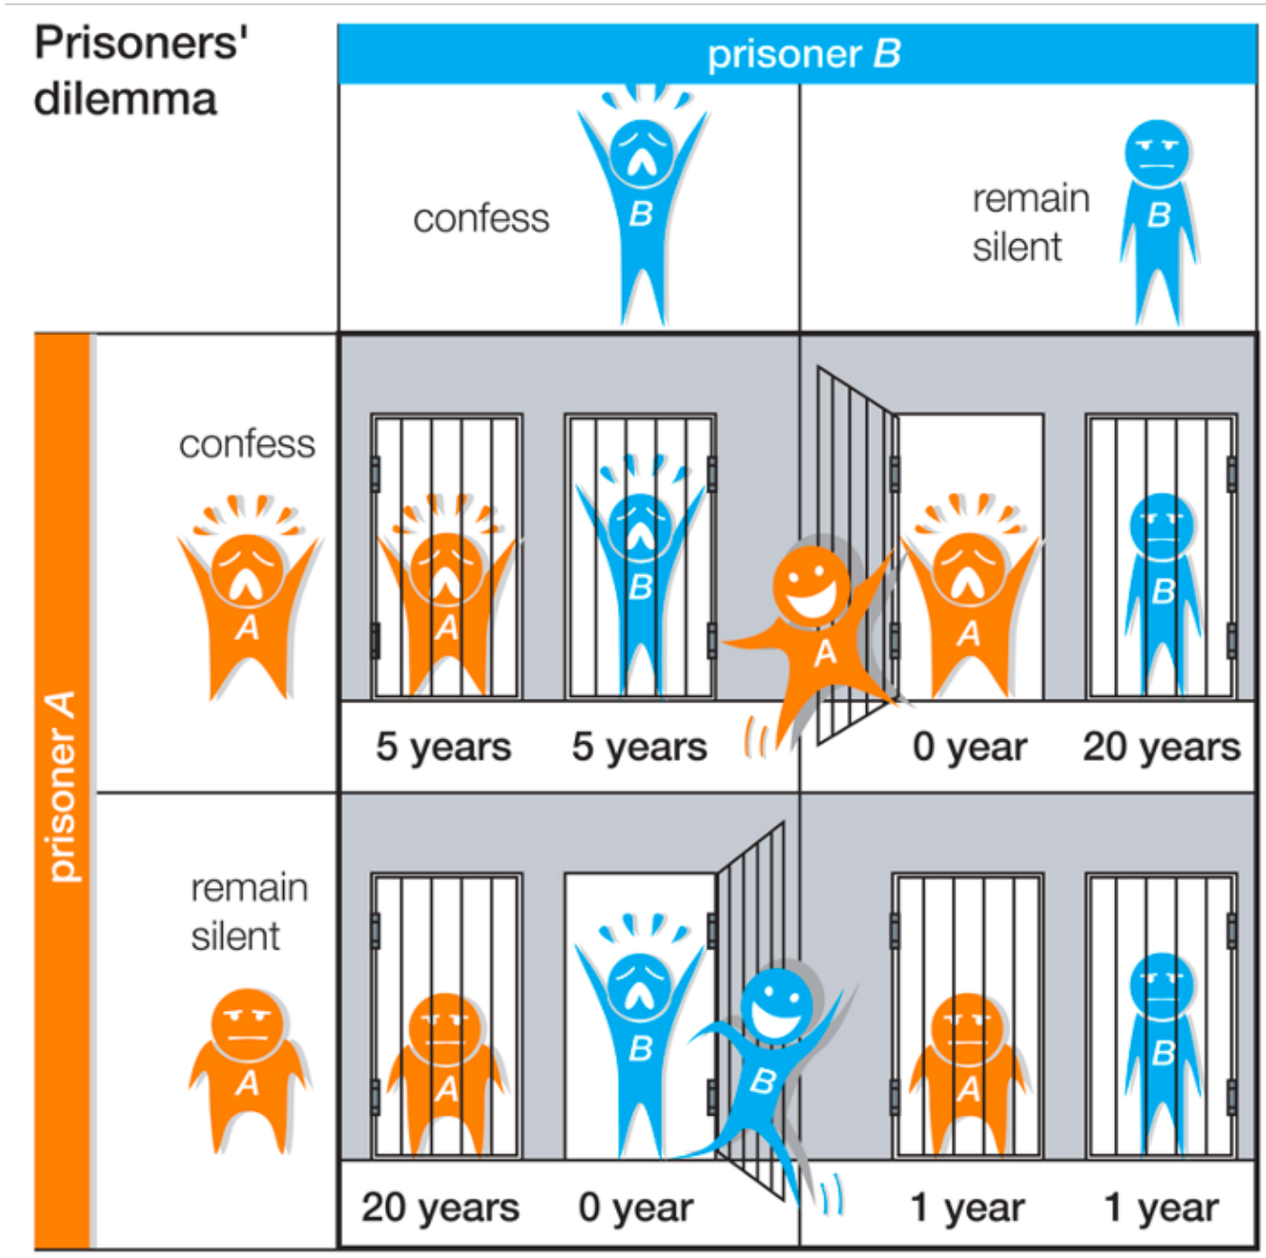

- Two prisoners, A and B, suspected of committing a robbery together, are isolated and urged to confess.
- Each is concerned only with getting the shortest possible prison sentence for himself; each must decide whether to confess without knowing his partner’s decision.
- Both prisoners, however, know the consequences of their decisions:
  - (1) if both confess, both go to jail for five years;
  - (2) if neither confesses, both go to jail for one year (for carrying concealed weapons);
  - (3) if one confesses while the other does not, the confessor goes free (for turning state’s evidence) and the silent one goes to jail for 20 years.

Prisoners' dilemma (cooperation)

- Although A cannot be sure what B will do, he knows that he does best to confess when B confesses (he gets five years rather than 20) and also when B remains silent (he serves no time rather than a year); analogously, B will reach the same conclusion.

- So the solution would seem to be that each prisoner does best to confess and go to jail for five years.

- Paradoxically, however, the two robbers would do better if they both adopted the apparently irrational strategy of remaining silent; each would then serve only one year in jail.

- The irony of PD is that when each of two (or more) parties acts selfishly and does not cooperate with the other (that is, when he confesses), they do worse than when they act unselfishly and cooperate together (that is, when they remain silent).

- PD is not just an intriguing hypothetical problem; real-life situations with similar characteristics have often been observed.

- For example, two shopkeepers engaged in a price war may well be caught up in a PD.

- Similarly, nations competing in an arms race and farmers increasing crop production can also be seen as manifestations of PD.

One-shot vs. repeated game

- People often interact in ongoing relationships. For example, most employment relationships last a long time. Countries competing over tariff levels know that they will be affected by each others' policies far into the future. Firms in an industry know that they are not playing a static game but one in which they compete everyday over time.

- In all of these dynamic situations, the way in which a party behaves at any given time is influenced by what this party and others did in the past. In other words, players "condition" their decisions on the history of their relationship. An employee may choose to work hard only if his employer gave him a good bonus in the preceding month. One country may set a low import tariff only if its trading partners had maintained low tariffs in the past. Repeated games help explain why ongoing economic phenomena produce behavior very different from those observed in a one-time interaction.

- If players believe that future behavior will be affected by the nature of current interaction, they may behave in ways that they would not otherwise. The prospect of reciprocity, either by way of rewards or punishments, is what separates a repeated game from a one-shot game. Rewards or punishments have to be credible in the sense that players will only believe them if they are part of a subgame perfect equilibrium.If a player believes that

    - "no good deed today will go unrewarded tomorrow", then he will have a greater reason to do a good deed

    - "no bad deed today will go unpunished tomorrow", he may be less inclined to do a bad deed today.

Oligopoly market structure

- The petrol/service station market structure is usually that of an oligopoly. An oligopoly has a small number of firms supplying the market, so there is some competition. In an oligopoly, a firm is concerned with how their rivals will react to any action it takes.
- Let's assume the following: The monopoly price is $2 and the quantity traded is 50,000 litres. The competitive price is $1 and the quantity traded is 100,000 litres. If there is only one petrol station, so it produces the monopoly outcome, and makes positive economic profits of $50,000.
- What happens if another petrol station enters the market? Either they collude on monopoly price, or the price war begins…
-

Station 1 Output	Station 2 Output	Total Output	Price	Station 1 Profit	Station 2 Profit
25000	25000	50000	$2	25000	25000
25000	35000	60000	$1.80	20000	28000
35000	35000	70000	$1.60	21000	21000
35000	45000	80000	$1.40	14000	18000

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

- 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

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.

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 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.

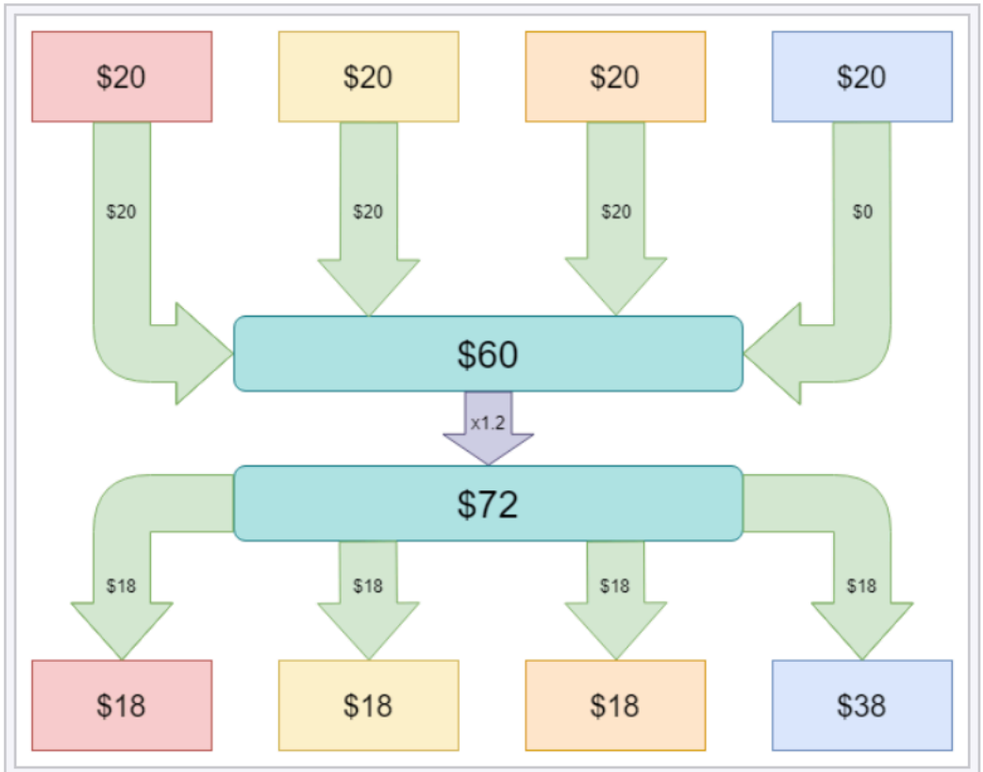

In this diagram of a public goods game, three players choose to contribute their full $20 while the fourth player chooses to contribute $0. The $60 is multiplied by a factor of 1.2 and the resulting $72 is distributed equally among the four players.

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.

- https://www.youtube.com/watch?v=S0qjK3TWZE8