Introduction to Game Theory - Lecture Notes

851-0588-00L

Spring 2020

Professors: Heinrich H. Nax & Bary S.R. Pradelski

Scribe: Han-Miru Kim

16. Mai 2020

Inhaltsverzeichnis

1 Introduction **1**
1.1 What is game theory? 1
1.1.1 Games and Non-Games 1
1.1.2 Prescriptive vs. descriptive agenda 1
1.2 Non-Cooperative vs. Cooperative Games 2
1.2.1 Non-cooperative Games 2
1.2.2 Nash Equilibrium 3
1.2.3 Cooperative Games 5

2 Cooperative Game Theory **6**
2.1 The Cooperative Game 6
2.1.1 Characteristic function form Games 6
2.2 The Core 7
2.2.1 Properties of the Core 7
2.2.2 Balancedness 8
2.3 Shapley Value 9
2.3.1 Finding the Shapley Value 9
2.3.2 Meaning of the Shapley Value 10
2.4 Other cooperative models 11
2.4.1 Matching problem 11

3 Non-Cooperative Game Theory **13**
3.1 Preferences and Utility 13
3.1.1 From Preference to Utility 13
3.2 Utility Function 14
3.2.1 Independence of irrelevant alternatives 15
3.2.2 Ordinal vs. Cardinal vs. Utils 16
3.2.3 Utility and Risk 17
3.3 Nash Equilibrium 18

4 Interactive Environments and Distributed Control **21**
4.1 Comparison to other agendas 21
4.1.1 Centralized vs. distributed control 21
4.2 Distributive Systems 22
4.2.1 Motivation 22
4.2.2 Solution Concepts 22
4.2.3 Descriptive Agenda analysis 22
4.2.4 Prescriptive agenda 23

Lecture 1 April 29

1 Introduction

Game theory is a relatively new field of mathematics. Its foundations were laid by the great mathematician John von Neumann(1903-1957) whose ideas were also valuable in other related fields, such as computer science and physics.

1.1 What is game theory?

As Myerson puts it: it is a mathematical language to express models of “conflict and cooperation between intelligent rational decision-makers”. It studies the interaction between decisionmakers whose decisions interact with each other, which is why Aumann calls it *interactive decision theory*. What is often criticized is how game theory models “rational” behaviour, but these assumptions lie in the model studied and are not inherent to the mathematical language of game theory. At its inception, game theory was applied to mundane concepts like board games, or “Gesellschaftsspiele”. Yet today, game theory sees its applications in much borader fields, such as Economics, Auctions and even Political decisionmaking.

1.1.1 Games and Non-Games

What is a game? And what is not a game? An example of a Non-Game would be roulette. The outcome of the ’Game’ is not depended on either your decision-making nor the decisions made by the casino. Your payoff is merely dependent on chance. In contrast, a ’Game’ like Poker could be considered a Game, since the decisions made by each Player directly affect the outcome of the Game.
But the boundary between Games and Non-Games isn’t always clear. For instance, if an everyday person bought or sold some Stonks on the stockmarket, the effects of your tradings would have a negligible impact on the stockmarktet. However, if big corporations, banks or nations decide to move currencies, buy stock or sell shares, their decisions will noticeably affect the situation in the market world. This is when game theory comes into play to analyse what is happening.

1.1.2 Prescriptive vs. descriptive agenda

Game theory can be used to both “reverse engineer” the mechanisms of a model, or to construct mechanisms, to create a most desirable output. Examples of descriptive game theory can be found in Biology, Social Sciences, Economics, Computer science etc., whereas prescriptive game theory can help us in the field of Law, medical research or Macroeconomics.
Its impact in economics is extremly outstanding, with over seven Nobel prizes given to game theorists.

1.2 Non-Cooperative vs. Cooperative Games

We find that we usally can describe games as being either non-cooperative, or cooperative games. What distinguishes the two is the rules in the game. In non-cooperative game theory, the player are individuals and they are unable to write binding contracts with eachothers. Some key concepts to solving these are the **Nash equilibrium**, or the **Strong equilibrium**.
By contrast, cooperative game theory is a formulation of rules, where players can be individuals or can merge into larger groups of individuals. The groups, or coalitions can be formed by binding contracts or other forms of leverage. These problems are drastically different from non-cooperative game theory. We use **Core** or the **Shapley value** to study these problems.
An prominent researcher in this field is JOHN NASH, who recieved a Nobel Prize, as well as other prestigious prizes for his outstanding research in non-cooperative game theory.

1.2.1 Non-cooperative Games

Definition 1.1: Non-cooperative game (normal-form)

A non-cooperative game consists of three main ingredients.

- A finite set of **players:** $N = \{1, 2, \ldots, n\}$.
- A finite set of **actions/strategies**, where the player i choses a strategy s_i from his stategy set S_i, which results in a stategy combination $s = (s_1, \ldots, s_n) \in (S_i)_{i\in\mathbb{N}}$.
- A **payoff** function u, which gives player i the payoff $u_i(s)$, depending of the outcome s.

To illustrate how game theory works, we will look into some 2-player examples. These games are usually represented in Matrix form.
To read these matrices, we distribute the points in the outcome as follows. The number of the left $(-,-)$ determintes the payoff for the Player 1 on the left, wheres the number on the right $(-,-)$ is the payoff for the Player 2 above.

- **Matching Pennies:** In this game, the two players have opposite goals. Player 1 always wants to match options with Player 2, and Player 2 wants to make the opposite choice of Player 1.

		Player 2	
		Heads	Tails
Player 1	Heads	1,-1	-1,1
	Tails	-1,1	1,-1

- **Prisoner's Dilemma:** In this Game, two robbers are caught commiting a crime together. They are being seperately interrogated, where each prisoner has the option to either confess or stay quiet. If both stay quiet, the police has no good evidence to keep the robbers at the police station for a long time. If one confesses, the police are lenient on the sentences, while punishing the other member.

		B	
		Confess	Stay quiet
A	Confess	-6,-6	0,-10
	Stay quiet	-10,0	-2,-2

- **Battle of the sexes:** In this classic 1950's game theory example, a man and a woman want to see eachother, and want to meet up in either a Boxing event or the shopping mall. While the man prefers the Boxing event and the woman prefers the shopping, they ultimately want to see eachother, which is why both get zero payoff in case they don't meet.

		Woman	
		Boxing	Shopping
Man	Boxing	2,1	0,0
	Shopping	0,0	1,2

- **Hawk and Dove Game:** Here, two players have the option to either chose a Hawk or Dove. If both chose Dove, they can live in harmony, but if one player is a Hawk, the Hawk can tyrannize the Dove and gets off better. However, if both chose Hawk, they have to fight eachother and get off worse.

		Player 2	
		Hawk	Dove
Player 1	Hawk	-2,-2	4,0
	Dove	0,4	2,2

- **Harmony Game:** Altough we are looking at non-cooperative Games, where binding contracts can be written, “Cooperation” can still arise. For example, consider two companies, which have the choice to cooperate, the payoff function can still dictate wether cooperation is achieved or not.

		Company 2	
		Cooperate	Not Cooperate
Company 1	Cooperate	9,9	4,7
	Not Cooperate	7,4	3,3

1.2.2 Nash Equilibrium

To analyse and solve these games just discussed, we introduce the concept of an **Equilibrium**.

Definition 1.2: Equilibrium

An **equilibrium** or **solution** of a game is a rule that maps the structure of a game into an equilibirum set of strategies s^*.

What game theory is concerned with is finding out what these equilibira are. To help find these, we first look at what a single Player might do given the knowledge of what the other players are doing. We denote the set of strategies *except* for player i with $s_{-i} = (s_1, \ldots, s_{i-1}, s_{i+1}, \ldots, s_n)$, i.e. the strategies which all the other players have chosen.

Definition 1.3: Best-response

Player i's **best-response** to the strategies s_{-i} played by all others is the strategy $s_i^* \in S_i$ such that

$$\mu_i(s_i^*, s_{-i}) \geq \mu_i(s_i', s_{-i}), \quad \forall s_i' \in S_i \text{ and } s_i' \neq s_i^*$$

Now, instead of focusing on a single player, we look at the strategies that *all* players might chose. To do so, we introduce the concept of a **Nash equilibrium**.

Definition 1.4: (Pure-strategy) Nash equilibrium

The **Nash equilibrium** is the outcome, where everybody choses the best strategy to the others best strategies. Which is to say that all strategies are *mutual best responses*, where no player has any interest in switching strategies.

$$\mu_i(s_i^*, s_{-i}) \geq \mu_i(s_i', s_{-i}), \quad \forall s_i' \in S_i \text{ and } s_i' \neq s_i^*$$

Using these few definitions, we can already try to solve some of the previously stated games.
NOTE: We will use the following convention to improve readability: In two-player games like before, we will call Player 1 a 'he', and Player 2 a 'she'. This has the benefit of simplifying the language when discussing the previous games.

Example 1.5: Prisoner's Dilemma

In this example, we look at the game from Player 1's perspective.

		B	
		Confess	Stay quiet
A	Confess	-6,-6	0,-10
	Stay quiet	-10,0	-2,-2

In the case, where Player 2 stays quiet, it makes sense for Player 1 to confess, because the coutcome (0,-10) is better for him than (-2,-2). So 'Confess' is Playe 1's **best response** to Player 2's staying quiet. In the other case, where she confesses, it still is better for him to confess, or else he would face harder consequences. Again, 'Confess' still is Player 1's **best response**. So regardless of what Player 2 does, in every possible situation s_{-1} it is always better for Player 1 to confess. We say that the strategy 'Confess' **strictly dominates** the other strategy, because the payoff function μ_1 is better in all cases s_{-1}.
Using the same argument for Player 2, we can see that the game will lead to the result, where both players confess (-6,-6).

Example 1.6: Battle of the sexes

In the Battle of the Sexes game, there appear to be two stable, pure strategy Nash Equilibria. It is when both the man and the woman chose the same option, because if one such Equilibrium is met, it doesn't make sense for either player to deviate from the Equilibrium, i.e. change strategies.

		Woman	
		Boxing	Shopping
Man	Boxing	2,1	0,0
	Shopping	0,0	1,2

Pure stategy Nash Equilibria don't always have to exist for every game. In the game **Matching Pennies**, there don't exist any pure strategies, which are best responses to another, because we see a form of anti-coordination in the rules of the game.
In these examples, we have seen that the Nash Equilibria can lead all kinds outcomes, both socially desireable and undesireable ones, aswell as Pareto-optimal ones, where the outcome of one party connot be increased without reducing the profit of another party by a equal or larger amount.

Prisoner's dilemma	Unique NE	socially undesirable outcome
Harmony Game	Unique NE	socially desirable outcome
Battle of the Sexes	Two NE	Pareto-optimal
Hawk vs. Dove	Two NE	Pareto-optimal
Matching pennies	No pure strategy NE	

An example of the counter intuitive-ness of Nash equlibris is presented in the following example, where the addition of seemingly good options can worsen the outcome in the Nash Equilibrium.

Example 1.7: Braess' Paradox

The story goes as follows: 60 people want to travel from S to D. From S, they have the option to either go to A or B. The road SB takes one hour, and the time to travel SA depends on how many drivers are on the road, same goes for BD.
In the following diagram, the numbers on the arrows indicate the travel times.

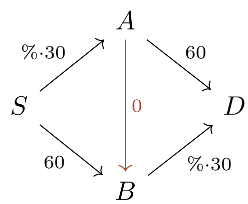

Initially, there is no middle road. The Nash Equilibrium is such that 30 people drive one way, and 30 people drive the other way, where both groups of drivers take 90 minutes.
But after the construction of the super efficient middle road, everyone will use it, but as a side effect, the congestion on the road BD will worsen, where the Nash Equilibrium will then be at 119/120 minutes.

1.2.3 Cooperative Games

We have seen that the Nash Equilibrium does not always lead to the most collectively desireable outcome. How can player overcome this problem? One way to achieve this would be to let player cooperate by writing **binding agreements** or **transferring unitilities**.
We can introduce the **cooperative value**, which takes the sum of all payoffs and try to maximize that. In the prisoner's dilemma, this would change the outcome from (-6,-6) to (-2,-2), since now, the cooperative value is now being considered. By defining $v_{12}(s) = \mu_1(s) + \mu_2(s)$, we would instead get this game matrix.

		B	
		Confess	Stay quiet
A	Confess	-12	-10
	Stay quiet	-10	-4

Lecture 2 April 30

2 Cooperative Game Theory

Cooperative Game theory concerns itself with players, where they may form coalitions or groups of individuals. Often the players may be individuals, corporations or nations and they are able to write binding contracts, transfer utility.

2.1 The Cooperative Game

The Cooperative game was introduced int the paper The Theory of Games and Economic behaviour (1944) published in Princeton University press written by John von Neumann (1903 - 1958) and Oskar Morgenstern (1902-1977), where the two defined the core notions still in use in modern game theory, which will take a look at in this chapter.

Definition 2.1: Cooperative game (normal-form)

Like a non-cooperative game, a **cooperative game** consists of three main ingredients

- A finite Population of **players:** $N = \{1, 2, \ldots, n\}$.
- **Coalitions** $C \subseteq N$ which form in the population and become players again, resulting in a coalition structure $\rho = \{C_1, C_2, \ldots, C_k\}$
- A finite set of **actions/strategies**, where the player i choses a strategy s_i from his stategy set S_i, which results in a stategy combination $s = (s_1, \ldots, s_n) \in (S_i)_{i \in \mathbb{N}}$.
- Compared to the non-cooperative game [1.1], **payoffs** $\varphi = \{\varphi_1, \ldots, \varphi_n\}$ are a bit more complicated using a **sharing rule** for players within a coalition.

Of particular interest is how the payoff function will behave given the set of coalitions. What determines how the shared payoffs will be distributed?

2.1.1 Characteristic function form Games

A common representation of a cooperative game is using the **characteristic function form** (CFG):

Definition 2.2: Characteristic function form

Here, the game is defined by a 2-tuple $G(v, N)$, where we again have a finite fixed population N together with *disjoint* **Coalitions** $C \subseteq N$ resulting in the **coalition partition** ρ. Some examples are the *empty coalition* $\emptyset$, the *grand coalition* N. We write 2^N for the set of all coalitions and ρ for the set of all partitions. The **characteristic function** v is the function form that assigns a *worth* $v(C)$ to each coalition.

$$v : 2^N \to \mathbb{R}, \quad C \mapsto v(C) \quad \text{and} \quad v(\emptyset) = 0$$

We can also think of the singleton coalitions ot be worth zero.

As an example, consider a 3-player game ($N = 1, 2, 3$), where individual players are "worth" nothing, but the coalitions (1,2) and (1,3) are worth 0.5, and (2,3) nothing, with the grand coalition being worth 1. Then, the characteristic function will be

$$v(i) = 0, \quad v(1,2) = v(1,3) = 0.5, \quad v(2,3) = 0, \quad v(N) = 1$$

What could this game represent? Let's think of player 1 as the owner of a machine that can produce a good, but requires workers. Players 2 and 3 could be workers, who alone can't produce anything but can work with the machine.

Definition 2.3: Transferable utility and feasibility

Let $G(v, N)$ be a CFG. The **outcome** is a coalition structure, consisting of

- The resulting **partiton** $\rho = \{C_1, \dots, C_k\}$ and
- **payoff allocation** $\varphi = \{\varphi_1, \dots, \varphi_n\}$

Importantly, $v(C)$ can be shared amongst $i \in C$ (transfer of utility). But the **feasibility constraint** is that the sum of all allocated utility can not exceed the worth of the coalition

$$\sum_{i \in C} \varphi_i \leq v(C), \forall C \in \rho$$

In the previous 3 player example, some feasible outcomes could be

- Outcome 1: $\{(1, 2), 3\}$ and $\{(0.25, 0.25), 0\}$, where the utility is spread evenly among the participants in the coalition.
- Outcome 2: $\{N\}$ and $\{0.25, 0.25, 0.25\}$, where everybody gets equal amounts in the grand coalition, but the whole worth of the coalition is not fully utilised.
- Outcome 3: $\{N\}$ and $\{0.8, 0.1, 0.1\}$, where player 1 gets the largest share.

Another assumption oftem made in CFG is the so-called **Superadditivity** assumption of coalitions. If two coalitions C and S are disjoint, then $v(C) + v(s) \leq V(C \cup S)$ i.e. mergers of coalitions weakly improve their worths. This implies that the grand coalition is the most *efficient* of them all.

2.2 The Core

The first of the solution concepts introduced is the idea of the **Core**, it is one of the most fundamental solution concepts in cooperative game theory and actually predates the earliest formulations of game theory, having been used in economic studies.

Definition 2.4: The Core

The **Core** of a superadditive CFG $G(v, N)$ consists of all outcomes where, where the *grand coalition* forms and where the payoff allocations φ^* are

- **Pareto-efficient:** $\sum_{i \in N} \varphi_i^* = v(N)$ i.e. the worth of the grand coalition is fully utilised
- **unblockable**: For all $C \subseteq N, \sum_{i \in C} \varphi_i^* \geq v(C)$. In other words, every player gets at least as much as they are worth themselves, $\varphi_i^* \geq v(i)$ and every coalition gets as much in total, as it would if it would form a coalition on its own.

This means that in the Core, every player and every coalition is incentivised to stay in the grand coalition, given the payoff allocation φ^* as it would not do better on their own.

In the 3 player example discussed earlier, Outcome 1: $\{(1, 2), 3\}$ and $\{(0.25, 0.25), 0\}$, is clearly not in the core, as the grand coalition is not formed. Outcome 2 : $\{N\}$ and $\{0.25, 0.25, 0.25\}$ aswell as Outcome 3: $\{N\}$ and $\{0.8, 0.1, 0.1\}$, are in the core, since no subcoalition could do better in other outcomes.

2.2.1 Properties of the Core

The core can be defined using a system of linear inequalities, $Ax \leq b$ or $By \geq c$, which means the the Core is **closed** and **convex**. The core could be anything from empty, non-empty or really large.

Example 2.5: Empty Core

In this 3 player game, we set

$$v(i) = 0, \quad v(i,j) = 0.9, \quad v(N) = 1$$

If the grand coalition formed, then two players can always break off and increase their total payoffs, leaving the third player empty handed. There always exists a pair of players, which would block the grand coalition from forming.

In the next example, the core is neither empty nor large. The core can sometimes consist of only a unique outcome.

Example 2.6: Unique Core

Here we again have a 3 player CFG and set

$$v(i) = 0, \quad v(i,j) = \frac{2}{3}, \quad v(N) = 1$$

The only stable payoff allocation is the one, where everybody gets an equal $\frac{1}{3}$. If the coutcome were otherwise, two players could break off and would do better *together* (Not necessarily individually!)

A trivial example of a very large core would be to have the characteristic function be zero for all small coalitions. Then, every possible payoff allocation for the grand coalition would be in the core, as any breaking off would leave everyone with zero payoff.

2.2.2 Balancedness

We saw that when the game had a nonempty core, the outcome seemed to show "balance". In the one case, where the core was empty, players could "betray" another by leaving the coalition and forming another one, leaving the left out player off worse. The next theorem gives us the ability to assess more easily and precisely when the core is empty or not.
To make sense of the term *Balancedness*, we introduce the following terms:

Definition 2.7: Balancedness

- **Balancing weight:** Let $\alpha(C) \in [0,1]$ be the balancing weight attached to any $C \in 2^N$.
- **Balanced family:** A set of balancing weights α is a balanced family, if for every i, $\sum_{C \in 2^N:\, i \in C} \alpha(C) = 1$

The Balancedness in a superadditive game then requires that, for all balanced families:

$$v(N) \geq \sum_{C \in 2^N} \alpha(C) v(C)$$

The balancing weight can be seen as how much time an individual spends in a coalition. If a game had empty core, a player could spend his time better in smaller coalitions than in the grand coalition, causing the instability of the grand coalition.
The next theorem was independently proven by the two mathematicians Olga Bondareva (1937 - 1991) and Lloyd Shapley (1923 - 2016), which states the following:

Theorem 2.8: Bondareva-Shapley

The core of a cooperative game is **nonempty** *if and only if* the game is **balanced**

Note that the balancedness of an outcome is not a measure of its equitability. In the 3 player game where

$$v(i) = v(2,3) = 0, \quad v(N) = v(1,2) = v(1,3) = 1$$

the unique core would be the outcome $(1, 0, 0)$. Although not fair, this outcome is balanced and therefore stable. This is again another example of game-theoretic predictions, where the resulting outcome is not always socially desirable. It shows us that stability might not actually be the best prescriptor of the decisions we ought to make.

2.3 Shapley Value

We can take another approach to cooperative games and try to look for the most just, fair, right outcomes. To find a method of finding these good outcomes, Lloyd Shapley introduced the concept of the **Shapley Value** in 1953.

Definition 2.9: Shapley value

Given some $G(v, N)$, the **Shapley value** is an outcome where the allocation $x^*(v)$ should satisfy:

- **Efficiency:** $\sum_{i \in N} x_i^*(v) = v(N)$.
- **Symmetry:** If for any two players i and j, $v(S \cup i) = v(S \cup j)$ for all coalitions S not including i and j, then $x_i^*(v) = x_j^*(v)$. Two players who are worth the same to all coalitions should be rewarded the same.
- **Dummy player:** if for any $i, v(S \cup I) = v(S)$ for all S not including i, then $x_i^*(v) = 0$. Freeloading players, who don't contribute to any coalition should not be rewarded anything.
- **Additivity:** If u and v are two characteristic functions, then $x^*(v + u) = x^*(v) + x^*(u)$.

2.3.1 Finding the Shapley Value

To find out what the Shapley Value must be, we can use the following function, which characterizes the Shapley Value.

Theorem 2.10: Shapley's characterization

Given a CFG $G(v, N)$, there exists a unique function satisfying all four conditions of the acceptable allocation for the set of all games and it is given by

$$\varphi_i(v) = \sum_{S \in N : i \in S} \frac{(|S| - 1)!(n - |S|)!}{n!} \left[v(S) - v(S \setminus \{i\}) \right] \tag{2.1}$$

An equivalent characterisation of the four axioms of the Shapley value was found by Peyton Young (1945*), which shows the normative appeal of the Shapley value.

Definition 2.11: Young (1985)

As in the axiomatization of Shapley, Young keeps the first two axioms, but replaces the latter two by a third one. Given some cooperative game, the outcome should satisfy

- **Efficiency:** $\sum_{i \in N} x_i^*(v) = v(N)$
- **Symmetry:** If for any two players i and j, $v(S \cup i) = v(S \cup j)$ for all coalitions S not including i and j, then $x_i^*(v) = x_j^*(v)$.
- **Monotonicity:** If u and v are two characteristic functions, then

$$\forall S : \; i \in S : \quad u(S) \geq v(S) \implies x_i^*(u) \geq x^*(v)$$

The Monotonicity is a more attractive phrasing of the Additivity and the Dummy player, as it requires that players who are worth more, should get a better payoff.

Example 2.12: Shapley value

- In the 3 player game, with empty core where the characteristic function is given by

$$v(i) = 0, \quad v(i,j) = \frac{5}{6}, \quad v(N) = 1$$

Then, the desireable outcome would yield the Shapley value $\left(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}\right)$.

- In our previous example with nonempty core, but still quite inequitable outcome, the characteristic function

$$v(i) = v(2,3) = 0, \quad v(N) = v(1,2) = v(1,3) = 1$$

The Shapley value would be the allocation $\left(\frac{2}{3}, \frac{1}{6}, \frac{1}{6}\right)$, which is more fair because players 2 and 3 still have some marginal contribution.

- In our example with the very large core $v(i) = v(i,j) = 0$ and $v(N) = 1$. The Shapley value again yields to the nice outcome $\left(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}\right)$.

2.3.2 Meaning of the Shapley Value

It might not be directly clear what the function [2.1] calculates. To better our understanding, we can, think of the *marginal contribution* of a player i in a coalition S. The marginal contribution of player i to the coalition S

$$MC_i(S) = v(S) - v(S \setminus \{i\})$$

is player i's contribution to the worth of the coalition S. Then we want the payoff for the player i to be the average of his marginal contribution to all the coalitions he is in. The term $\frac{(|S|-1)!(n-|S|)!}{n!}$ is counting in how many possible ways that particular combination/permutation of players can be arranged.
Consider a set of n players entering a room in some order.

- Whenever a player i enters a room, and the players $S \setminus \{i\}$ are already there, the player i is paid his marginal contribution $MC_i(S) = v(S) - v(S \setminus \{i\})$.
- Suppose all $n!$ orders or permutations are equally likely. Then there are $(|S| - 1)!$ different orders in which these players $S \setminus \{i\}$ can be there before i. And there can be $(n - |S|)!$ different orders in which the remaining players can enter the room.
- This means, that there are $(|S| - 1)!(n - |S|)!$ orders out of the $n!$ possible orders, in which player i enters the Room at step i.

Using this reasoning, we end up with the formula

$$x_i^*(v) = \sum_{S \in \rho, i \in S} \frac{(|S| - 1)!(N - |S|)!}{n!} \cdot MC_i(S)$$

We have to remember that the Shapley Vale is normative, i.e. it is an outcome which we, as humans find pleasant. Finding the Shapley value therefore is more on the prescriptive side of game theory than on the descriptive side. By contrast, the Core is more descriptive in nature, as it is looking for stable outcomes, rather than desirable ones. Some key characterisitcs are

- When the Core is non-empty, the Shapley Value may lie inside the core or not.
- When the Core is empty, the Shapley Value is uniquely determined.

2.4 Other cooperative models

There are other ways to model cooperative games. One such model is the **Non-transferable-**utility cooperative game. As before, we have a CFG defined by the tuple $G(v, N)$, where in the **Outcome**, the partition $\rho = \{C_1, \dots, C_k\}$ *directly implies*, how the payoff is allocated. Which means that $\varphi_i = f_i(C_i)$ is fixed for every coalition C_i. In this model, individual player have preferences over the Coalitions, where the utility cannot be re-negotiated within each coalition.

2.4.1 Matching problem

An example of such a cooperative game with non-transferable utility would be the **Stable Marriage/Matching problem**.

Example 2.13: Matching Problem

In this problem, we have two sets of players. Men $M = \{m_1, \dots, m_n\}$ and Women $W = \{w_1, \dots, w_n\}$. Each man has preferences on how to match with a person from W, which give us a strict ordering, i.e. $m_i(w_j)$ forms a strict order. The same goes for the women.
It is clear how there can be no transfer of utilty, as one only cares about their own match.

Ideally, we want to establish a **stable matching**. A formation of couples (man-woman) such that there exists no alternative couple where *both* partners prefer to be matched with each other rather than with their current partners. In a paper published by Gale and Shapley (College Admissions and the Stability of Marriage), they found a theorem to adress these stable matchings. There are two results in this theorem. That a stable matching always exists and that there is an algorithm to obtain such a matching.

Theorem 2.14: Gale-Shapley 1962

For any marriage problem, one can make all matchings stable using the deferred acceptance algorithm

Algorithm 1: Stable Matching algorithm

```
Result: Stable matching between M and W.
Initialize: all m_i ∈ M and all w_i ∈ W are single. ;
while There exists a single man do
    Each single man m ∈ M proposes to his preferred woman w to whom he as not yet proposed.;
    if w is single. then
        w will become engaged with her preferred proposer;
    else
        w is already engaged with m';
        if w prefers her preferred proposer m over her current engagement m' then
            (m, w) become enaged;
            m' becomes single;
        end
    end
end
Marry all engagements.
```

This alogrithm is often seen in practice in varous fields. Following a widespread textbook by Roth & Sotomayor 1990 (Two-Sided Matching). For example, the algorithm is used in the organ exchange market, where in particular it is prevalent in the kidney-exchange market in the United States, or college admission procedures. It is also used to assign users to servers in their internet services etc.
Why does this algorithm work? Notice these properties of the algorithm.

- **Tradeup:** Women can *trade up* until every woman is engaged. That is, they always get better matches, as the algorithm repeads and never get diched. When everyone is engaged, they all get married.
- **Termination:** The algorithm must end, because it's impossible that there always are some single men. At that point, every single man would have proposed to every woman, and every single woman would have become engaged after being proposed to.
- **Stability:** Is the resulting matching stable? Yes.
  Suppose that the algorithm terminates such that exist a pair (m, w) whose partners are engaged to $w' \neq w$ and $m' \neq m$. Then, it is not possible for both m and w to prefer each other over their current partner, because if m prefers w over w', then he would have proposed to w before he proposed to w'.

  - If at that time, w would have engaged with m and traded up later with some m'', then she would prefer m'' over m and also m', which means that she can't have married m'.
  - If however w was already engaged and rejected m's proposal, then she was already with some m'' who she also prefers over m and m'.

  Therfore, either m prefers w' over w, or w prefers m' over m.

Lecture 3 May 11th

3 Non-Cooperative Game Theory

We now will focus our attention to games, where the sets of actions available to *individual* players are the main component of the games. Here, the strategic interactions between self-interested and independet agents are considered more in depth. Some examples where we may see non-cooperative games play out in the real world can vary from games like Chess or Rock-paper-scissors to more impactful situations like the Cold War.

3.1 Preferences and Utility

In order to be able to rigourously analyze these strategic interactions we need to be able to extract the players **Preferences** or goals into some real-valued function. We call this function the **utility function** which has to incorporate the players preferences.

Definition 3.1: Binary Relation

A **binary relation** $\succeq$ on a set X is a non-empty subset $P \subseteq X \times X$. We write $x \succeq y$ if and only if $(x, y) \in P$

To state preferences, we will write $x \succeq y$ to mean "the player weakly prefers x over y, $x \succ y$ for "the player *strictly* prefers x over y and $x \sim y$ for "the player is *indifferent* between x and y.
In order to meaningfully work with these preferences, we will have to make some assumptions about what the properties of preferences are.

3.1.1 From Preference to Utility

- **Completenes:** If a consumer is chosing between two bundles x and y, one of the following possibilites hold:

  ()i) $x \succ y$: they prefer x to y
  ()ii) $y \succ x$: they prefer y to x
  ()iii) $x \sim y$: they are indifferent between x and y

  At first glace, it might appear that this assumptions seems perfectly logical as one should be able to compare options. Consider yourself at a Chinese Market and you are given the choice between ####. In this case, you can't fully tell which one you prefer or not because you don't know what either of those options really are.

  More generally consumers/agents often find it impossible to rank some option without having a sense of something being not quite right.

  Decision making takes time and we are often uninformed, uncertain, subject to biases or just unable to evaluate what a product or choice is or does.

- **Transitivity:** If a consumer is chosing between trhee bundles x, y, z with $x \succ y$ and $y \succ z$, then $x \succ z$. Again, choices are not always as simple as that. Consumers/agents find it difficult how to rank choices, as our needs are manifold and having to make a choice often leads to the unfulfillment of some desires.

- **Continuity:** If a consumer is chosing between trhee bundles, x, y, z with $x \succ y$ and y is *very similar* to z, then $x \succ z$. To define this more clearly we introduce the following concept:
  Let $\succeq$ be a rational preference ordering on X. For $x \in X$ define the subsets of alternatives that are (weakly) *worse/better* than x to be

  $$W(x) = \{y \in X \ : \ x \succeq y\}, \qquad B(x) = \{y \in X \ : \ y \succeq x\}$$

  Continuity then means, that the sets $W(x)$ and $B(x)$ are *closed*.

  To understand the axiom of continuity, we can think of the rate of consumption of a good as an example. Suppose 100g Müsli per day $\succ$ 200g bananas per day. Then we should also prefer 100g Müsli per day over 201g of bananas per day.

We can write these assumptions as a list of **Axioms**

Axiom 3.2: Axioms of Preference

Completeness: $\forall x, y \in X \ : \ x \succeq y$ or $y \succeq x$ or both.
Transivitiy: $\forall x, y, z \in X \ :$ if $x \succeq y$ and $y \succeq z$, then $x \succeq z$.
Continuity: $\forall x \in X \ : \ B(x)$ and $W(x)$ are closed sets.

3.2 Utility Function

Now that we have an idea on how preferences should work, we must try to model these into a real valued function, whose expected value the agent aims to maximize.

Definition 3.3: Utility Function

A **utility function** for a binary relation $\succeq$ on a set X is a function $u : X \to \mathbb{R}$ such that

$$x \succeq y \Leftrightarrow u(x) \geq u(y)$$

The existence of such a function is ensured by the following proposition.

Proposition 3.4

Version a: There exists a utility function for each complete, transitive, postively measureable and continuous preference ordering on any closed set.
Version b: There exists a utility function for every transitive and complete preference ordering on any countable set.

We will see that in the real world, humans will often behave in ways irrational when looked at from a utility point of view. The frequent emergence of this can be explained by the difference between *utility* and *payoff*

Example 3.5: Coin-Toss Game

In this game, a fair coin is tossed until it shows head for the first time.

- If head turns up at the first toss, you win 1 CHF.
- If head turns up first at the second toss, you win 2 CHF.
- If head turns up first at the third toss you win 4 CHF and generally, if head turns up at the k-th toss, you win 2^{k-1} CHF.

You have a ticket for this lottery. For which price would you sell it?

If you try to calculate the expected gain, we see that $E[]\text{lottery}] = \frac{1}{2} + \frac{1}{2} + \frac{1}{2} + \ldots = \infty$.
An effect that came into play in the previous game is that the more money you win, each *additional* CHF might be *worth* less to you. Bernoulli saw this suggested in 1738 the theory of *diminishing marginal utility* of wealth. Later, the need for utility characterization under uncertainty arose, which laid the foundation for **expected utility theory**.
Let $T = \{\tau_1, \ldots, \tau_m\}$ be a finite set and let X consist of all probability distributions on T. The **unit simplex** in $\mathbb{R}^m$ is the set $X = \Delta(T)$ defined by

$$X = \{x = (x_1, \ldots, x_m) \in \mathbb{R}^m_+ \ : \ \sum_{k=1}^{m} x_k = 1\}$$

We can interpret this set as the set of all *probability distributions* on T.
With the set X, we don't directly chose between options, but probability distributions over those options and the question arises: Can we define a utility function in this setting?

3.2.1 Independence of irrelevant alternatives

Consider a consumer choosing between two bundles x and y with $x \succ y$. If we want every bundle to be independent of eachother, then for any z, the **Indepence of irrelevant alternatives** requires that $x + z \succ y + z$. We can generalize a bit further and use it as an axiom when constructing our utility function.

Axiom 3.6: Independence of irrelevant alternatives

$\forall x, y, z \in X, \forall \lambda \in (0, 1)$ the preference must satisfy

$$x \succ y \implies (1 - \lambda)x + \lambda(z) \succ (1 - \lambda)y + \lambda z$$

We can interpret this in the following way. If you can chose to take a bit less of x or y and instead take an equivalent amount of z, your preference still holds.
It is clear that in many situations this is not the case, as some options "go well" with eachother as one might prefer Budweiser over Corona, but if they have the option of adding a lemon, they go for the combination Corona plus a lemon.
The independence of irrelevant alternatives assumes that any decision can be broken down into its smallest parts.
The assumption that preferences can be expressed in this form is called the **expected utility hypothesis**. One example of this in action is the **Bernoulli function** ν:

Example 3.7: Bernoulli Function/ von Neumann-Morgenstern utility function

Here, we look at utility functions over lotteries. If $\succeq$ is a binary relation on X (lotteries) representing the agent's preferences over lotteries over T. If there is a function $\nu : T \to \mathbb{R}$ such that

$$x \succeq y \Leftrightarrow \sum_{k=1}^{m} x_k \nu(\tau_k) \geq \sum_{k=1}^{m} y_k \nu(\tau_k)$$

then the function $u : X \to \mathbb{R}$

$$u(x) = \sum_{k=1}^{m} x_k \nu(\tau_k)$$

defines a utility function for $\succeq$ on X.

Using the four Axioms from 3.2 and 3.6, von Neumann and Morgensterm were able to prove the existence of such a utility function in their book (The theory of games and Economic Behaviour)[citation].

Theorem 3.8: von Neumann-Morgenstern utility function

Let $\succeq$ be a complete, transitive and continuous preference relation on $X = \Delta(T)$, for any finite set T.
Then $\succeq$ admits a utility function u of the expected-utility from if and only if $\succeq$ meets the axiom of independence of irrelevant alternatives.

The introduction of the fourth axiom however does bring in some problems with our model when analyzing human behaviour.

Example 3.9: Allais paradox

People are given set of prices $T = \{0, 1'000'000, 5'000'000\}$ they like and are asked which proability distribution they prefer.

$$x_1 = (0.00, 1.00, 0.00) \quad \text{or} \quad x_2 = (0.01, 0.089, 0.10)$$

And in a second question they can chose

$$x_3 = (0.90, 0.00, 0, 10) \quad \text{or} \quad x_4 = (0.89, 0.11, 0.00)$$

Most people report $x_1 \succ x_2$ and $x_3 \succ x_4$, which seems a bit problematic:

Suppose (v_0, v_{1M}, v_{5M}) is a Bernoulli function for $\succeq$. Then the preference $x_1 \succ x_2$ implies

$$v_{1M} > 0.01 \cdot v_0 + 0.89 \cdot v_{1M} + 0.1 \cdot v_{5M}$$
$$0.11 \cdot v_{1M} - 0.1 \cdot v_0 > 0.1 \cdot v_{5m}$$

but if we add $0.9 \cdot v_0$ to both sides we would get

$$0.11 \cdot v_{1M} + 0.89 \cdot v_0 > 0.1 \cdot v_{5M} + 0.9 \cdot v_0$$

which implies $x_4 \succ x_3$, which contradicts the second preference most people give.

In order to remedy this problem, the economist Savage defined the **Sure thing principle**:
A decision maker who would take a certain action if he knew that event B happens and also if he knew that *not* B happens, should also take the same action if he knew nothing about B.
This leads to the following lemma:

Lemma 3.10: Sure thing principle and independence of irrelevant alternatives

Assume that everything the decision maker knows is true. Then the **sure thing principle** is equivalent to the **independence of irrelevant alternatives**.

In a quotation by Savage (1954) he states that he knows "of no other extralogical principle governing decisions that finds such ready acceptance".

3.2.2 Ordinal vs. Cardinal vs. Utils

Given a Bernoulli function ν for given preferences $\succeq$ let:

$$\nu' = \alpha + \beta\nu, \quad \alpha \in \mathbb{R}, \beta \in \mathbb{R}^+$$

Then ν' is also a Bernoulli function for another utility function $u' = \alpha + \beta u$.
Using this characterisation, it follows that utility functions are unique up to a positive affine transformation.
This makes intuitive sense, since the utility function should only express the relationships between the outcomes and not the absolute value or difference of the options.
We can exploit the indifference to absolute differences to somehow normalize the utility.

Definition 3.11: Ordinal, cardinal utility functions and Utils

- **Ordinal utility function:** A utility function where differences between $u(x)$ and $u(y)$ are meaningless. Only the fact that $u(x) \geq u(y)$ are meaningful. An ordinal utility function can be subjected to any increasing transformation $f(u)$ which will represent the same preferences $\succeq$.
- **Cardinal utility function:** A utility function where the differences between $u(x)$ and $u(y)$ are meaningful as they reflect the *intensity* of the preferecen. Cardinal utility functions are only invariant to positive affine transformations.
- **Utils:** An even stronger statement would be that there is a fundamental *unit* or *measure* of utility. Such a utility function is not invariant to any transformation.

Comparing these, we can visualize the differences between the three in the chart 1 for the statements

(1) She likes x less than z

(2) She likes x over z twice as much as y over z.

(3) She likes x five times more than y.

	(1)	(2)	(3)
Ordinal utility function	yes	no	no
Cardinal utility function	yes	yes	no
Utils	yes	yes	yes

Table 1: Differences bettween Ordinal & Cardinal utility functions and Utils

Note that Ordinal utility functions do not contain more information that the preference relation itself.
In addition, we can also compare the utilities of different people using **interpersonal comparability** (**IC**). Since this involves some kind of measure, this generally only makes sense when talking about utility functios which are **utils**. In the other two cases, interpersonal comparability might not be guaranteed.

Suppose we have cardinal utility functions u_1, u_2 that are **IC** for agent 1 and 2. We can transform them both by some non-affine increasing transformation resulting in $v_1 = f(u_1), v_2 = f(u_2)$. Then, v_1 and v_2 are no longer cardinal but are **IC**.
Again, we can compare some interpersonal utilty functions and see that they are almost always impossible.

(1) Elisabeth values 1000 CHF less than a starving child values 1000 CHF.

(2) Eve would pay 10 utils for the chocolate, whereas Sarah would only pay 5 utils.

(3) Mother loves d_1 more than d_2 and Father loves d_2 more than d_1.

We see that comparing utilities between agents implies some (social) welfare statement/judgement.

3.2.3 Utility and Risk

Consider a lottery where you recieve τ_1 with probability α and τ_2 with probability $1 - \alpha$. We call the lottery a fair gamble, if and only if $\alpha \cdot \nu(\tau_1) = (1 - \alpha) \cdot \nu(\tau_2)$
We can categorize different types of agents by their aversion or affinity to risks.

Definition 3.12: Risk behaviour

- An agent is **risk-neutral** if and only if he is indifferent between accepting and rejecting all fair gambles, that is for all $\alpha \in [0, 1], \tau_1, \tau_2 \in T$

$$\mathbb{E}[u(\text{lottery})] = \alpha \cdot \nu(\tau_1) + (1 - \alpha) \cdot \nu(\tau_2) = u\big(\alpha \cdot \tau_1 + (1 - \alpha) \cdot \tau_2\big)$$

  An agent is risk-neutral if and only if he has a linear von Neumann-Morgenstern utility function.

- An agent is **risk averse** if and only if he rejects all fair gambles, that is for all $\alpha \in (0, 1), \tau_1, \tau_2 \in T$

$$\mathbb{E}[u(\text{lottery})] = \alpha \cdot \nu(\tau_1) + (1 - \alpha) \cdot \nu(\tau_2) < u\big(\alpha \cdot \tau_1 + (1 - \alpha) \cdot \tau_2\big)$$

  Note that the above is similar to the definition of concave functions, which means that an agent is risk averse if and only if he has a strictly **concave** utility function.

- An agent is **risk seeking** if and only if he strictly prefers all fair gambles, that is for all $\alpha \in (0, 1), \tau_1, \tau_2 \in T$

$$\mathbb{E}[u(\text{lottery})] = \alpha \cdot \nu(\tau_1) + (1 - \alpha) \cdot \nu(\tau_2) > u\big(\alpha \cdot \tau_1 + (1 - \alpha) \cdot \tau_2\big)$$

  which means that an agent is risk seeking if and only if he has a strictly *convex* utility function

Lecture 4 May 15th

3.3 Nash Equilibrium

Lecture 5 Datum

Lecture 6 Date 2

Lecture 7 May 7th

4 Interactive Environments and Distributed Control

In this lecture, we will look at how game theory is used in distrubed control and give an overwiew on the apporaches that applied game theory take and compare this approach in other disciplines that use game theory.
What makes distributed control appealing from a game theoretic standpoint is that it gives us another perspecive on control. Rather than taking top-down approaches, we will often see bottom-up, emerging algorithms in use.
Although it is a relatively new area of game theory, it is often interconnected with other factes of game theory like Behavioural or Algorithmic game theory and gives us surprisingly relevant insights into social sciences and human interactions.

4.1 Comparison to other agendas

In the categorisation between prescriptive and descriptive game theory, Distributed Control leans more on the prespriptive side as we are able to manipulate certain aspects of the game

Keep in mind that game theory describes interactions between *agents* in an *environment* that chose *strategies* given some *information* and get some **outcome**. For example, these may be traders in a stock-market with a set of strategies to chose from to gain favourable outcomes (money).
In Biology, we might study bees in nature, where they have to chose foraging strategies which result in different outcomes, i.e. survival. Here, we typically are not interested to change the system in any way and merely want to describe and explain behavioural patterns.
In Mechanism design, an example would be the auction house, where we, the auctioneer are able to change the game rules to our favour and can give out information to the buyers to gain an advantage.
In Control Theory, an example for the agents would be turbines, which can chose to orient themselves relative to the wind to maximize the energy obtained as the outcome.
In a generic distributed control application, we are looking at mutiplie decision making elements, where we have some very specific **interdependency** between the individual elements. In the turbine example, the front line of turines have special roles. There usually is **no central authority**. We are not thinking of the elements as being parts of a whole unified agent, since the information is distributed between the agents. There is however a **collective performance** being measured in these systems.

	Biology	Social Systems	Mechanism Design	Distributed Control
Game structure	given	given	manipulabe	manipulable
Actions	given	given	given	given
Payoffs	give	given	given	manipulable
Information	given	given	manipulable	given

Table 2: Comparison between Distributed Control and other agendas

4.1.1 Centralized vs. distributed control

In contrast to the optimized approach to control, where we look at the system as a whole, distributive control allows for decentralization, where each system component can act on their own to find out the best outcome.
One reason why we might this decentralisation or loss of control is that the distribution of information can be very costly or generally impossible to analyze, as many systems do not have graph structure to analyze.
This however comes at a cost, where the best outcome or performance in a centralized approach may not be found in a decentralized approach.

4.2 Distributive Systems

Saksena, Oreilly describes distributed systems as being "characterized by decentralization in available information, mutiplicity of decision makers and individuality of objective functions for each decision maker". Compared to Myerson's description of game theory as being "the study of mathematical models of conflict and cooperation between intelligent rational decision-makers", we see that the application of game theory to distributed systems seems very natural, as both carry similar structures of characterization.

4.2.1 Motivation

Recall [Braess' Paradox 1.7], where we saw that local objectives of individual components may lead to behaviour that worsens the collective performance of the system.
How do we get the agents to behave in a way to benefit the overall performance of the group?

4.2.2 Solution Concepts

The solution concept in a distributed environment is to find out what to expect given a certain interaction and then try to manipulate the interactions such that the group behaves such that they achieve the outcome we want.
One such **solution concept** is the [Nash Equilibrium1.4], where people chose the best response given other people's best responses.

In the Keynesian beauty contest game, where we had to choose a number between 0 and 100 such that we get closest to half of the average, the **rational** best reply would be to pick 0. However, since the **percieved** best reply differs, we have to instead pick half of what we think others will play.
In the repeated beauty contest, we see that a repetition of the game *decreases* the average of all guesses, as people's perception of the game changes.

We can therefore see the differences between **Rationality** and **Perception**, which can change over iterations of the game as the strategies undergo **Evolution**.
We see that the shift of focus moves away from the static solution concepts like the Nash Equilibrium towards a more dynamic approach: how players might arrive to a solution.
We can therefore give rules to the system so that the system as a whole evolves towards a goal we want to establish.

Example 4.1: Ficticious play (1951)

This game procedes in **Stages** $t \in T$. Each player can maintain empirical frequencies of the actions the opponents take. The individual will (incorrectly) assume that others will play according to how the played in the past and will select an action that maximizes their expected payoff.

The **Bookkeeping** will be written as $x^i(\cdot)$ = evolving empirical frequency of player i.
We also differentiate between **Discrete Time**, $T = \{0, 1, 2, \dots\}$ and **Continuous time**, $T = [0, \infty)$.

$$\textbf{Discrete time}: \quad x^i(t+1) = x^i(t) + \frac{1}{t+1}\left(x^i(t) - \textbf{rand}[\beta^i(x^{-i}(t))]\right)$$

$$\textbf{Continuous time:} \quad \frac{dx^i}{dt} = -x^i + \beta^i(x^{-i})$$

4.2.3 Descriptive Agenda analysis

Descriptive agenda analysis of these games found various interesting results. For different classes of games, different outcomes will be realized. We therefore can pick and chose the classes depending on what behaviour we find nice.
Since we don't have the formal definitions yet, we can write down some of the findings in an informal manner:

- **Meta-theorem:** For [special structure games] under [specific dynamics], players exhibit **asymptotic** behaviour.

- **Theorem:** For *zero sum games* under **ficticious play**, empirical frequencies converge to the **Nash Equilibrium**.
- **Theorem:** For **matching markets** under *random blocking by pairs*, outcomes converge to **stable matchings**.
- **Theorem:** For **cooperative games** under *random blocking by coalitions*, outcomes may not converge (if the core is empty).

A lot of results in game theory follow the structure of the Meta-theorem.

4.2.4 Prescriptive agenda

In the prescriptive agenda that distributive control adopts, we can use evolutionary dynamics to feed the collective objective into the system. This means that we want to manipulate the individual agents in order to establish a favourable outcome.

Theorem 4.2: Potential games

For **potential games** under *restricted movement log linear learning*, joint actions "linger" at potential maximier

The restricted movement describes that information between agents is restricted and can't move easily.
Now, we want to appropriate the best dynamics to code the indiviual components in a robust way without a central authority, as coordination is sometimes extremely hard to achieve. An example of this would be the **Wind Farm**.

Example 4.3: Wind Farm

Each windmill takes a directional orientation and a blade angle. Depending on the wind direction, this willl lead to energy production for each windmill.
We want to maximize the total energy production, but how do we achieve this, when we don't have a central authority to coordinate the windmills? The centralized approach has been proven unsuccessful, because each turbine does not have acess to the functional form of the power generated by the wind farm. This is because the aerodynamic interaction between turbines is poorly understood.

-23:31