



MAS and Games

Prof. Dr. Jürgen Dix Prof. Dr. Jörg Müller



Time: Monday (13-15), Wednesday: 10–12

Place: NN

Website

<http://www.in.tu-clausthal.de/abteilungen/cig/cigroot/teaching>

Visit regularly!

There you will find important information about the lecture, documents, labs et cetera.

Check also StudIP, which contains updated slides and exercise sheets.

Lecture: Prof. Dix, Prof. Müller

Exam: oral examinations

About this lecture (1)

This MSc course is about **Decision Making in Multi Agent Systems** and is held alternately by Profs. Dix and M \tilde{A} $\frac{1}{4}$ ller. We look at decision making mainly from a **game-theoretical perspective**.

The lecture can be roughly divided into four parts:

- **Foundations, Prof. Müller:** In the first two chapters, the agent paradigm is introduced and the notion of **intelligent agent** is put forward.
- **Game Theory, Prof. Dix:** The third chapter introduces complete information games and the foundations of classical game theory by various notions of equilibria.

About this lecture (2)

- **MAS and Programming, Prof. Müller:** In Chapters 4 and 5 we look at the foundations of **Multiagent systems** and how to program them using the **BDI** framework.
- **Coalitions and decision making, Prof. Dix:** Chapter 6 considers the problem of *how tasks are allocated and how coalitions are forming*.

While agents are essentially self interested, a MAS needs to collectively take decisions and therefore needs protocols to do so. We therefore introduce decisions based on social choice, ranking systems (Chapter 7)).

Main References (1)



Bordini, R. and Dix, J. (2013).

Chapter 13: Programming Multi Agent Systems.

In G. Weiß (Editor), *MultiAgent Systems*, pages 587–639.

MIT Press.



Dix, J. and Fisher, M. (2013).

Chapter 14: Verifying Multi Agent Systems.

In G. Weiß (Editor), *MultiAgent Systems*, pages 641–693.

MIT Press.

Main References (2)



Fudenberg, D. and J. Tirole (1991).

Game Theory.

MIT Press.



Shoham, Y. and Leyton-Brown, K. (2009).

Multiagent Systems - Algorithmic, Game-Theoretic, and Logical Foundations.

Cambridge University Press.



Lecture Overview

Foundations: MAS as a **paradigm** for decentralized systems.
5 lectures.

Game Theory: Introducing equilibria (Nash's theorem) for **complete** information games.
3 lectures, 1 exercise class.

MAS and Programming: BDI framework for programming MAS.
6 lectures, 4 labs

Coalitions and decision making: Task allocation and forming of coalitions. Social choice and voting systems.
6 lectures, 1 exercise class.

Exercises: **2 exercise classes** and 4-5 labs covering theory and practice.



Outline

- 1 Overview
- 2 Intelligent Agents
- 3 Complete Information Games
- 4 MAS: Basic Concepts
- 5 Agent-Oriented Programming
- 6 Coalitional Games
- 7 Social Choice and Auctions



1. Overview

1 Overview



2. Intelligent Agents

2 Intelligent Agents

3. Complete Information Games

- 3 Complete Information Games
 - Examples and Terminology
 - Normal Form Games
 - Extensive Form Games
 - An Example from Economics

Outline (1)

We illustrate the difference between classical AI and MAS. We present several **evaluation criteria** for comparing protocols.

We then introduce the formal machinery of game theory assuming we have **complete information**:

- **normal form (NF) games**, where players play simultaneously,
- **extensive form (tree form) games**, where players play one after another. Here the history plays a role and players come up with strategies depending on the past.
- We also distinguish between **perfect** and **imperfect recall** and discuss various notions of equilibria.
- Finally we consider the existence of equilibria for **market mechanisms**.

Classical DAI: System Designer fixes **Interaction-Protocol** which is uniform for all agents. The **designer also fixes a strategy for each agent.**

Outcome

What is a the **outcome**, assuming that the given **protocol** is followed and the agents follow the given strategies?

MAI: Interaction-Protocol is given. Each agent determines its own strategy (maximising its own good, via a utility function, without looking at the global task).

Global optimum

Find a **protocol** such that if each agent chooses its best local strategy, the **outcome** is a global optimum.



3.1 Examples and Terminology

We need to **compare protocols**. Each such protocol leads to a solution. So we determine how good these solutions are.

Social Welfare: Sum of all utilities

Pareto Efficiency: A solution x is Pareto-optimal, if

there is no solution x' with:

- (1) \exists agent ag : $ut_{ag}(x') > ut_{ag}(x)$
- (2) \forall agents ag' : $ut_{ag'}(x') \geq ut_{ag'}(x)$.

Individual rational: The payoff should be higher than not participating at all.

Stability:

Case 1: Strategy of an agent depends on the others.

The profile $s_{\mathbf{A}}^* = \langle s_1^*, s_2^*, \dots, s_{|\mathbf{A}|}^* \rangle$ is called a **Nash-equilibrium**, iff $\forall i : s_i^*$ is the best strategy for agent i if all the others choose

$$\langle s_1^*, s_2^*, \dots, s_{i-1}^*, s_{i+1}^*, \dots, s_{|\mathbf{A}|}^* \rangle.$$

Case 2: Strategy of an agent does not depend on the others.

Such strategies are called **dominant**.

Example 3.1 (Prisoners Dilemma, Type 1)

Two prisoners are suspected of a crime (which they both committed). They can choose to (1) **cooperate** with each other (not confessing to the crime) or (2) **defect** (giving evidence that the other was involved). Both cooperating (not confessing) gives them a shorter prison term than both defecting. But if only one of them defects (the betrayer), the other gets maximal prison term. The betrayer then has maximal payoff.

		Prisoner 2	
		cooperate	defect
Prisoner 1	cooperate	(3,3)	(0,5)
	defect	(5,0)	(1,1)

- **Social Welfare:** Both cooperate,
- **Pareto-Efficiency:** All are Pareto optimal, except when both defect.
- **Dominant Strategy:** Both defect.
- **Nash Equilibrium:** Both defect.

Prisoners dilemma revisited: $c \geq a \geq d \geq b$

		Prisoner 2	
		cooperate	defect
Prisoner 1	cooperate	(a, a)	(c, b)
	defect	(b, c)	(d, d)

Example 3.2 (Trivial mixed-motive, Type 0)

		Player 2	
		C	D
Player 1	C	(4, 4)	(2, 3)
	D	(3, 2)	(1, 1)

Example 3.3 (Battle of the Bismarck Sea)

In 1943 the northern half of New Guinea was controlled by the Japanese, the southern half by the allies. The Japanese wanted to reinforce their troops. This could happen using two different routes: (1) **north** (rain and bad visibility) or (2) **south** (weather ok). Trip should take 3 days.

The allies want to bomb the convoy as long as possible. If they search north, they can bomb 2 days (independently of the route taken by the Japanese). If they go south, they can bomb 3 days **if the Japanese go south too**, and only 1 day, **if the Japanese go north**.

		Japanese	
		Sail North	Sail South
Allies	Search North	$\frac{2 \text{ days}}{1 \text{ day}}$	$\frac{2 \text{ days}}{3 \text{ days}}$
	Search South		

Allies: **What is the largest of all row minima?**

Japanese: **What is smallest of the column maxima?**

Battle of the Bismarck sea:

largest row minimum = smallest column maximum.

This is called a **saddle point**.



3.2 Normal Form Games

Definition 3.4 (n -Person Normal Form Game)

A finite n -person normal form game is a tuple $\langle \mathbf{A}, \text{Act}, O, \varrho, \mu \rangle$, where

- $\mathbf{A} = \{1, \dots, i, \dots, n\}$ is a finite set of **players** or **agents**.
- $\text{Act} = A_1 \times \dots \times A_i \times \dots \times A_n$ where A_i is the finite set of actions available to player i . $\vec{a} \in \text{Act}$ is called **action profile**. Elements of A_i are called **pure strategies**.
- O is the set of outcomes.
- $\varrho : \text{Act} \rightarrow O$ assigns each action profile an outcome.
- $\mu = \langle \mu_1, \dots, \mu_i, \dots, \mu_n \rangle$ where $\mu_i : O \rightarrow \mathbb{R}$ is a real-valued **utility (payoff) function** for player i .

Note that we distinguish between **outcomes** and **utilities assigned to them**. Often, one **assigns utilities directly to actions**.

Games can be represented graphically using an **n -dimensional payoff matrix**. Here is a generic picture for 2-player, 2-strategy games:

		Player 2	
		a_2^1	a_2^2
Player 1	a_1^1	$(\mu_1(a_1^1, a_2^1), \mu_2(a_1^1, a_2^1))$	$(\mu_1(a_1^1, a_2^2), \mu_2(a_1^1, a_2^2))$
	a_1^2	$(\mu_1(a_1^2, a_2^1), \mu_2(a_1^2, a_2^1))$	$(\mu_1(a_1^2, a_2^2), \mu_2(a_1^2, a_2^2))$

We often forget about ϱ (thus we are making no distinction between **actions and outcomes**). Thus we simply write $\mu_1(a_1^1, a_2^1)$ instead of the more precise $\mu_1(\varrho(\langle a_1^1, a_2^1 \rangle))$. However, there are situations where we need to distinguish between the two, in particular when talking about **mechanism design** (in Chapter ??, and **auctions** (in Chapter 7, Section ??).

Definition 3.5 (Common Payoff Game)

A **common payoff game (team game)** is a game in which for all action profiles $\vec{a} \in A_1 \times \dots \times A_n$ and any two agents i, j the following holds: $\mu_i(\vec{a}) = \mu_j(\vec{a})$.

In such games agents have no conflicting interests. Their graphical depiction is simpler than above (the second component is not needed).

While a team game is on one side of the spectrum, there is another type of games which is on the opposite side:

Definition 3.6 (Constant Sum Game)

A 2-player n -strategy normal form game is called **constant sum game**, if there exists a constant c such that for each action profile $\vec{a} \in A_1 \times A_2$: $\mu_1(\vec{a}) + \mu_2(\vec{a}) = c$.

We usually set wlog $c = 0$ (**zero sum games**).

Constant sum games can also be visualised with a simpler matrix, missing the second component (like common payoff games) $\mu_2(a_1^2, a_2^2)$ in each entry:

		Player 2	
		C	D
Player 1	C	4	2
	D	3	1

Of course, we then have to state whether it is a **common payoff** or a **zero-sum** game (they are completely different).

Pure vs. mixed strategies

What we are really after are **strategies**.

Definition 3.7 (Pure strategy)

A **pure strategy** for a player is a particular action that is chosen and then played constantly.

A **pure strategy profile** is just an action profile

$$\vec{a} = \langle a_1, \dots, a_n \rangle.$$

Are pure strategy profiles sufficient?

Example 3.8 (Rochambeau Game)

Also known as paper, rock and scissors: paper covers rock, rock smashes scissors, scissors cut paper.

		Min		
		P	S	R
Max	P	0	-1	1
	S	1	0	-1
	R	-1	1	0

What about **pure** vs **mixed** strategies?

Definition 3.9 (Mixed Strategy for NF Games)

Let $\langle \mathbf{A}, \text{Act}, O, \rho, u \rangle$ be normal form game. For a set X let $\Pi(X)$ be the set of all **probability distributions** over X . The set of **mixed strategies for player i** is the set $S_i = \Pi(A_i)$. The set of mixed strategy profiles is $S_1 \times \dots \times S_n$. This is also called the **strategy space** of the game.

Note: Some books use

- S_i , with elements s_i , to denote the set of **pure** strategies,
- Σ_i with elements σ to denote the set of mixed strategies,
- u to denote utilities, and
- N to denote the set of agents.

The **support** of a mixed strategy is the **set of actions** that are assigned non-zero probabilities.

What is the payoff of such strategies? We have to take into account the probability with which an action is chosen. This leads to the expected utility $\mu^{expected}$.

Definition 3.10 (Expected Utility for player i)

The **expected utility** for player i of the mixed strategy profile (s_1, \dots, s_n) is defined as

$$\mu^{expected}(s_1, \dots, s_n) = \sum_{\vec{a} \in \text{Act}} \mu_i(\varrho(\vec{a})) \prod_{j=1}^n s_j(a_j).$$

What is the optimal strategy (maximising the expected payoff) for an agent in an 2-agent setting?

Example 3.11 (Fighters and Bombers)

Consider fighter pilots in WW II. A good strategy to attack bombers is to swoop down from the sun: **Hun-in-the-sun strategy**. But the bomber pilots can put on their sunglasses and stare into the sun to watch the fighters. So another strategy is to attack them from below **Ezak-Imak strategy**: if they are not spotted, it is fine, if they are, it is fatal for them (they are much slower when climbing). The table contains the **survival probabilities of the fighter pilot**.

		Bomber Crew	
		Look Up	Look Down
Fighter Pilots	Hun-in-the-Sun	$\frac{0.95}{1}$	$\frac{1}{0}$
	Ezak-Imak		

Example 3.12 (Battle of the Sexes, Type 2)

Married couple looks for evening entertainment. They prefer to go out together, but have different views about what to do (say going to the theatre and eating in a gourmet restaurant).

		Wife	
		Theatre	Restaurant
Husband	Theatre	(4, 3)	(2, 2)
	Restaurant	(1, 1)	(3, 4)

Example 3.13 (Leader Game, Type 3)

Two drivers attempt to enter a busy stream of traffic. When the cross traffic clears, each one has to decide whether to concede the right of way of the other (C) or drive into the gap (D). If both decide for C, they are delayed. If both decide for D there may be a collision.

		Driver 2	
		C	D
Driver 1	C	(2,2)	(3,4)
	D	(4,3)	(1,1)

Example 3.14 (Matching Pennies Game)

Two players display one side of a penny (head or tails). Player 1 wins the penny if they display the same, player 2 wins otherwise.

		Player 2	
		Head	Tails
Player 1	Head	$(1, -1)$	$(-1, 1)$
	Tails	$(-1, 1)$	$(1, -1)$

Definition 3.15 (Maxmin strategy)

Given a game $\langle \{1, 2\}, \{A_1, A_2\}, \{\mu_1, \mu_2\} \rangle$, the **maxmin strategy** of player i is a **mixed strategy** that maximises the guaranteed payoff of player i , no matter what the other player $-i$ does:

$$\arg \max_{s_i} \min_{s_{-i}} \mu_i^{expected}(s_i, s_{-i})$$

The **maxmin value** for player i is $\max_{s_i} \min_{s_{-i}} \mu_i^{expected}(s_i, s_{-i})$.

The **minmax strategy** for player i is

$$\arg \min_{s_i} \max_{s_{-i}} \mu_{-i}^{expected}(s_i, s_{-i})$$

and its **minmax value** is $\min_{s_i} \max_{s_{-i}} \mu_{-i}^{expected}(s_i, s_{-i})$.

Lemma 3.16

In each finite normal form 2-person game (not necessarily constant sum), the maxmin value of one player is never strictly greater than the minmax value for the other.

We illustrate the maxmin strategy using a 2-person 3-strategy constant sum game:

		Player B		
		B-I	B-II	B-III
Player A	A-I	0	$\frac{5}{6}$	$\frac{1}{2}$
	A-II	1	$\frac{1}{2}$	$\frac{3}{4}$

We assume Player A's optimal strategy is to play strategy

- A-I with probability x and
- A-II with probability $1 - x$.

In the following we want to determine x .

Thus Player A's expected utility is as follows:

- 1 when playing against B-I: $0x + 1(1 - x) = 1 - x$,
- 2 when playing against B-II: $\frac{5}{6}x + \frac{1}{2}(1 - x) = \frac{1}{2} + \frac{1}{3}x$,
- 3 when playing against B-III: $\frac{1}{2}x + \frac{3}{4}(1 - x) = \frac{3}{4} - \frac{1}{4}x$.

This can be illustrated with the following picture (see blackboard). Thus B-III does not play any role.

Thus the maxmin point is determined by setting

$$1 - x = \frac{1}{2} + \frac{1}{3}x,$$

which gives $x = \frac{3}{8}$. The **value of the game** is $\frac{5}{8}$.

The strategy for Player B is to choose B-I with probability $\frac{1}{4}$ and B-II with probability $\frac{3}{4}$.

More in accordance with the minmax strategy let us compute

$$\arg \max_{s_i} \min_{s_{-i}} \mu_i^{expected}(s_1, s_2)$$

We assume Player A plays (as above) A-I with probability x and A-II with probability $1 - x$ (strategy s_1). Similarly, Player B plays B-I with probability y and B-II with probability $1 - y$ (strategy s_2).

We compute $\mu_1^{expected}(s_1, s_2)$

$$0 \cdot x \cdot y + \frac{5}{6}x(1 - y) + 1 \cdot (1 - x)y + \frac{1}{2}(1 - x)(1 - y)$$

thus

$$\mu_1^{expected}(s_1, s_2) = y\left(-\frac{4}{3}x + \frac{1}{2}\right) + \frac{1}{3}x + \frac{1}{2}$$

According to the minmax strategy, we have to choose x such that the minimal values of the above term **are maximal**. For each value of x the above is a straight line with some gradient. Thus we get the maximum when the line does not slope at all!

Thus $x = \frac{3}{8}$. A similar reasoning gives $y = \frac{1}{4}$.

Theorem 3.17 (von Neumann (1928))

In any finite 2-person constant-sum game the following holds:

- 1** The **maxmin value** for one player **is equal** to the **minmax value** for the other. The maxmin of player 1 is usually called **value of the game**.
- 2** For each player, the set of maxmin strategies coincides with the set of minmax strategies.
- 3** The maxmin strategies are **optimal**: if one player does not play a maxmin strategy, then its payoff goes down.

From now on we use just $\mu_1(s_1, s_2)$ instead of the more precise $\mu_1^{expected}(s_1, s_2)$. It will be clear from context whether the argument is a profile (and thus it is the expected utility $\mu^{expected}$) or it is the utility of an outcome (and thus it is defined in the underlying game with μ).

What is the optimal strategy (maximising the expected payoff) for an agent in an n -agent setting?




Figure 1: A saddle.

Definition 3.18 (Notation s_{-i}, S_{-i})

Note that from now on, for $\vec{s} = \langle s_1, s_2, \dots, s_n \rangle$ we use the notation $s_{-i}^{\vec{s}}$ to denote the strategy profile

$\langle s_1, s_2, \dots, s_{i-1}, s_{i+1}, \dots, s_n \rangle$: the strategies of all opponents of agent i are fixed. For a set of strategies S , we denote by $S_{-i} = \{s_{-i}^{\vec{s}} \mid \vec{s} \in S\}$.

For ease of notation, we also use $\mu_i(s_i, s_{-i}^{\vec{s}})$ to denote $\mu_i(\langle s_1, s_2, \dots, s_n \rangle)$, thus

$$\langle s_1, s_2, \dots, \mathbf{s}_i, \dots, s_n \rangle = \langle \mathbf{s}_i, s_{-i}^{\vec{s}} \rangle.$$

In the last vector, although the \mathbf{s}_i is written in the first entry, we mean it to be inserted at the i 'th place.

Definition 3.19 (Best Response to a Profile)

Given a strategy profile

$$s_{-i}^{\rightarrow} = \langle s_1, s_2, \dots, s_{i-1}, s_{i+1}, \dots, s_n \rangle,$$

a **best response of player i to s_{-i}^{\rightarrow}** is any mixed strategy $s_i^* \in S_i$ such that

$$\mu_i(s_i^*, s_{-i}^{\rightarrow}) \geq \mu_i(s_i, s_{-i}^{\rightarrow})$$

for all strategies $s_i \in S_i$.

Is a best response **unique**?

Example 3.20 (Responses for Rochambeau)

How does the set of best responses look like?

- 1 Player 2 plays the pure strategy *paper*.
- 2 Player 2 plays *paper* with probability .5 and *scissors* with probability .5.
- 3 Player 2 plays *paper* with probability $\frac{1}{3}$ and *scissors* with probability $\frac{1}{3}$ and *rock* with probability $\frac{1}{3}$.

Observation

Is a non-pure strategy in the best response set (say a strategy (a_1, a_2) with probabilities $\langle p, 1 - p \rangle$, $p \neq 0$), then so are all other mixed strategies with probabilities $\langle p', 1 - p' \rangle$ where $p \neq p' \neq 0$.

- Consider the set of best responses.
- Either this set is a **singleton** (namely when it consists of a **pure strategy**), or
- the set is **infinite**.

Definition 3.21 (Nash Equilibrium (NE))

A strategy profile $\vec{s}^* = \langle s_1^*, s_2^*, \dots, s_n^* \rangle$ is a **Nash equilibrium** if for any agent i , s_i^* is a best response to $\vec{s}_{-i}^* = \langle s_1^*, s_2^*, \dots, s_{i-1}^*, s_{i+1}^*, \dots, s_n^* \rangle$.

What are the Nash equilibria in the Battle of sexes?
What about the matching pennies?

Example 3.22 (Cuban Missile Crisis, Type 4)

This relates to the well-known crisis in October 1962.

		USSR	
		Withdrawal	Maintain
U. S.	Blockade	(3, 3) Compromise	(2, 4) USSR victory
	Air strike	(4, 2) U.S. victory	(1, 1) Nuclear War

Theorem 3.23 (Nash (1950))

Every finite normal form game has a Nash equilibrium.

Corollary 3.24 (Nash implies maxmin)

In any finite normal form 2-person constant-sum game, the Nash equilibria are exactly all pairs $\langle s_1, s_2 \rangle$ of maxmin strategies (s_1 for player 1, s_2 for player 2).

All Nash equilibria have the same payoff: the value of the game, that player 1 gets.

Proof

We use Kakutani's theorem: Let X be a nonempty subset of n -dimensional Euclidean space, and $f : X \rightarrow 2^X$. The following are sufficient conditions for f to have a fixed point (i.e. an $x^* \in X$ with $x^* \in f(x^*)$):

- 1 X is compact: any sequence in X has a limit in X .
- 2 X is convex: $x, y \in X, \alpha \in [0, 1] \Rightarrow \alpha x + (1 - \alpha)y \in X$.
- 3 $\forall x : f(x)$ is nonempty and convex.
- 4 For any sequence of pairs (x_i, x_i^*) such that $x_i, x_i^* \in X$ and $x_i^* \in f(x_i)$, if $\lim_{i \rightarrow \infty} (x_i, x_i^*) = (x, x^*)$ then $x^* \in f(x)$.

Let X consist of all mixed strategy profiles and let f be **the best response set**: $f(\langle s_1, \dots, s_n \rangle)$ is the set of all best responses $\langle s'_1, \dots, s'_n \rangle$ (where s'_i is player i 's best response to $\langle s_1, \dots, s_{i-1}, s_{i+1}, \dots, s_n \rangle$).

Why is that a subset of an n -dimensional Euclidean space?

A mixed strategy over k actions (pure strategies) is a $k - 1$ dimensional simplex (namely the one satisfying $\sum_{i=1}^k p_i = 1$).

Therefore X is the cartesian product of n simplices. X is compact and convex (why?).

The function f satisfies the remaining properties listed above. Thus there is a fixed point and this fixed point is a Nash equilibrium. □

Theorem 3.25 (Brouwer's Fixed Point Theorem)

Let B^n be the unit Euclidean ball in \mathbb{R}^n and let $f : B^n \rightarrow B^n$ be a continuous mapping. Then **there exists a fixed point of f** : there is a $x \in B^n$ with $f(x) = x$.

Proof

Reduction to the C^1 -differentiable case: Let $r : \mathbb{R}^n \rightarrow \mathbb{R}$ be the function

$$r(x) = r(x_1, x_2, \dots, x_n) = \sqrt{x_1^2 + x_2^2 + \dots + x_n^2}.$$

Let $\phi : \mathbb{R}^n \rightarrow \mathbb{R}$ given by:

$$\phi(x) = a_n(1/4r(x)^4 - 1/2r(x)^2 + 1/4)$$

on D and equal to 0 in the complement of B^n , where the constant a_n is chosen such that the integral of ϕ over \mathbb{R}^n equals 1.

Proof (cont.)

Let $f : B^n \rightarrow B^n$ be continuous. Let $F : \mathbb{R}^n \rightarrow B^n$ be the extension of f to \mathbb{R}^n , for which we have

$$F(x) = f\left(\frac{x}{\|x\|}\right) \text{ on } \mathbb{R}^n \setminus B^n.$$

For $k \in \mathbb{N}, k \neq 0$, put $f_k(x) := \int_{\mathbb{R}^n} k^n \phi(ky) F(x - y) dy$. Show that the restriction of f_k to B^n maps B^n into B^n . Show that the mappings f_k are continuously differentiable and approximate in the topology of uniform convergence the mapping F . Show that if there exists a continuous mapping $f : B^n \rightarrow B^n$ without fixed points, then there will also exist a continuously differential mapping without fixed points. It follows, that it suffices to prove the Brouwer Fixed Point Theorem only for continuously differentiable mappings.

Proof (cont.)

Proof for C^1 -differentiable mappings:

Assume, that the continuously differentiable mapping $f : B^n \rightarrow B^n$ has no fixed points. Let $g : B^n \rightarrow \partial B^n$ the mapping, such that for every point $x \in B^n$ the points $f(x), x, g(x)$ are in that order on a line of \mathbb{R}^n . ∂B^n is the surface of the unit ball B^n , it is also denoted by S^{n-1} . The mapping g is also continuously differentiable and satisfies $g(x) = x$ for $x \in \partial B^n$. We write $g(x) = (g_1(x), g_2(x), \dots, g_n(x))$ and get (for $x \in \partial B^n$ and $i = 1 \dots n$) $g_i(x_1, x_2, \dots, x_n) = x_i$. Note $dg_1 \wedge dg_2 \wedge \dots \wedge dg_n = 0$ since $g_1^2 + g_2^2 + \dots + g_n^2 = 1$. Then:

$$\begin{aligned} 0 &\neq \int_{B^n} dx_1 \wedge dx_2 \wedge \dots \wedge dx_n &= \int_{\partial B^n} x_1 \wedge dx_2 \wedge \dots \wedge dx_n \\ &= \int_{\partial B^n} g_1 \wedge dg_2 \wedge \dots \wedge dg_n &= \int_{B^n} dg_1 \wedge dg_2 \wedge \dots \wedge dg_n \\ &= \int_{B^n} 0 = 0 \end{aligned}$$

Example 3.26 (Majority Voting)

Consider agents 1, 2, 3 and three outcomes A, B, C . Agents vote simultaneously for one outcome (no abstaining). The outcome with most votes wins. If there is no majority, then A is selected. The payoff functions are as follows:

$\mu_1(A) = \mu_2(B) = \mu_3(C) = 2$, $\mu_1(B) = \mu_2(C) = \mu_3(A) = 1$ and $\mu_1(C) = \mu_2(A) = \mu_3(B) = 0$. What are the Nash equilibria and what are their outcomes?

Example 3.27 (Unique Equilibrium)

The following game has exactly one Nash equilibrium.

	L	C	R
U	$\langle 1, -2 \rangle$	$\langle -2, 1 \rangle$	$\langle 0, 0 \rangle$
M	$\langle -2, 1 \rangle$	$\langle 1, -2 \rangle$	$\langle 0, 0 \rangle$
D	$\langle 0, 0 \rangle$	$\langle 0, 0 \rangle$	$\langle 1, 1 \rangle$

Minmax value and feasible payoffs

Definition 3.28 (Minmax value and feasible payoffs)

In a n -person normal form game $\langle \mathbf{A}, \text{Act}, O, \varrho, \mu \rangle$ we define the **minmax value** of player i as follows

$$\underline{v}_i = \min_{s_{-i}} \max_{s_i} \mu_i^{\text{expected}}(s_i, s_{-i}).$$

We call a payoff profile $\langle r_1, \dots, r_i, \dots, r_n \rangle$ **feasible** if there exist rational values $\alpha_{\vec{a}} \geq 0$ such that for all i

$$r_i = \sum_{\vec{a} \in \text{Act}} \alpha_{\vec{a}} \mu_i(\vec{a}) \text{ where } \sum_{\vec{a} \in \text{Act}} \alpha_{\vec{a}} = 1.$$

Geometrically speaking, this is the (rational) convex hull of all possible payoffs of pure action profiles.

Feasible payoffs

What is the idea behind feasible payoffs?

- The minmax value is important, as this is the **minimal payoff in any Nash equilibrium**.
- Not any feasible profile can be obtained as payoff. In the leader game in Example 3.13 on Slide 38, the profile $\langle \frac{7}{2}, \frac{7}{2} \rangle$ is feasible but can not be realised in one game. **What if the game is played infinitely often (or just twice)?**
- In general, convex combinations of pure-strategy payoffs can only be obtained by **correlated** strategies, not by **independent randomizations**.

In fact, feasible payoffs (and their convexity) will play a role later in Section ?? on Slide ?? (Theorem ??). They can be realized in **repeated games**.

It is obvious that **there is not always** a strategy that is **strictly dominating** all others (this is why the Nash equilibrium has been introduced).

Reducing games

However, often games can be **reduced** and the computation of the equilibrium considerably simplified.

A rational player would never choose a strategy that is strictly dominated.

Definition 3.29 (Dominating Strategy: Weakly, Strictly)

A pure strategy a_i is **strictly dominated** for an agent i , if there exists some other (mixed) strategy s'_i that strictly dominates it, i.e. for all profiles

$\vec{a}_{-i} = \langle a_1, \dots, a_{i-1}, a_{i+1}, \dots, a_n \rangle$, we have

$$\mu_i(\langle s'_i, \vec{a}_{-i} \rangle) \geq \mu_i(\langle a_i, \vec{a}_{-i} \rangle).$$

We say that a pure strategy a_i is **weakly dominated** for an agent i , if in the above inequality we have \geq instead of \geq and the inequality is strict for **at least one** of the other \vec{a}_{-i} .

Definition 3.30 (Reduced Sets A_i^∞, S_i^∞)

For an arbitrary normal form game with A_i the set of pure strategies and S_i the set of mixed strategies for agent i , we define ($A_i^0 := A_i, S_i^0 := S_i$)

$$A_i^n := \{a_i \in A_i^{n-1} \mid \text{there is no } s'_i \in S_i^{n-1} \text{ s.t.} \\ \mu_i(s'_i, a_{-i}^{\rightarrow}) \geq \mu_i(a_i, a_{-i}^{\rightarrow}) \\ \text{for all } a_{-i}^{\rightarrow} \in A_{-i}^{n-1}\}$$

$$S_i^n := \{s' \in S_i \mid s'(a_i) \geq 0 \text{ only if } a_i \in A_i^n\}$$

Finally, $A_i^\infty := \bigcap_{n=0}^\infty A_i^n$ and S_i^∞ is the set of all **mixed strategies** s_i , such that there is no s'_i with $\mu_i(s'_i, a_{-i}^{\rightarrow}) \geq \mu_i(s_i, a_{-i}^{\rightarrow})$ for all $a_{-i}^{\rightarrow} \in A_{-i}^\infty$.

Some Comments to A_i^∞ and S_i^∞

- What are the S_i^n ? They are the sets of mixed strategies over only the **pure** strategies in A_i^n .
- The A_i^n are sets of pure strategies from which we remove those, that are **strictly dominated** by **certain** other mixed strategies.
- Therefore only the strategies in A_i^∞ are those that we have to keep.
- Note that S_i^∞ is defined wrt. A_i^∞ , not wrt. S_i^n .
- S_i^∞ is the set of mixed strategies that are **not strictly dominated** by pure action profiles from A_i^∞ .
- Note that S_i^∞ can be strictly smaller than the set of all mixed strategies over A_i^∞ .

Theorem 3.31 (Solvable by Iterated Strict Dominance)

*If for a finite normal form game, the sets A_i^∞ are all singletons (such a game is called **solvable by iterated strict dominance**), then this strategy profile is the **unique Nash equilibrium**.*

Lemma 3.32 (Church-Rosser)

Given a 2-person normal form game. All **strictly dominated** columns, as well as all **strictly dominated** rows can be **eliminated** without changing the Nash equilibria (or similar solution concepts). This results in a **finite series of reduced games**. The final result **does not depend on the order of the eliminations**.

Note: the last lemma is not true for **weakly** dominated strategies. There, the order **does** matter.

Note that we eliminate only **pure** strategies. Such a strategy might be **dominated** by a **mixed** strategy.

	L	C	R
U	$\langle 3, 2 \rangle$	$\langle 2, 1 \rangle$	$\langle 3, 1 \rangle$
M	$\langle 1, 1 \rangle$	$\langle 1, 1 \rangle$	$\langle 2, 2 \rangle$
D	$\langle 0, 1 \rangle$	$\langle 4, 2 \rangle$	$\langle 0, 1 \rangle$

- 1** Eliminate row **M**.
- 2** Eliminate column **R**.

This leads to

	L	C
U	$\langle 3, 2 \rangle$	$\langle 2, 1 \rangle$
D	$\langle 0, 1 \rangle$	$\langle 4, 2 \rangle$

Elimination of Weakly Dominated actions

We consider the normal form game

	L	R
T	$\langle 1, 1 \rangle$	$\langle 0, 0 \rangle$
M	$\langle 1, 1 \rangle$	$\langle 2, 1 \rangle$
B	$\langle 0, 0 \rangle$	$\langle 2, 1 \rangle$

- 1 If we first eliminate **T**, and then **L** we get the outcome $\langle 2, 1 \rangle$.
- 2 If we first eliminate **B**, and then **R** we get the outcome $\langle 1, 1 \rangle$.



3.3 Extensive Form Games

We have previously introduced **normal form games** (Definition 3.4 on Slide 25). This notion does not allow to deal with sequences of actions that are **reactions to actions** of the opponent.

Extensive form (tree form) games

Unlike games in normal form, those in **extensive form** do not assume that all moves between players are made simultaneously. This leads to a **tree form**, and allows to introduce **strategies**, that take into account the **history** of the game.

We distinguish between **perfect** and **imperfect** information games. While the former assume that the players have **complete** knowledge about the game, the latter do not: a player might **not know** exactly which node it is in.

The following definition covers a game as a tree:

Definition 3.33 (Perfect Extensive Form Games)

A **finite perfect information game in extensive form** is a tuple $\Gamma = \langle \mathbf{A}, \text{Act}, H, Z, \alpha, \rho, \sigma, \mu_1, \dots, \mu_n \rangle$ where

- \mathbf{A} is a set of n players, Act is a set of actions
- H is a set of non-terminal nodes, Z a set of terminal nodes, $H \cap Z = \emptyset$, $H \cup Z$ form a **tree**,
- $\alpha : H \rightarrow 2^{\text{Act}}$ assigns to each node a set of actions,
- $\rho : H \rightarrow \mathbf{A}$ assigns to each non-terminal node a player who chooses an action at that node,
- $\sigma : H \times \text{Act} \rightarrow H \cup Z$ assigns to each (node,action) a successor node ($h_1 \neq h_2$ implies $\sigma(h_1, a_1) \neq \sigma(h_2, a_2)$),
- $\mu_i : Z \rightarrow \mathbb{R}$ are the utility functions.

Such games can be visualised as trees. Here is the famous “Sharing Game”.

Example 3.34 (Sharing Game)

The game consists of two rounds. In the first, player 1 offers a certain share (namely (1) 2 for player 1, 0 for player 2, (2) 1 for player 1, 1 for player 2, (3) 0 for player 1, 2 for player 2). Player 2 can only accept, or refuse. In the latter case, nobody gets anything.




Figure 2: The Sharing game.

Strategies in extensive form games

Definition 3.35 (Strategies in Extensive Form Games)

Let $\Gamma = \langle \mathbf{A}, \text{Act}, H, Z, \alpha, \rho, \sigma, \mu_1, \dots, \mu_n \rangle$ be a finite perfect information game in extensive form.

A **strategy** for player i in Γ is any **function** that assigns a legal move to each history owned by i .

The **pure strategies** of player i are the elements of $\prod_{h \in H, \rho(h)=i} \alpha(h)$. These are also functions: whenever player i can do a move, it chooses one of the actions available. Thus we can write a pure strategy as a vector $\langle a_1, \dots, a_r \rangle$, where a_1, \dots, a_r are i 's choices at the respective moves.

In the sharing game, a pure strategy for player 2 is $\langle \text{no}, \text{yes}, \text{no} \rangle$. A (better) one is $\langle \text{no}, \text{yes}, \text{yes} \rangle$.

Why don't we introduce **mixed** strategies?

Best response, Nash Equilibrium

Note that the definitions of **best response** and **Nash equilibrium** carry over (literally) to games in extensive form.

Note that in the following we are talking only about **pure strategy profiles**.

What are the NE's in the sharing game?

$\langle 1, \langle y, y, y \rangle \rangle$, $\langle 1, \langle n, n, n \rangle \rangle$, $\langle 1, \langle n, n, y \rangle \rangle$, $\langle 1, \langle y, n, n \rangle \rangle$,
 $\langle 1, \langle y, n, y \rangle \rangle$, $\langle 1, \langle y, y, n \rangle \rangle$ are NE's, $\langle 1, \langle n, y, y \rangle \rangle$ is not. Also
 $\langle 2, \langle n, y, n \rangle \rangle$, $\langle 2, \langle n, y, y \rangle \rangle$ and $\langle 3, \langle n, n, y \rangle \rangle$ are NE's.

We claim that only $\langle 1, \langle y, y, y \rangle \rangle$ and $\langle 2, \langle n, y, y \rangle \rangle$ make sense.

Transforming extensive form games into normal form

Lemma 3.36 (Extensive form \leftrightarrow Normal form)

*Each game Γ in **perfect information extensive form** can be transformed to a game $\text{NF}(\Gamma)$ in **normal form** (such that the pure strategy spaces correspond).*

Proof.

A **strategy profile** determines a **unique** path from the root \emptyset of the game to one of the terminal nodes (and hence also a single profile of payoffs). Therefore one can construct the corresponding normal form game $NF(\Gamma)$ by **enumerating** all strategy profiles and filling the payoff matrix with the resulting payoffs. □

Sharing Game in normal form

1 \ 2	nnn	nny	nyn	nyy	ynn	yny	yyn	yyy
(2, 0)	(0, 0)	(0, 0)	(0, 0)	(0, 0)	(2, 0)	(2, 0)	(2, 0)	(2, 0)
(1, 1)	(0, 0)	(0, 0)	(1, 1)	(1, 1)	(0, 0)	(0, 0)	(1, 1)	(1, 1)
(0, 2)	(0, 0)	(0, 2)	(0, 0)	(0, 2)	(0, 0)	(0, 2)	(0, 0)	(0, 2)




Figure 3: A generic game.

Example 3.37 (Generic Game in normal form)

We consider the game in Figure 3. The pure strategies of player 1 are $\{\langle A, E \rangle, \langle A, F \rangle, \langle B, E \rangle, \langle B, F \rangle\}$. The pure strategies of player 2 are $\{C, D\}$.

	<i>C</i>	<i>D</i>
<i>AE</i>	<i>W</i>	<i>Y</i>
<i>AF</i>	<i>X</i>	<i>Y</i>
<i>BE</i>	<i>Z</i>	<i>Z</i>
<i>BF</i>	<i>Z</i>	<i>Z</i>

Note that $\langle B, E \rangle, \langle B, F \rangle$ are pure strategies that have to be considered.



Is there a converse of Lemma 3.36?

We consider prisoner's dilemma and try to model a game in extensive form with the same payoffs and strategy profiles.

In fact, it is not surprising that we do not succeed in the general case:

Theorem 3.38 (Zermelo, 1913; Kuhn)

*For each perfect information game in extensive form **there exists** a pure strategy NE.*

*The theorem can be strengthened (Kuhn's theorem): For each perfect information game in extensive form **there exists** a pure strategy **subgame perfect** NE.*

In fact, this was the reason that we do not need mixed strategies for perfect information extensive games (question on Slide 81).

We will later introduce **imperfect information games (in extensive form): Slide 101.**

Example 3.39 (Unintended Nash equilibria)

Consider the following game in extensive form.




Figure 4: Unintended Equilibrium.

The game depicted in Example 3.39 has two equilibria: $\langle A, R \rangle$ and $\langle B, L \rangle$. The latter one is not intuitive (while the first one is).

Can we refine the notion of NE and rule out this unintended equilibrium?

This leads to the notion of **subgame perfect NE**:

Definition 3.40 (Subgame Perfect NE (SPE))

Let Γ be a perfect information game in extensive form.

Subgame: A **subgame of G rooted at node h** is the restriction of Γ to the descendants of h .

SPE: The **subgame perfect Nash equilibria (SPE)** of a perfect information game Γ in extensive form are those Nash equilibria of Γ , that are also Nash equilibria for all subgames Γ' of Γ .

- What are the SPE's in the Sharing game (Example 3.34)?
- What are the SPE's in the following instance of the generic game (Example 3.37):

	<i>C</i>	<i>D</i>
<i>AE</i>	$\langle 2, 0 \rangle$	$\langle 1, 1 \rangle$
<i>AF</i>	$\langle 0, 2 \rangle$	$\langle 1, 1 \rangle$
<i>BE</i>	$\langle 3, 3 \rangle$	$\langle 3, 3 \rangle$
<i>BF</i>	$\langle 3, 3 \rangle$	$\langle 3, 3 \rangle$

Theorem 3.41 (Existence of SPE (Kuhn))

For each **finite perfect information** game in extensive form **there exists a SPE**.

The proof is by induction on the length of histories. The SPE is therefore defined constructively.

Proof.


By backward induction we construct a subgame perfect NE:

- Let h be a terminal history and h' be the history with the last action removed, say $h = h'a$. Moreover, let it be player i 's move in h' . Then, we define $s_i(h')$ to be the action which maximizes i 's payoff. We proceed like this for all such histories h' .
- Suppose now that s is a subgame perfect NE for all histories of a certain length, say k . Consider a history $h' = ha$ of length $k + 1$. As before the player whose move it is in h' chooses an action which maximizes its payoff assuming that all other players follow s . We proceed like this for all histories h' .
- The constructed strategy s is a subgame perfect NE.



Example 3.42 (Centipede Game)

This is a two person game which illustrates that even the notion of SPE can be critical.



Centipede Revisited

- The Centipede game has just one SPE: **All players always choose D .**
- This is rational, but humans often do not behave like that.
- Experiments show, that humans start with going **across** and do a **down** only towards the end of the game.

Imperfect Information

- In an extensive game with **perfect** information, the player does know all previous moves (and also the payoffs that result).
- In an extensive game with **imperfect** information, a player might not be completely informed about the past history.
- Or some moves in the past may have been **done randomly**, so in the future, even under the same circumstances, other actions might be taken.

- An **extensive game** is nothing else than a **tree**. Thus each node is unique and carries with it the path from the root (the history that lead to it).
- In order to model that a player does not perfectly know the past events, we introduce an **equivalence relation** on the nodes. That two nodes are **equivalent**, means that the **player cannot distinguish** between them.
- All nodes in one equivalence class must be **assigned the same actions**: otherwise the player could distinguish them.

Definition 3.43 (Information set I_i , Partition)

For a set W (nodes, worlds, games) and a set of agents \mathbf{A} , we define a **partition I_i of agent $i \in \mathbf{A}$ over it** (or the **information set** of i) as an **equivalence relation** over W . Its classes I_{ij} are also called **partition classes**. Thus the following holds.


A **partition I_i** is a set of subsets W_{i_1}, \dots, W_{i_s} of W such that: (1) $\bigcup_i W_{ij} = W$ and (2) $W_{ij} \cap W_{ij'} = \emptyset$ for $j \neq j'$.

Definition 3.44 (Extensive Games, Imperfect Inf.)

A **finite imperfect information game in extensive form** is a tuple $G = \langle \mathbf{A}, \text{Act}, H, Z, \alpha, \rho, \sigma, \mu_1, \dots, \mu_n, I_1, \dots, I_n \rangle$ where

- $\langle \mathbf{A}, \text{Act}, H, Z, \alpha, \rho, \sigma, \mu_1, \dots, \mu_n \rangle$ is a perfect information game in the sense of Definition 3.33 on Slide 77,
- I_i are partitions on $\{h \in H : \rho(h) = \mathbf{i}\}$ such that $h, h' \in I_{i,j}$ implies $\alpha(h) = \alpha(h')$.

Example 3.45



(Utility of the rightmost leaf node is $(2, 1)$.)

- Player 1 can not distinguish between the nodes connected by a dotted line.
- Therefore player 1 can not play the **right** move.
- It could play a mixed strategy: with probability $\frac{1}{2}$ choose l .

Now we need mixed strategies, to deal with the uncertainty.

Definition 3.46 (Pure strategy in Extensive Form)

Given an imperfect information game in extensive form, a **pure strategy for player i** is a vector $\langle a_1, \dots, a_k \rangle$ with $a_j \in \alpha(I_{ij})$ where I_{i1}, \dots, I_{ik} are the k equivalence classes for agent i . Note that this vector is just a function assigning an action to each node owned by player i .

Can we model prisoner's dilemma as an extensive game with imperfect information?




Figure 5: Prisoner's dilemma in extensive form.

There is a pure strategy Nash equilibrium.

But we could have chosen to switch player 1 with player 2.

NF game \leftrightarrow Imperfect game

For pure strategies we have the following:

- Each **game in normal form** can be transformed into an **imperfect information game in extensive form** (but this is not one-to-one).
- Each imperfect information game in extensive form can be transformed into a game in normal form (this is one-to-one).

What are **mixed** strategies for an imperfect information game?

Mixed Strategy: First try.

- Let Γ be an imperfect information game in extensive form.
- Assign the normal form game (for any i) as usual, by enumerating the pure strategies.
- Now we can take the usual set of mixed strategies **in the normal form game** as the set of **mixed strategies of the original game Γ** .

Behavioral Strategy: Second try.

- We consider the game from Figure 3 on Slide 86.
- Consider the following strategy for player 1. A is chosen with probability $.7$, B with $.3$ and E with probability $.4$ and F with $.6$. Such strategies are called **behavioral**: at each node, the (probabilistic) choice is made **independently** from the other nodes.
- Consider the following **mixed** strategy for player 1. $\langle A, E \rangle$ is chosen with probability $.6$ and $\langle B, F \rangle$ with probability $.4$. Thus, here we have a strong **correlation**: $\langle A, F \rangle$ is not possible!

Mixed vs. Behavioral

Definition 3.47 (Mixed and Behavioral strategies)

Let $G = \langle \mathbf{A}, \text{Act}, H, Z, \alpha, \rho, \sigma, \mu_1, \dots, \mu_n, I_1, \dots, I_n \rangle$ be an imperfect information game in extensive form.

Mixed: A **mixed** strategy of player i is **one single** probability distribution **over i 's pure strategies**.

Behavioral: A **behavioral** strategy of player i is a **vector** of probability distributions $P(I_{ij})$ over the set of actions $\alpha(I_{ij})$ **for $I_{ij} \in I_i$** . We define by $P(h)(a)$ the probability $P(I_{ij})(a)$ for the action a for player i if $h \in I_{ij}$.

Mixed vs. Behavioral (2)

- The main difference is that for **behavioral** strategies, at each node the probability distribution is started freshly.
- Even if a player ends up in the same partition, she can choose independently of her previous choice.
- Whereas for **mixed** strategies, this choice is **not** independent: there is **just one single** distribution that relates the possible choices.


Are behavioral strategies more general?

We consider a one player game. At the start node, the player can choose L or R . There result two nodes which can not be distinguished by the player. Again, L or R can be played and result in the four outcomes o_1, o_2, o_3, o_4 .

- What is the outcome of the mixed strategy $\langle \frac{1}{2}\mathbf{LL}, \frac{1}{2}\mathbf{RR} \rangle$?
- It is $\langle \frac{1}{2}, 0, 0, \frac{1}{2} \rangle$
- **No behavioral strategy** results in this distribution.
- Therefore mixed strategies **are not necessarily** behavioral.

Example 3.48 (A game of imperfect recall)

We consider the following game



For **mixed** strategies, $\langle R, D \rangle$ is the unique NE. But for **behavioral** strategies, the following mixed strategy is a better response of player 1 to D : $(\frac{98}{198}L, \frac{100}{198}R)$.

Are mixed strategies more general? (2)

- For mixed strategies, once decided, the pure strategy is consistently chosen. Therefore the outcome $\langle 100, 100 \rangle$ is **not reachable**.
- This is not true for behavioral strategies, where at **each node**, the probabilistic choice is done **independently**.
- **What is the best response of player 1 to D in behavioral strategies?** Consider a mixed behavioral strategy: choose L with probability p and R with $1 - p$. A little computation shows that the maximal payoff is obtained for $p = \frac{98}{198}$.

Behavioral vs. mixed strategies?

- We have just seen that there are mixed strategies for which there are no behavioral strategies with the same outcome and vice versa.
- Therefore we introduce two concepts of Nash equilibria on the next page.
- Is there a **class of games** where both concepts are **equivalent**?

Definition 3.49 (NE for Mixed/Behavioral Strategies)

A **NE in mixed strategies** for an extensive game G is a mixed strategy profile $s^* = \langle s_1^*, s_2^*, \dots, s_n^* \rangle$, s.t. for any agent i :

$$\mu_i(\langle s_i^*, s_{-i}^* \rangle) \geq \mu_i(\langle s_i, s_{-i}^* \rangle)$$

for all **mixed** strategies s_i of player i .

A **NE in behavioral strategies** is a behavioral strategy profile $s^* = \langle s_1^*, s_2^*, \dots, s_n^* \rangle$, s.t. for any agent i :

$$\mu_i(\langle s_i^*, s_{-i}^* \rangle) \geq \mu_i(\langle s_i, s_{-i}^* \rangle)$$

for all **behavioral** strategies s_i of player i .




Definition 3.50 (Perfect Recall)

Let Γ be an imperfect information game in extensive form. We say that player i has **perfect recall** in Γ , if the following holds. If h, h' are two nodes in the same I_{ij} (for a j), and $h_0, a_0, h_1, a_1, \dots, h_n, a_n, h$ resp. $h'_0, a'_0, h'_1, a'_1, \dots, h'_m, a'_m, h'$ are paths from the root of the tree to h (resp. h'), then



- 1 $n = m$,
- 2 for all $0 \leq j \leq n$: h_j and h'_j are in the same partition class,
- 3 for all $0 \leq j \leq n$: if $\alpha(h_j) = i$ then $a_j = a'_j$.

Γ is a **game of perfect recall**, if all players have perfect recall. Otherwise it is called of **imperfect recall**.

A few games: Which have perfect recall?



Do they model the same situation?



Perfect Recall: Behavioral strategies suffice?

Theorem 3.51 (Behavioral = Mixed (Kuhn, 1953))

Let Γ be a **game of perfect recall** (perfect or imperfect information). Then for any mixed strategy of agent i there is a behavioral one such that both strategies induce the same probabilities on outcomes for all fixed strategy profiles of the other agents.


Corollary 3.52

In a game of perfect recall, it suffices to compute the Nash equilibria **based on behavioral strategies**.

SPE: What about **subperfect** equilibria (analogue of Definition 3.40 on Slide 92 for imperfect games?)

First try: In each information set, we have a set of subgames (a **forest**). Why not asking that a strategy should be a **best response in all subgames of that forest**?

Example 3.53 (A Game with no SPE's)



- Nash equilibria: (L, U) and (R, D) . Can we see them as **subgame perfect**?
- In one subtree, U dominates D , in the other D dominates U .
- **But (R, D) seems to be the unique choice: both players can put themselves into the others place and reason accordingly.**
- **Requiring that a strategy is best response to all subtrees might be too strong.**

There are two prominent refinements of SPE's.

- One is the **trembling hand perfect** equilibrium. It is defined for normal form, and extensive form games.
- The other the **sequential** equilibrium is defined for extensive form games of perfect recall.

We use a **belief system** μ : functions which assign to each information set I_{ij} a probability measure over nodes in I_{ij} . $\mu(I_{ij})(h)$ is the probability that player **i** assigns to history h , provided that I_{ij} is being reached.

So, a belief system captures the **probability of being in a specific node** of an information set.

Given a behavioral strategy profile β which is **completely mixed** (i.e. assigns non-zero probability to all actions), a belief system μ can be **uniquely assigned to β** .

Definition 3.54 (Sequential Equilibrium)

A (behavioral) strategy profile $\beta^* = \langle \beta_1^*, \beta_2^*, \dots, \beta_n^* \rangle$ is a **sequential equilibrium** of an extensive form game Γ if there exist probability distributions $\mu(h)$ for each information set I_i such that

- 1 $(\beta^*, \mu) = \lim_{n \rightarrow \infty} (\beta^n, \mu^n)$ for some sequence where β^n is completely mixed (where μ^n is uniquely determined by β^n),
- 2 for any I_i of agent i and any alternative strategy β'_i of agent i : $\mu_i(\beta^* | h, \mu(h)) \geq \mu_i((\beta', \beta_{-i}) | h, \mu(h))$.

The first assumption (consistency) is of a rather technical nature to enable the consistent definition of expectations in case of behavioral strategies which assign probability 0 to actions.

Theorem 3.55 (Sequential Equilibrium)

*For each imperfect information **game in extensive form with perfect recall** there **exists a sequential equilibrium**.*

For perfect information games, each SPE is a sequential equilibrium but not vice versa.



3.4 An Example from Economics

A theory for efficiently allocating goods and resources among agents, based on market prices.

Goods: Given $n > 0$ goods g_1, \dots, g_n (coffee, mirror sites, parameters of an airplane design). We assume $g_i \neq g_j$ for $i \neq j$ but within each g_i everything is indistinguishable.

Prices: The market has prices $\mathbf{p} = [p_1, p_2, \dots, p_n] \in \mathbb{R}^n$: p_i is the price of the good i .

Consumers: Consumer i has $\mu_i(\mathbf{x})$ encoding its preferences over consumption bundles $\mathbf{x}_i = [x_{i1}, \dots, x_{in}]^t$, where $x_{ig} \in \mathbb{R}^+$ is consumer i 's allocation of good g . Each consumer also has an initial endowment $\mathbf{e}_i = [e_{i1}, \dots, e_{in}]^t \in \mathbb{R}$.

Producers: Use some commodities to produce others: $\mathbf{y}_j = [y_{j1}, \dots, y_{jn}]^t$, where $y_{jg} \in \mathbb{R}$ is the amount of good g that producer j produces. \mathbf{Y}_j is a set of such vectors \mathbf{y} .
Profit of producer j : $\mathbf{p} \times \mathbf{y}_j$, where $\mathbf{y}_j \in \mathbf{Y}_j$.

Profits: The profits are divided among the consumers (given predetermined proportions Δ_{ij}): Δ_{ij} is the fraction of producer j that consumer i owns (stocks). Profits are divided according to Δ_{ij} .

Definition 3.56 (General Equilibrium)

$(\mathbf{p}^*, \mathbf{x}^*, \mathbf{y}^*)$ is in **general equilibrium**, if the following holds:

- I. The markets are in equilibrium:

$$\sum_i \mathbf{x}_i^* = \sum_i \mathbf{e}_i + \sum_j \mathbf{y}_j^*$$

- II. Producer j maximises profit wrt. the market

$$\mathbf{y}_j^* = \arg \max_{\{\mathbf{y}_j \in \mathbf{Y}_j\}} \mathbf{P}^* \times \mathbf{y}_j$$

III. Consumer i maximises preferences according to the prices

$$\mathbf{x}_i^* = \arg \max_{\{\mathbf{x}_i \in \mathbb{R}_+^n \mid \text{cond}_i\}} \mu_i(\mathbf{x}_i)$$

where cond_i stands for

$$\mathbf{p}^* \times \mathbf{x}_i \leq \mathbf{p}^* \times \mathbf{e}_i + \sum_j \Delta_{ij} \mathbf{p}^* \times \mathbf{y}_j.$$

Theorem 3.57 (Pareto Efficiency)

Each general equilibrium is pareto efficient.

Theorem 3.58 (Coalition Stability)

*Each general equilibrium **with no producers** is coalition-stable: **no subgroup can increase their utilities by deviating from the equilibrium and building their own market.***

Theorem 3.59 (Existence of an Equilibrium)

Let the sets Y_j be closed, convex and bounded above. Let μ_i be continuous, strictly convex and strongly monotone. Assume further that at least one bundle \mathbf{x}_i is producible with only positive entries x_{il} .

Under these assumptions a general equilibrium exists.

Meaning of the assumptions

Formal definitions: \rightsquigarrow **blackboard**.

Convexity of Y_j : Economies of scale in production do not satisfy it.

Continuity of the μ_i : Not satisfied in bandwidth allocation for video conferences.

Strictly convex: Not satisfied if **preference increases when one gets more of this good** (drugs, alcohol, dulce de leche).

In general, there exist more than one equilibrium.

Theorem 3.60 (Uniqueness)

*If the society-wide demand for each good is **non-decreasing in the prices** of the other goods, then a **unique equilibrium exists**.*

This condition is called **gross substitutes property** and comes in many variants.

Positive example: increasing price of meat forces people to eat potatoes (pasta).

Negative example: increasing price of bread implies that the butter consumption decreases.

How to find market equilibria?

We describe an algorithm using **steepest descent**.

Theorem 3.61

The **price tâtonnement algorithm**, explained on the next few pages, **converges to a general equilibrium** if for all \mathbf{p} that are not proportional to an equilibrium vector \mathbf{p}^* , the following holds:

$$\sum_i (\mathbf{x}_i(\mathbf{p}) - \mathbf{e}_i) - \sum_j \mathbf{y}_j(\mathbf{p}) \succeq 0$$

Price tâtonnement process (1)

- This is a **decentralized** algorithm that performs steepest descent (can be improved by Newton-method).
- The main part is a **price adjustor**, that suggests a price and receives production plans from the producers and consumption plans from the consumers.
- Based on these plans, a new price is calculated and the cycle starts again.
- Producer **j** takes the current price and develops a production plan maximizing its profit. This plan is sent to the adjustor.
- Consumer **i** takes the current price and production plans from the producers and develops a consumption plan maximizing its utility given budget constraints. This plan is sent to the adjustor.

Price tâtonnement: The adjustor

for $g = 1$ **to** n **do**

$p_g \leftarrow 1$

end for

for $g = 1$ **to** n **do**

$\lambda_g \leftarrow$ a positive number

end for

repeat

Broadcast \mathbf{p} to consumers and producers.

Receive a production plan \mathbf{y}_j from each producer j .

Broadcast the plans \mathbf{y}_j to consumers.

Receive a consumption plan \mathbf{x}_i from each consumer i .

for $g = 1$ **to** n **do**

$p_g \leftarrow p_g + \lambda_g (\sum_i (x_{ig} - e_{ig}) - \sum_j y_{jg})$

end for

until $|\sum_i (x_{ig} - e_{ig}) - \sum_j y_{jg}| \leq \epsilon$ for all $1 \leq g \leq n$

Inform consumers and producers that an equilibrium has been reached.

Price tâtonnement: Consumer i

repeat

Receive \mathbf{p} from the adjustor.

Receive a production plan \mathbf{y}_j for all j from the adjustor.

Announce to the adjustor a consumption plan \mathbf{x}_i that maximizes i 's utility given the budget constraint (see Condition III on Slide 132).

until Informed by adjustor that equilibrium has been reached.

Exchange and consume.

Price tâtonnement: Producer j

repeat

Receive \mathbf{p} from the adjustor.

Receive a production plan for all j from the adjustor.

Announce to the adjustor a production plan $\mathbf{y}_j \in Y_j$
that maximizes $\mathbf{p} \times \mathbf{y}_j$

until Informed by adjustor that equilibrium has been
reached.

Exchange and produce.



3.5 References



Fudenberg, D. and J. Tirole (1991).

Game Theory.

MIT Press.



Nash, J. (1950).

Equilibrium points in n-person games.

Proceedings of the National Academy of Sciences of the United States of America 36, 48–49.



Osborne, M. and A. Rubinstein (1994).

A Course in Game Theory.

MIT Press.



4. MAS: Basic Concepts

4 MAS: Basic Concepts



5. Agent-Oriented Programming

5 Agent-Oriented Programming

6. Coalitional Games

6 Coalitional Games

- Coalition Formation in CFG's
- General Contract Nets
- Classes of Games
- The Core and its refinements
- Payoff Division: Shapley value and Banzhaf Index

Outline

What happens if agents decide to **team up** and work together to solve a problem more efficiently? We consider

- **abstract coalition formation** for **characteristic function games (CFG)**;
- algorithms for searching the **coalition structure graph**;
- the **task allocation problem** and present different types of contracts between agents (**no IR-contract leads to the global optimum**, even if all types are allowed), and
- how to distribute the profit among the agents: **core of a CFG** and **Shapley value**.



6.1 Coalition Formation in CFG's

Definition 6.1 (Characteristic Function Game (CFG))

A **characteristic function game** is a tuple $\langle \mathbf{A}, \mathbf{v} \rangle$ where \mathbf{A} is a finite set (of agents) and $\mathbf{v} : 2^{\mathbf{A}} \rightarrow \mathbb{R}_0^+$; $S \mapsto \mathbf{v}(S)$.

We assume $\mathbf{v}(\emptyset) = 0$ and call $\mathbf{v}(S)$ the **value of coalition S** .

Thus the value is independent of the nonmembers. But

- 1 **Positive Externalities:** Overlapping goals.
Nonmembers perform actions and move the world closer to the coalition's goal state.
- 2 **Negative Externalities:** Shared resources.
Nonmembers may use up the resources.

Coalition Structure

Definition 6.2 (Coalition Structure CS)

A **coalition structure** CS over the set \mathbf{A} is any partition $\{C^1, \dots, C^k\}$ of \mathbf{A} , i.e. $\bigcup_{j=1}^k C^j = \mathbf{A}$ and $C^i \cap C^j = \emptyset$ for $i \neq j$. We denote by \mathbf{CS}_M the set of all coalition structures CS over the set $M \subseteq \mathbf{A}$.

Finally, we define the **social welfare** of a coalition structure CS by

$$v(\mathbf{CS}) := \sum_{C \in \mathbf{CS}} v(C).$$

Definition 6.3 (Coalition Formation in CFG's)

Coalition Formation in CFG's consists of:

Forming CS: formation of coalitions such that within each coalition agents coordinate their activities.

Solving Optimisation Problem: For each coalition in a CS the tasks and resources of the agents have to be pooled. **Maximise monetary value.**

Payoff Division: **Divide the value** of the generated solution among agents.

Maximising Social Welfare

Maximise the social welfare of the agents **A** by finding a coalition structure

$$CS^* = \arg \max_{CS \in \mathcal{CS}_A} v(CS),$$

where

$$v(CS) := \sum_{S \in CS} v(S).$$

How many coalition structures are there?

A lot: $\Omega(|\mathbf{A}|^{\frac{|\mathbf{A}|}{2}})$. Enumerating is feasible for $|\mathbf{A}| < 15$.




Figure 6: Number of Coalition (Structures).

Approximation of $v(\text{CS})$.

How can we approximate $v(\text{CS})$?

Choose set $\mathbf{N} \subseteq \mathbf{CS}_A$ and pick the best coalition seen so far:

$$\mathbf{CS}_N^* = \arg \max_{\text{CS} \in \mathbf{N}} v(\text{CS}).$$

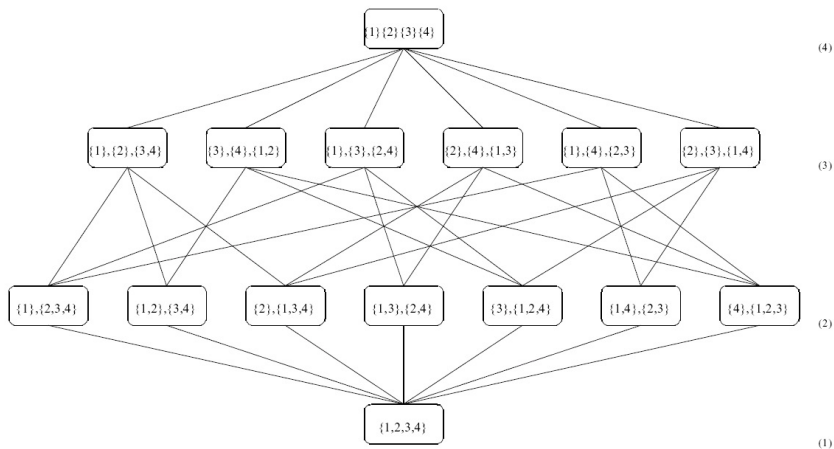


Figure 7: Coalition Structure Graph CS_A .

We want our approximation as good as possible.

We want to find a small k and a small N such that

$$\frac{v(\mathbf{CS}^*)}{v(\mathbf{CS}_N^*)} \leq k.$$

k is the **bound** (best value would be 1) and N is the **part of the graph** that we have to search exhaustively.

We consider 3 search algorithms:

MERGE: Breadth-first search **from the top**.

SPLIT: Breadth first **from the bottom**.

Coalition-Structure-Search (CSS1): First the bottom 2 levels are searched, then a breadth-first search from the top.

MERGE might not even get a bound, without looking at **all** coalitions.

SPLIT gets a good bound ($k = |\mathbf{A}|$) after searching the bottom 2 levels (see below). But then it can get slow.

CSS1 combines the good features of MERGE and SPLIT.

Theorem 6.4 (Minimal Search to get a bound)

To bound k , it suffices to **search the lowest two levels** of the CS-graph. Using this search, the bound $k = |\mathbf{A}|$ can be taken. This bound is tight and the number of nodes searched is $2^{|\mathbf{A}|-1}$.

No other search algorithm can establish the bound $|\mathbf{A}|$ while searching through less than $2^{|\mathbf{A}|-1}$ nodes.

Proof.

There are at most $|\mathbf{A}|$ coalitions included in \mathbf{CS}^* . Thus

$$v(\mathbf{CS}^*) \leq |\mathbf{A}| \max_S v(S) \leq |\mathbf{A}| \max_{\mathbf{CS} \in \mathbf{N}} v(\mathbf{CS}) = |\mathbf{A}| v(\mathbf{CS}_{\mathbf{N}}^*)$$

Number of **coalitions** at the second lowest level: $2^{\mathbf{A}} - 2$.

Number of **coalition structures** at the second lowest level:

$$\frac{1}{2}(2^{\mathbf{A}} - 2) = 2^{\mathbf{A}-1} - 1.$$

Thus the number of nodes visited is: $2^{\mathbf{A}-1}$. □

What exactly does the last theorem mean? Let n_{min} be the smallest size of \mathbf{N} such that a bound k can be established.

Positive result: $\frac{n_{min}}{\text{partitions of } \mathbf{A}}$ approaches 0 for $|\mathbf{A}| \rightarrow \infty$.

Negative result: To determine a bound k , one needs to search through exponentially many coalition structures.

Algorithm (CS-Search-1)

The algorithm comes in 3 steps:

- 1 Search the bottom two levels of the CS-graph.
- 2 Do a breadth-first search from the top of the graph.
- 3 Return the CS with the highest value.

This is an **anytime algorithm**.

Theorem 6.5 (CS-Search-1 up to Layer l)

With the algorithm **CS-Search-1** we get the following bound for k after searching through layer l :

$$\begin{cases} \lfloor \frac{|\mathbf{A}|}{h} \rfloor & \text{if } |\mathbf{A}| \equiv h - 1 \pmod{h} \text{ and } |\mathbf{A}| \equiv l \pmod{2}, \\ \lfloor \frac{|\mathbf{A}|}{h} \rfloor & \text{otherwise.} \end{cases}$$

where $h =_{\text{def}} \lfloor \frac{|\mathbf{A}| - l}{2} \rfloor + 2$.

Thus, for $l = |\mathbf{A}|$ (check the top node), k switches from $|\mathbf{A}|$ to $\frac{|\mathbf{A}|}{2}$.

Experiments

6-10 agents, values were assigned to each coalition using the following alternatives

- 1 values were uniformly distributed between 0 and 1;
- 2 values were uniformly distributed between 0 and $|A|$;
- 3 values were superadditive;
- 4 values were subadditive.




Figure 8: MERGE.




Figure 9: SPLIT.

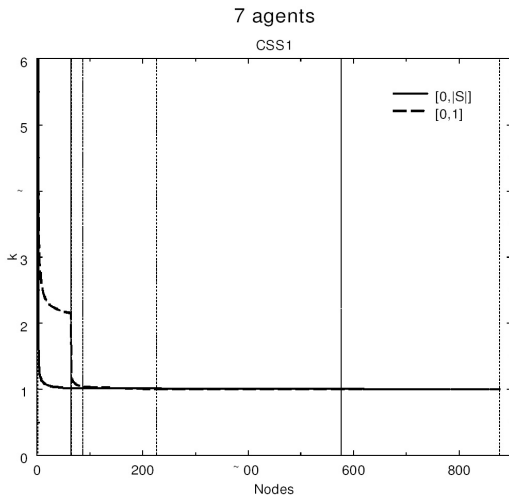


Figure 10: CS-Search-1.

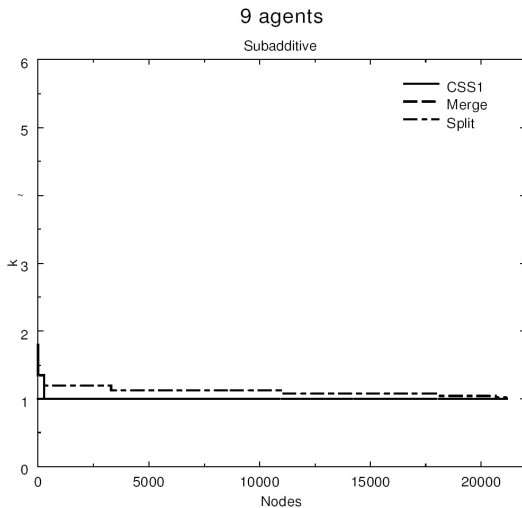


Figure 11: Subadditive Values.




Figure 12: Superadditive Values.




Figure 13: Coalition values chosen uniformly from $[0, 1]$.




Figure 14: Coalition values chosen uniformly from $[0, |S|]$.




Figure 15: Comparing CS-Search-1 with SPLIT.

- 1 Is **CS-Search-1** the **best anytime algorithm**?
- 2 The search for best k for $n' > n$ is perhaps not the same search to get best k for n .
- 3 **CS-Search-1** does not use any information while searching. Perhaps k can be made smaller by not only considering $v(\text{CS})$ but also $v(S)$ in the searched CS' .



6.2 General Contract Nets

How to distribute tasks?

- Global Market Mechanisms. Implementations use a **single centralised mediator**.
- **Announce, bid, award** -cycle. **Distributed Negotiation**.

We need the following:

- 1 **Define a task allocation problem in precise terms.**
- 2 **Define a formal model for making bidding and awarding decisions.**

Definition 6.6 (Task-Allocation Problem)

A **task allocation problem** is given by

- 1 a set of tasks T ,
- 2 a set of agents \mathbf{A} ,
- 3 a cost function $\mathbf{cost}_i : 2^T \rightarrow \mathbb{R} \cup \{\infty\}$ (stating the costs that agent i incurs by handling some tasks), and
- 4 the initial allocation of tasks

$$\langle T_1^{init}, \dots, T_{|\mathbf{A}|}^{init} \rangle,$$

where $T = \bigcup_{i \in \mathbf{A}} T_i^{init}$, $T_i^{init} \cap T_j^{init} = \emptyset$ for $i \neq j$.

Definition 6.7 (Accepting Contracts, Allocating Tasks)

A contractee q **accepts a contract** if it gets paid more than the marginal cost of handling the tasks of the contract

$$MC^{add}(T^{contract}|T_q) =_{def} \text{cost}_q(T^{contract} \cup T_q) - \text{cost}_q(T_q).$$

A contractor r is willing to **allocate the tasks** $T^{contract}$ from its current task set T_r to a contractee, if it has to pay less than it saves by handling them itself:

$$MC^{remove}(T^{contract}|T_r) =_{def} \text{cost}_r(T_r) - \text{cost}_r(T_r - T^{contract}).$$

Definition 6.8 (The Protocol)

Agents suggest contracts to others and make their decisions according to the above MC^{add} and MC^{remove} sets.

Agents can be both contractors and contractees. Tasks can be recontracted.

- The protocol is **domain independent**.
- Can only improve at each step: **Hill-climbing in the space of all task allocations**. Maximum is social welfare: $-\sum_{i \in A} \mathbf{cost}_i(T_i)$.
- **Anytime algorithm!**




Figure 16: TSP as Task Allocation Problem.

Definition 6.9 (O-, C-, S-, M- Contracts)

A contract is called of type

O (Original): only one task is moved: $\langle T_{i,j}, \rho_{i,j} \rangle, |T_{i,j}| = 1$.

C (Cluster): a set of tasks is moved: $\langle T_{i,j}, \rho_{i,j} \rangle, |T_{i,j}| \geq 1$.

S (Swap): if a pair of agents swaps a pair of tasks:

$$\langle T_{i,j}, T_{j,i}, \rho_{i,j}, \rho_{j,i} \rangle, |T_{i,j}| = |T_{j,i}| = 1.$$

M (Multi): if more than two agents are involved in an atomic exchange of tasks: $\langle \mathbf{T}, \rho \rangle$, both are $|A| \times |A|$ matrices. At least 3 elements are non-empty, $|T_{i,j}| \leq 1$.

Lemma 6.10 (O-Path reaches Global Optimum)

A path of O-contracts always exists from any task allocation to the optimal one. The length of the shortest such path is at most $|T|$.

Does that solve our problem? (Lookahead)

Definition 6.11 (Task Allocation Graph)

The **task allocation graph** has as vertices all possible task allocations (i.e. $|A|^{|T|}$) and directed edges from one vertex to another if there is a possible contract leading from one to the other.

For O-contracts, searching the graph using breadth-first search takes how much time?

Lemma 6.12 (Allocation graph is sparse)

We assume that there are at least 2 agents and 2 tasks. We consider the **task allocation graph for O contracts**. Then the fraction

$$\frac{\text{number of edges}}{\text{number of edges in the fully connected graph}}$$

converges to 0 both for $|T| \rightarrow \infty$ as well as $|A| \rightarrow \infty$

Lemma 6.13 (No IR-Path to Global Optimum)

*There are instances where **no path of IR O contracts exists** from the initial allocation to the optimal one. The length of the shortest IR path (if it exists) may be greater than $|T|$.*

But the shortest IR path is never greater than $|A|^{|T|} - (|A| - 1)|T|$.

Problem: local maxima.

A contract may be individually rational but the task allocation is not globally optimal.

Lemma 6.14 (No Path)

*There are instances where **no path of C-contracts** (IR or not) exists from the initial allocation to the optimal one.*

k -optimal

A task allocation is called k -optimal if no beneficial C-contract with clusters of k tasks can be made between any two agents.

Let $m \lesseqgtr n$. Does

- m optimality imply n optimality;
- n optimality imply m optimality?

Lemma 6.15 (No Path)

*There are instances where **no path of S-contracts** (IR or not) exists from the initial allocation to the optimal one.*

Lemma 6.16 (No Path)

*There are instances where **no path of M-contracts** (IR or not) exists from the initial allocation to the optimal one.*

Lemma 6.17 (Reachable allocations for S contracts)

We assume that there are at least 2 agents and 2 tasks. We consider the **task allocation graph for S contracts**. Given any vertex v the fraction

$$\frac{\text{number of vertices reachable from } v}{\text{number of all vertices}}$$

converges to 0 both for $|T| \rightarrow \infty$ as well as for $|A| \rightarrow \infty$.

Proof.

S contracts preserve the number of tasks of each agent. So any vertex has certain allocations $t_1, \dots, t_{|A|}$ for the agents. How many allocations determined by this sequence are there? There are exactly $|T|!$ many. This is to be divided by $|A|^{|T|}$. □

Theorem 6.18 (All Types necessary)

*For each of the 4 types there exist task allocations where **no IR contract with the remaining 3 types is possible**, but an IR contract with the fourth type is.*

Proof.

Consider O contracts (as fourth type). One task and 2 agents: $T_1 = \{t_1\}$, $T_2 = \emptyset$.

$c_1(\emptyset) = 0$, $c_1(\{t_1\}) = 2$, $c_2(\emptyset) = 0$, $c_2(\{t_1\}) = 1$. The O contract of moving t_1 would decrease global cost by 1. No C-, S-, or M-contract is possible.

To show the same for C or S contracts, two agents and two tasks suffice.

For M-contracts 3 agents and 3 tasks are needed. □

Theorem 6.19 (O-, C-, S-, M- $\not\Rightarrow$ Global Optima)

*There are instances of the task allocation problem where **no IR sequence from the initial task allocation to the optimal one exists** using O-, C-, S-, and M- contracts.*

Proof.

Construct cost functions such that the deal *where agent 1 gives one task to agent 2 and agent 2 gives 2 tasks to agent 1* is the only one increasing welfare. This deal is not possible with O-, C-, S-, or M-contracts. □

Corollary 6.20 (O-, C-, S-, M- $\not\Rightarrow$ Global Optima)

There are instances of the task allocation problem where no IR sequence from the initial task allocation to the optimal one exists using any pair or triple of O-, C-, S-, or M- contracts.

Definition 6.21 (OCSM Nets)

A **OCSM-contract** is a pair $\langle \mathbf{T}, \boldsymbol{\rho} \rangle$ of $|\mathbf{A}| \times |\mathbf{A}|$ matrices. An element $T_{i,j}$ stands for the set of tasks that agent i gives to agent j . $\rho_{i,j}$ is the amount that i pays to j .

- How many OCSM contracts are there?
- How much space is needed to represent one?

Theorem 6.22 (OCSM-Nets Suffice)

Let $|A|$ and $|T|$ be finite. If a protocol allows OCSM-contracts, any hill-climbing algorithm finds the globally optimal task allocation in a finite number of steps without backtracking.

Proof.

An OCSM contract can move from any task allocation to any other (in one step). So moving to the optimum is IR. Any hill-climbing algorithm strictly improves welfare. As there are only finitely many allocations, the theorem follows. \square

Theorem 6.23 (OCSM-Nets are Necessary)

If a protocol does not allow a certain OCSM contract, then there are instances of the task allocation problem where no IR-sequence exists from the initial allocation to the optimal one.

Proof.

If one OCSM contract is not allowed, then the task allocation graph contains two vertices without an edge. We let the initial and the optimal allocation be these two vertices. We construct it in such a way, that all adjacent vertices to the vertex with the initial allocation have lower social welfare. So there is no way out of the initial allocation. \square

[?] consider the multiagent version of the TSP problem and apply different sorts of contracts to it.

Several salesmen visit several cities on the unit square. Each city must be visited by exactly one salesman. They all have to return home and want to minimise their travel costs.

salesman = agent, task = city

- Experiments with up to 8 agents and 8 tasks. Initial allocation randomly chosen.
- **Ratio bound** (welfare of obtained local optimum divided by global optimum) and **mean ratio bound** (over 1000 TSP instances) were computed (all for fixed number of agents and tasks). Global optimum was computed using IDA*.
- A **protocol consisting of 5 intervals** is considered. In each interval, a particular contract type was considered (and all possible contracts with that type): 1024 different sequences.
- In each interval a particular order for the agents and the tasks is used. First all contracts involving agent 1. Then all involving agent 2 etc. Thus it makes sense to have subsequent intervals with the same contract type.

Order No.	Sequence	Social Welfare
1	OCOCO	1.03113
2	Oocco	1.03268
3	OCCOC	1.03276
4	OOCOC	1.03279
5	OCOOc	1.03413
6	SOCOC	1.03488
7	SOCCO	1.03536
8	COCOC	1.03755
9	COCCc	1.03857
10	MCOCO	1.03945
11	OCCCO	1.03954
12	MOCCO	1.03988
13	MOCOC	1.04001
14	MCCOC	1.04304
15	COCCO	1.04407

Table 1: The best Contract sequences.

Order No.	Sequence	Social Welfare
1	OCOCO	1.03113
2	Oocco	1.03268
3	OCCOC	1.03276
4	OOCOC	1.03279
375	<i>C-local</i>	<i>1.13557</i>
565	<i>O-local</i>	<i>1.2025</i>
579	OOOOO	1.21298
696	CCCCC	1.23515
1021	CSSSS	1.61181
1022	CMMMM	1.65965
1023	MMMMM	1.76634
1024	SSSSS	1.89321

Table 2: Best, average and worst Contracts.

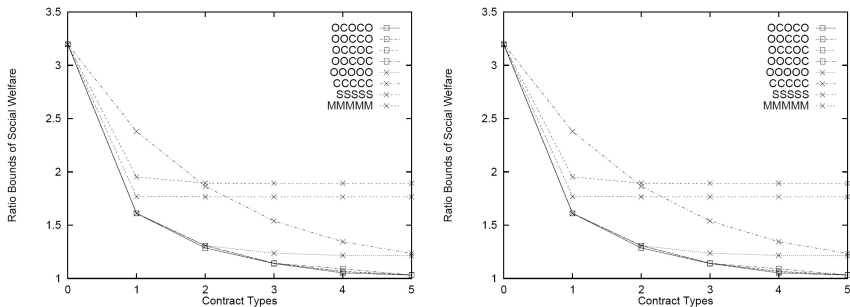




Figure 18: Ratio bounds for the 4 best and single type sequences.



- 1 Best protocol is 3.1% off the optimum.
- 2 12 best protocols are between 3% and 4% off the optimum.
- 3 Single type contracts do not behave well.
- 4 More contracts are tried and performed in the second interval (compared with the first).
- 5 The more mixture between contract types, the higher the social welfare.



6.3 Classes of Games

Idea: Consider a protocol (to build coalitions) as a game and consider Nash-equilibrium.

Problem: Nash-Eq is too weak!

Definition 6.24 (Strong Nash Equilibrium)

A profile is in **strong Nash-Eq** if there is no subgroup that can deviate by changing strategies jointly in a manner that increases the payoff of all its members, given that nonmembers stick to their original choice.

This is often too strong and does not exist.

Definition 6.25 (Monotone Games)

A CFG $\langle \mathbf{A}, \mathbf{v} \rangle$ is called **monotone**, if

$$\mathbf{v}(C) \leq \mathbf{v}(D),$$

for every pair of coalitions $C, D \subseteq \mathbf{A}$ such that $C \subseteq D$.

Many games have this property, but there may be **communication/coordination costs**. Or some players hate others and do not want to be in the same coalition. The next slide introduces a strictly stronger condition.

Definition 6.26 (Superadditive Games)

A CFG $\langle \mathbf{A}, \mathbf{v} \rangle$ is called **superadditive**, if

$$\mathbf{v}(S \cup T) \geq \mathbf{v}(S) + \mathbf{v}(T),$$

where $S, T \subseteq \mathbf{A}$ and $S \cap T = \emptyset$.

Lemma 6.27

Coalition formation for superadditive games is trivial.

Conjecture

All games are superadditive.

The conjecture is wrong, because the **coalition process** is not for free:
communication costs, penalties, time limits.

Definition 6.28 (Subadditive Games)

A CFG $\langle \mathbf{A}, \mathbf{v} \rangle$ is called **subadditive**, if

$$\mathbf{v}(S \cup T) \leq \mathbf{v}(S) + \mathbf{v}(T),$$

where $S, T \subseteq \mathbf{A}$ and $S \cap T = \emptyset$.

Coalition formation for subadditive games is trivial.

Superadditive Cover

Definition 6.29 (Superadditive Cover)

Given a game $G = \langle \mathbf{A}, \mathbf{v} \rangle$ that is not superadditive, we can transform it to a superadditive game $G^* = \langle \mathbf{A}, \mathbf{v}^* \rangle$ as follows

$$\mathbf{v}^*(C) := \max_{CS \in \mathcal{CS}_C} \mathbf{v}(CS)$$

This game is called the **superadditive cover** of G .

Convex Games

Definition 6.30 (Convex Game)

A CFG $\langle \mathbf{A}, v \rangle$ is **convex**, if for all coalitions T, S with $T \subseteq S$ and each player $i \in \mathbf{A} \setminus S$:

$$v(S \cup \{i\}) - v(S) \geq v(T \cup \{i\}) - v(T)$$

- **Convex games are superadditive.**
- **Superadditive games are monotone.**
- The other directions do not hold.

Example 6.31 (Treasure of Sierra Madre Game)

There are n people finding a treasure of many gold pieces in the Sierra Madre. Each piece can be carried by two people, not by a single person.

Example 6.32 (3-player majority Game)

There are three people that need to agree on something. If they all agree, there is a payoff of 1. If just 2 agree, they get a payoff of α ($0 \leq \alpha \leq 1$). The third player gets nothing.

Example 6.33 (Parliament)

Suppose there are four parties and the result of the elections is as follows:

- 1 A : 45 %,
- 2 B : 25 %,
- 3 C : 15 %,
- 4 D : 15 %.

We assume there is a 100 Mio Euro spending bill, to be controlled by (and distributed among) the parties that win.

Version 1: Simple majority wins ($\geq 50\%$).

Version 2: Any majority over 80% wins.

How do the $v(S)$ look like in the last three examples?

Definition 6.34 (Payoff Vector)

A **payoff vector** for a CFG and a coalition structure CS is a tuple $\langle x_1, \dots, x_n \rangle$ such that

- 1 $x_i \geq 0$ and $\sum_{i=1}^{|\mathbf{A}|} x_i = v(\mathbf{A})$,
- 2 $\forall C \in \text{CS} : \sum_{i \in C} x_i \leq v(C)$.

Note that the last condition is only supposed to hold for all **coalitions in the given coalition structure**.

If $\sum_{i \in C} x_i \geq v(C)$ would mean that a coalition gets more than its value, that is not possible.



6.4 The Core and its refinements

Consider a CFG that is not necessarily superadditive (so the grand coalition does not necessarily form). Assume that a certain coalition structure CS forms.

Definition 6.35 (Core of a CFG)

The **core of a CFG** is the **set of all pairs** $\langle \text{CS}, \langle x_1, \dots, x_n \rangle \rangle$ of coalition structures ($\text{CS} \in \mathcal{CS}_{\mathbf{A}}$) and payoff vectors such that the following holds:

$$\forall S \subseteq \mathbf{A} : \sum_{i \in S} x_i \geq v(S)$$

Here, the condition is supposed to hold **for all** S . We do not want any set of agents to form a new coalition. It ensures that only the grand coalition forms.

If $\sum_{i \in S} x_i \not\geq v(S)$, then these agents would form a coalition and get a higher payoff than in CS.

When the grand coalition forms, we can simplify the last definition.

Definition 6.36 (Core of Superadditive Games)

The **core of a superadditive CFG** is the set of all payoff vectors $\langle x_1, \dots, x_n \rangle$ such that the following holds:

$$\forall S \subseteq \mathbf{A} : \sum_{i \in S} x_i \geq v(S)$$

Thus the **core** corresponds to the **strong Nash equilibrium** mentioned in the beginning.

What about the core in the above examples?

Lemma 6.37

If $\langle \text{CS}, \langle x_1, \dots, x_n \rangle \rangle$ is in the core of a CFG $\langle \mathbf{A}, \mathbf{v} \rangle$, then $\mathbf{v}(\text{CS}) \geq \mathbf{v}(\text{CS}')$ for all coalition structures $\text{CS}' \in \mathcal{CS}_{\mathbf{A}}$.

Proof.

We can write $\mathbf{v}(\text{CS}) = \sum_{i \in \mathbf{A}} x_i = \sum_{C' \in \text{CS}'} x(C')$ and $\mathbf{v}(\text{CS}') = \sum_{C' \in \text{CS}'} \mathbf{v}(C')$.

Because of the definition of the core, $x(C') \geq \mathbf{v}(C')$ for all C' and therefore $\mathbf{v}(\text{CS}) \geq \mathbf{v}(\text{CS}')$. □

Theorem 6.38 (Core and superadditive cover)

Let a CFG $G = \langle \mathbf{A}, \mathbf{v} \rangle$ be given (not necessarily superadditive). Then G has a non-empty core if and only if its superadditive cover G^* has a non-empty core.

Proof \rightsquigarrow **exercise**

Theorem 6.39 (Convex games and their cores)

Each **convex** CFG $G = \langle \mathbf{A}, \mathbf{v} \rangle$ has a **non-empty core**.

Proof.

Let π be a permutation of \mathbf{A} and let $S_\pi(i)$ be the set of all **predecessors** of i wrt. π .

We claim that for $x_i := \mathbf{v}(S_\pi(i) \cup \{i\}) - \mathbf{v}(S_\pi(i))$, the core of G contains $\langle x_1, \dots, x_n \rangle$.

It is easy to show that all x_i are greater or equal to 0 and that they all sum up to the value of the game (\rightsquigarrow **exercise**). \square

(Proof of Theorem 6.39, cont.)

Assume there is a coalition $C = \{i_1, \dots, i_s\}$ such that $v(C) \not\geq x(C)$. Wlog we assume $\pi(i_1) \leq \dots \leq \pi(i_s)$. Obviously

$$v(C) = v(\{i_1\}) - v(\emptyset) + v(\{i_1, i_2\}) - v(\{i_1\}) + \dots + v(C) - v(C \setminus \{i_s\})$$

Because of convexity (apply convexity to $T_j := \{i_1, \dots, i_{j-1}\}$ and $S_j := \{1, 2, \dots, i_j - 1\}$) for all j :

$$v(T_j \cup \{i_j\}) - v(T_j) \leq v(S_j \cup \{i_j\}) - v(S_j) = x_{i_j}$$

Adding these pairs up, we get $v(C) \leq x(C)$, which is a contradiction.



6.5 Payoff Division: Shapley value and Banzhaf Index

We now assume w.l.o.g. that the grand coalition forms.
The payoff division should be fair between the agents, otherwise they would leave the coalition.

Definition 6.40 (Dummies, Interchangeable)

Agent i is called a **dummy**, if for all coalitions S with $i \notin S$:

$$v(S \cup \{i\}) - v(S) = v(\{i\}).$$

Agents i and j are called **interchangeable**, if for all coalitions S with $i \in S$ and $j \notin S$:

$$v(S \setminus \{i\} \cup \{j\}) = v(S)$$

Marginal Contribution

The marginality axiom, introduced by Young in the 80'ies, concentrates on the **marginal contributions** of a player in two different games.

Definition 6.41 (Marginal Contribution in two games)

We consider two CFG games over the same coalition structure, with values v and w . We say that agent i is **marginally indifferent between v and w** , if its marginal contributions in all coalitions is the same in both games: for all $S \subseteq A \setminus \{i\}$

$$v(S \cup \{i\}) - v(S) = w(S \cup \{i\}) - w(S).$$

Axioms for Payoff Division

Efficiency: $\sum_{i \in A} x_i = v(A)$.

Symmetry: If i and j are interchangeable, then $x_i = x_j$.

Dummies: For all dummies i : $x_i = v(\{i\})$.

Additivity: For any two games v, w :

$$x_i^{v \oplus w} = x_i^v + x_i^w,$$

where $v \oplus w$ denotes the game defined by
 $(v \oplus w)(S) = v(S) + w(S)$.

Marginality: If an agent i is **marginally indifferent** between two games v, w , then it should get the **same payoff** in both of them:

$$x_i^v = x_i^w.$$

Theorem 6.42 (Shapley-Value: 1st Characterisation)

For a CFG $G = \langle \mathbf{A}, \mathbf{v} \rangle$ there is **only one payoff division satisfying the first four axioms**. It is called the **Shapley value** of agent \mathbf{i} and is defined by

$$\phi_{\mathbf{i}}(G) = \frac{1}{|\mathbf{A}|!} \sum_{S \subseteq \mathbf{A} \setminus \{\mathbf{i}\}} (|\mathbf{A}| - |S| - 1)! |S|! (\mathbf{v}(S \cup \{\mathbf{i}\}) - \mathbf{v}(S))$$

Theorem 6.43 (Shapley-Value: 2nd Characterisation)

For a CFG $G = \langle \mathbf{A}, \mathbf{v} \rangle$ there is **only one payoff division satisfying the efficiency, symmetry and marginality**. It is the **Shapley value**.

The **expected gain** can be computed by taking a random joining order and computing the Shapley value.

- $(|\mathbf{A}| - |S|)!$ is the number of all possible joining orders of the agents (to form a coalition).
- There are $|S|!$ ways for S to be built before \mathbf{i} 's joining. There are $(|\mathbf{A}| - |S| - 1)!$ ways for the remaining agents to form S (after \mathbf{i}).
- $(v(S \cup \{\mathbf{i}\}) - v(S))$ is \mathbf{i} 's marginal contribution when added to set S .
- The Shapley value sums up the marginal contributions of agent \mathbf{i} **averaged over all joining orders**.

We have shown in Theorem 6.39 that convex games have non-empty cores. In fact, we can show a stronger statement.

Theorem 6.44 (Shapley and Core for convex games)

In each convex game, at least the Shapley value is contained in the core.

Definition 6.45 (Banzhaf Index)

For a CFG $G = \langle \mathbf{A}, v \rangle$ the **Banzhaf Index** of agent \mathbf{i} is

$$\beta_{\mathbf{i}}(G) = \frac{1}{2^{|\mathbf{A}|-1}} \sum_{S \subseteq \mathbf{A} \setminus \{\mathbf{i}\}} (v(S \cup \{\mathbf{i}\}) - v(S))$$

The Banzhaf Index satisfies all axioms but **efficiency**.

The following **normalised** Banzhaf index $\eta_{\mathbf{i}}(G)$ is also often considered:

$$\eta_{\mathbf{i}}(G) := \frac{\beta_{\mathbf{i}}(G)}{\sum_{\mathbf{i} \in \mathbf{A}} \beta_{\mathbf{i}}(G)} v(\mathbf{A}).$$

7. Social Choice and Auctions

- 7 Social Choice and Auctions
 - Classical Voting Systems
 - Formal model for social choice
 - Social Choice Functions
 - Social Choice Correspondences
 - Social Welfare Functions

Outline (1)

We deal with **voting systems** and discuss

- some **classical** approaches;
- an **abstract framework** to describe arbitrary voting mechanisms: **social choice theory**, and
- **Arrow's theorem** and some variants thereof in this framework.

We also consider **auctions** (deals between two agents). They constitute one of the most important frameworks for resource allocation problems between selfish agents. They can be seen as an important application of **mechanism design**, dealt with in Chapter ??.



7.1 Classical Voting Systems

Voting procedure

Agents give input to a **mechanism**: The outcome is taken as a **solution** for the agents.

Non-ranking voting: Each agent **votes for exactly one** candidate. Winners are those with a majority of votes.

Approval voting: Each agent can **cast a vote for as many candidates** as she wishes (at most one for each candidate). Winners are those with the highest number of approval votes.

Ranking voting: Each agent expresses his **full preference over the candidates**. Computing the winning can be complicated.

Candidates and Voters

From now on we assume there is a fixed set of alternatives (**candidates, outcomes**) O , in addition to the set of agents A , the elements of which we call now **voters**.

Definition 7.1 (Beat and tie)

We say that a candidate o **beats** another candidate o' (in direct comparison) if the number of voters that **strictly prefer** o to o' is **strictly greater** than the number of voters that **strictly prefer** o' to o .

If both numbers are equal, we say that o **ties** o' (in direct comparison).

Often e.g. in the french elections, **runoff** systems are considered: two candidates are singled out in the first round, one of which is then selected in the second round.

Definition 7.2 (Condorcet- winner, -Set)

- 1 A candidate o is a **Condorcet winner** if o **beats** any other candidate o' ($o' \neq o$).
- 2 A candidate o is a **weak Condorcet winner** if o **beats or ties** any other candidate o' ($o' \neq o$).
- 3 The **Condorcet set** is the set of weak Condorcet winners.

Note that sometimes the Condorcet winner is called **strict Condorcet winner**.

Condorcet helps

	1	2	3
1	A	B	C
2	B	C	A
3	C	B	A

Figure 20: A Tie, but Condorcet helps.

Condorcet does not help

	1	2	3
1	A	B	C
2	B	C	A
3	C	A	B

Figure 21: Strict Condorcet rules out all candidates.

Comparing A and B: majority for A.

Comparing A and C: majority for C.

Comparing B and C: majority for B.

Desired Preference ordering: $A > B > C > A$

Another interesting set is the following

Definition 7.3 (Smith set)

The **Smith set** is the **smallest**, non-empty set $S \subseteq O$ of candidates such that for each candidate $o \in S$ and each candidate $o' \notin S$ the following holds: o beats o' .

There is a strong relation between the **Condorcet set** and the **Smith set**. This will be treated in more detail in the exercise class.

Approval Voting

Definition 7.4 (Winner in Approval Voting)

A **winner in approval voting** is any candidate that received **at least as many votes as any other candidate**.

Can we model the situation in Figure 21 in approval voting?

Lemma 7.5 (Approval Voting and Condorcet)

*In approval voting, at least one of the winners is a **weak Condorcet winner**.*

Variants of Borda

Definition 7.6 (Borda Protocol: Standard and Nauru)

In **standard** Borda, each voter gives its best candidate $|O|$ points, the second best gets $|O| - 1$ points, etc.

In **Nauru** Borda, any first preference is assigned **1** point, any second just $1/2$, any third just $1/3$ and so on.

After all votes have been cast, they are **summed up, across all voters**. Winners are those with the **highest count**.

Originally, Jean-Charles de Borda wanted his system to determine a single winner. He also assumed all voters are **honest**.

Truncated Ballots

What, if voters do not want to **express full preference on all** candidates: **truncated ballots**.

Nauru: All candidates **must** be ranked.

Kiribati: Rank only a subset (but this subset completely), and all others get 0 points.

Modified BC: Points given depend on the **number of candidates ranked** (in each individual ballot).

Problem with Kiribati: Tactical voting. Bullet votes are more effective than fully ranked ballots.

Ties between candidates

Later we introduce **weak orders**: although they are **total**, they allow for ties (candidates among which the voter is indifferent).

Winner turns loser and loser turns winner.

Agent	Preferences
1	$A \succ B \succ C \succ D$
2	$B \succ C \succ D \succ A$
3	$C \succ D \succ A \succ B$
4	$A \succ B \succ C \succ D$
5	$B \succ C \succ D \succ A$
6	$C \succ D \succ A \succ B$
7	$A \succ B \succ C \succ D$
Borda count	C wins: 20, B: 19, A: 18, D loses: 13
Borda count without D	A wins: 15, B: 14, C loses: 13

Figure 22: Winner turns loser and vice versa.

Absolute Majority, but not elected

Example 7.7 (Nauru/Standard Borda, Plurality Voting)

	1	2	3	4
51 voters	A	C	B	D
5 voters	C	B	D	A
23 voters	B	C	D	A
21 voters	D	C	B	A

Who wins and who should win???

- 1 Plurality voting:
- 2 Standard Borda:
- 3 Nauru Borda:

Borda and Tactical Voting

Example 7.8 (Compromising and Burying)

We assume four cities M, N, K, and C want to become the capital of the state. M has most voters but is far away from the others. N, K, and C are close to the centre of the state.

	1	2	3	4
42 % (M)	M	N	C	K
26 % (N)	N	C	K	M
17 % (K)	K	C	N	M
15 % (C)	C	K	N	M

Borda would elect N . **How could voters of K change their poll to ensure C is chosen and not N ?**

Binary protocol: **Pairwise comparison.**

Take any two candidates and determine the winner. The winner enters the next round, where it is compared with one of the remaining candidates.

Which ordering should we use?

35% of agents have preferences $C \succ D \succ B \succ A$

33% of agents have preferences $A \succ C \succ D \succ B$

32% of agents have preferences $B \succ A \succ C \succ D$




Figure 23: Four different orderings and four alternatives.

Last ordering:

D wins, but all agents prefer C over D.

Coomb's method: Each voter ranks all candidates in **linear order**. If there is no candidate ranked first by a **majority of all voters**, the candidate which is ranked last (by a majority) is eliminated. The last remaining candidate wins.

d'Hondt's method: Each voter cast his votes. Seats are allocated according to the quotient $\frac{V}{s+1}$ (V the number of votes received, s the number of seats already allocated).

Nanson's method: Compute the Borda scores of all candidates and eliminate the candidate with the lowest Borda score (using some tie breaking mechanism). Then, proceed in the same way with the remaining candidates, recomputing the Borda score.

Proportional Approving voting: Each voter gives points $1, \frac{1}{2}, \frac{1}{3}, \dots, \frac{1}{n}$ for her candidates (she can choose as many or few as she likes). The winning candidates are those, where the **sum of all points is maximal** (across all voters).



7.2 Formal model for social choice

Ballots of the voters

How should a **general model for voting** look like? What are **fair** elections based on it?

Before determining such a model we need to answer the question

How should voters express their intentions?

Voters often have preferences over candidates: “A” is better than “B”, but “C” is better than both of them. Only if “C” is not a candidate, I would vote for “B”. What are properties of such an ordering?

First try: Here are two “obvious” properties.

transitive: If “A” is better than “B” and “B” is better than “C”, then “A” should be considered better than “C”.

no cycles: A voter should not be able to express that “A” is a strictly better candidate than “B” and, at the same time, “B” is strictly better than “A”.

Ballots of the voters (2)

Thus we could express the ballot of an agent as a **dag**, a directed acyclic graph.




Figure 24: Two Examples of dags.

In all voting systems considered in Subsection 4.1, the voters' ballot can be modelled with dags.

Dags are just **strict partial orders**: irreflexive and transitive binary relations. Most ballots are even based on **linear orders** (no indifference between any pair of candidates).

Cycles might make sense for non-strict orderings (a cycle might model that all elements in it are of equal standing).

Ballots of the voters (3)

⇒ We use **binary relations** for the ballot of an agent.

- **A** set of agents, O set of possible outcomes.
(O could be **A**, a set of laws, or a set of candidates).

Preferences based on binary orderings \prec_i

The **preference order** or **ranking** of agent **i** is described by a binary relation

$$\prec_i \subseteq O \times O.$$

- Which properties should we assume from such a binary relation?



Some Terminology on Posets

Here are some important properties of binary relations \prec :

Reflexivity: For all x : $x \prec x$.

Transitivity: For all x, y, z : if $x \prec y, y \prec z$ then $x \prec z$.

Antisymmetry: For all x : $x \prec y$ and $y \prec x$ implies $x = y$.

Totality: For all x, y : either $x = y$, or $x \prec y$ or $y \prec x$.

Irreflexivity: For all x : not $x \prec x$.

Asymmetry: If $x \prec y$ then not $y \prec x$.

A (non-strict) **partial order**, denoted by \preceq or \preccurlyeq or \leq or \leqslant or \leqq , is any relation satisfying reflexivity, transitivity and antisymmetry.

A **strict partial order**, denoted by \succneq or \succneqq or $\not\leq$ or $\not\leqq$, is any relation satisfying irreflexivity, and transitivity (and therefore also asymmetry).

Often, one simply writes \prec or $<$ to denote a strict order.

Ballots of the voters (4)

What are the **right** properties for \prec_i ?

- To express **dags**: irreflexive, transitive (both imply asymmetry) (**strict partial orders** \preceq).

We could also use non-strict partial orders \leq , which are in one-to-one correspondence via:

- \leq is the **reflexive closure** of \preceq ,
- \preceq is the **irreflexive kernel** of \leq .

How to express ties?

A partial order allows for **incomparability**: a and b might simply not be ordered at all (in the first graph on Slide 252 elements C and E are tied). **Voters have more freedom when they are allowed not to order candidates.**

Ballots of the voters (5)

- **Second try:** We could also consider **total (or linear) orders**: they are transitive and **strictly total** (for all $a \neq b$ either $a \prec_i b$ or $b \prec_i a$).

In that case, ties are not allowed: voters have to take a decision for each pair of candidates. So voters have less possibilities to express their ballot.

Ballots of the voters = Weak Orders

Orderings **inbetween** partial and total orders are the following:

Definition 7.9 (Weak Order, $L(O)$)

Any binary relation \succsim satisfying **transitivity** and **totality** (for all a, b : $a \succsim b$ or $b \succsim a$) is called a **weak order** (**total preorder**).

We denote by $L(O)$ the set of all such binary relations over O (we omit O if it is clear from context).

While weak orders rank all pairs of candidates, they **allow ties**: it is perfectly possible that there are pairs $a \neq b$ with $a \succsim b$ and $b \succsim a$.

Thus a and b are **indifferent**: the weak order treats them as equivalent. This can not happen with linear partial orders, because they are antisymmetric.

Using partial orders, such ties have to be modelled as **incomparable**. However, note the subtle differences between the two concepts.

Ballots of the voters = Weak Orders (cont.)

Any **weak order** induces a **strict weak order** \prec_i :

- $a \prec_i b$ iff $a \succsim_i b$ and not $b \succsim_i a$: “**i strictly prefers** b over a ”.

\prec_i is irreflexive and transitive, but not **total** anymore.

So we allow elements to be “equivalent” (or **indifferent**):

- $a \sim_i b$: $a \succsim_i b$ and $b \succsim_i a$. “**i is indifferent** between a and b ”.

However, **not any strict partial order can be obtained as a strict weak order.**

Attention

It is important not to confuse $=$ and \sim .

Dag, weak orders, strict weak order

Here we illustrate the use of dags compared with weak orders.




Figure 25: Equivalent modelings as dag (left), as weak order (middle), and as strict weak order.

Definition 7.10 (Preference Profile)

A **preference profile** for agents $1, \dots, |\mathbf{A}|$ is a tuple

$$\langle \succsim_1, \dots, \succsim_{|\mathbf{A}|} \rangle \in L^{|\mathbf{A}|} \quad (:= \prod_{i=1}^{|\mathbf{A}|} L)$$

We often write \succsim for $\langle \succsim_1, \dots, \succsim_{|\mathbf{A}|} \rangle$ if the set of agents is clear from context.

So the weak orders \succsim_i do allow for ties (although they are total). They are also called **preferences** or **rankings**.

- Often, not all subsets of O are *votable*, only a subset $V \subseteq 2^O \setminus \{\emptyset\}$. The simplest scenario is for $V = \{O\}$.
Each $v \in V$ represents a possible “set of candidates”. The voting model then has to select some of the elements of v .
- Each agent votes independently of the others. But we also allow that only a subset is considered. Let therefore be

$$U \subseteq \prod_{i=1}^{|\mathbf{A}|} L(V).$$

The set U represents the set of agents (and their preferences over the candidates) participating at the election and casting their votes.

Election systems

In the following sections we introduce three different election systems. The characterization of the voters preferences is different in all of them:

Social choice functions: Voters express their preferences by strict total orders \prec .

Social choice correspondences: Voters express their preferences by strict total orders \prec .

Social welfare functions: Voters express their preferences by weak orders \succsim .

Election Systems

We have now a good idea about how to model the voters and their preferences. **How about the outcome of an election?**

Social Choice Function (SCF): We define any

$$C^* : V \times U \rightarrow O; (v, \succsim) \mapsto o$$

as an election. The outcome is **one single winner**.

Social Choice Correspondence (SCC): We define any

$$W^* : V \times U \rightarrow 2^O; (v, \succsim) \mapsto v'$$

as an election. The outcome is a **set of winners**.

Social Welfare Function (SWF): This approach views a

$$f^* : U \rightarrow L; \succsim \mapsto \succsim^*$$

as an election: The outcome is a **weak order**, which determines the winners of the election (the maximal (or top) elements, for example.)

Dictators

In the next sections we show that under reasonable conditions on the election systems, there simply does not exist a fair election. More specifically we show results of the form

Dictators always exist

For each of the election systems defined, under reasonable assumptions on the election process, **there is only a dictatorship possible.**

What are the **most preferred** elements of \succsim ?

- $top(\succsim) = \{o \in O \mid \forall o' \in O : o \succsim o' \Rightarrow o' \succsim o\}$
- $bot(\succsim) = \{o \in O \mid \forall o' \in O : o \succsim o' \Rightarrow o \succsim o'\}$

For a strict total order \prec , $top(\prec)$ and $bot(\prec)$ are singletons.

Dictators (2)

Informally, a dictator is an agent i , such that whatever the profile of all voters \succsim looks like, the result of the election is always the one that agent i puts forward.

Let now $U \subseteq \prod_{i=1}^{|\mathbf{A}|} L(V)$ be given.

Social Choice Function: A dictator is an agent i if
for all \succsim : $top(\prec_i|_v) = \{C^*(v, \succsim)\}$

Social Choice Correspondence: A dictator is an agent i if
for all \succsim : $top(\prec_i|_v) \subseteq W^*(v, \succsim)$

Social Welfare Function: A dictator is an agent i if
for all \succsim : for all $o, o' \in O$, $o \prec_i o' \Rightarrow o \prec^* o'$
where \prec^* is the **strict** version of the weak order
 $f^*(\langle \succsim_1, \dots, \succsim_{|\mathbf{A}|} \rangle)$



7.3 Social Choice Functions

This section: ballots are **strict total orders** ↖.

Definition 7.11 (Social choice function (SCF))

A **social choice function** is any function

$$C^* : V \times U \rightarrow O; (v, \succ) \mapsto o$$

where $o \in v$.

A SCF returns exactly one “winner”. E.g. plurality voting where ties are broken in a predefined way (e.g. lexicographic ordering).

Definition 7.12 (Unanimity, Monotonicity)

A SCF C^* satisfies

surjectivity: for any $v \in V$ and outcome $o \in v$ there is a profile \succsim such that $C^*(v, \succsim) = o$.

unanimity: for any profile \succsim , $v \in V$ and outcome $o \in v$:
If for all $1 \leq i \leq |A|$: $o = \text{top}(\succsim_i|_v)$ then $C^*(v, \succsim) = o$,

Unanimity implies surjectivity.

strong monotonicity: for any \succsim , $v \in V$ and $o = C^*(v, \succsim) \in v$:
if \succsim' is a different profile such that
for all $o' \neq o$ and i : " $o' \prec_i o$ implies $o' \prec_i o$ ",
then $C^*(v, \succsim') = o$.

Strong monotonicity means: **any additional support for a winning alternative, should only benefit that alternative.**

Example 7.13 (Plurality with runoff)

This is the system used in the french elections. Assume that

6 voters support $B \succ C \succ A$

5 voters support $C \succ A \succ B$

6 voters support $A \succ B \succ C$

Under plurality with runoff, A and C make it to the second round, where A wins with 11 to 6.

Suppose 2 voters of the last group change their preferences and behave like the first group. **Thus there is additional support for A .**

But now, A and B make it to the second round, where B beats A with 9 to 8.

Theorem 7.14 (May (1952))

*If there are only two candidates, there is a SCF which is **not dictatorial but yet satisfies unanimity and strong monotonicity**.*

An example of such a SCF is **simple majority voting**, called **plurality voting** when there are more than two candidates.

In fact, we get a complete characterization if we assume two more, very natural properties expressing that a choice function should be symmetric wrt. (1) individuals (**anonymity**), and (2) alternatives (**neutrality**). As this result holds in the general case for correspondences, we refer to Slide 285.

In an exercise, you have to show that surjectivity together with strong monotonicity implies unanimity.

Theorem 7.15 (Muller-Satterthwaite (1977))

If there are at least 3 candidates, then any SCF satisfying surjectivity and strong monotonicity must be dictatorial.

What about **plurality voting**?

Suppose we fix a $o_{\text{fix}} \in O$, and define a SCF by mapping any ranking profile to o_{fix} . **Is that SCF strongly monotone? Is it non-dictatorial? Is it unanimous?** What sort of properties does it **not** satisfy?

Decisive sets

An important proof technique is that of **decisive** sets. Intuitively, a set of agents G is **decisive** on a pair (p, o) , if p cannot be a winner, if all agents in G put p below o . **Note: this is among all \succsim .**

Definition 7.16 (Decisive sets $G \subset A$)

A set $G \subset A$ is called **decisive** for the pair $(p, o) \in O^2$, if for all \succsim : “for all $i \in G: p \prec_i o$ ” implies $C^*(O, \succsim) \neq p$.

Definition 7.17 (Independence)

A SCF C^* satisfies the **Independence** property, if the following holds: If $C^*(O, \succsim) = o$, then $C^*(O, \succsim') \neq p$ for all $p \neq o$ and all \succsim' with “for all $i: p \prec_i o$ iff $p \prec'_i o$ ”.

Contraction-Lemma

Two observations are important:

- The grand coalition \mathbf{A} is decisive on any pair (p, o) (this follows from **surjectivity** and **strong monotonicity**).
 \rightsquigarrow Exercise.
- A singleton set $\{\mathbf{d}\}$ is decisive on any pair (p, o) if and only if \mathbf{d} is a dictator.

Contraction-Lemma

Let $G \subseteq \mathbf{A}$ with $|G| \geq 2$ be **decisive on all pairs** $(p, o) \in O^2$ and $G = G_1 \dot{\cup} G_2$ (meaning $G_1 \cap G_2 = \emptyset$ and $G = G_1 \cup G_2$). Then **G_1 or G_2 is decisive** on all pairs $(p, o) \in O^2$.

If we can prove the contraction lemma under some assumptions, then this implies that there is a dictator (using our two observations above).

Theorem of Muller-Satterthwaite (1977) (1)

Proof of Theorem 7.15.

- **Claim:** $G \subset \mathbf{A}$ is decisive on any pair $(p', o') \in O^2$, if the following holds: $G = \{i \in \mathbf{A} : p \prec_i o'\}$ s.t. their ranking ensures $C^*(O, \succsim) \neq p$. \rightsquigarrow whiteboard.
- We then show that strong monotonicity implies **independence**. Hint: Given a profile \succsim consider the following profile \succsim'' and apply strong monotonicity (twice): " $o \prec_i'' o'$ " iff: $o \prec_i o'$ and o, o' are in each \prec_i'' ranked among the two top places (o, o' are the two outcomes in the independence property).
- Then we show the Contraction-Lemma (in the proof we need the independence property).



Theorem of Muller-Satterthwaite (1977) (2)

Proof (cont.)

We prove the Contraction lemma. So $G = G_1 \dot{\cup} G_2$ is given. We construct a ranking profile \succsim as follows.

- for all $i \in G_1$: $q \succ_i p \succ_i o$,
- for all $i \in G_2$: $o \succ_i q \succ_i p$,
- for all $i \in \mathbf{A} \setminus G$: $p \succ_i o \succ_i q$,
- all other $o' \in O$ are ranked below o, p, q by all agents i .

Because G is **decisive on all pairs**, q can not be the winner. So either o or p is the winner. □

Theorem of Muller-Satterthwaite (1977) (3)

Proof (cont.)

p is the winner: Only those in G_2 rank o below p .
Therefore G_2 is decisive (using our **Claim** on Slide 274).

o is the winner: Only those in G_1 rank q below o .
Independence implies that q loses against o in each profile in which exactly the agents from G_1 rank q below o .
Therefore G_1 is decisive.



Manipulation

Can voters **hide their true preferences** and use another ballot to achieve better results?

Example 7.18 (Plurality Voting)

49 %:	Nader	⤵	Gore	⤵	Bush
20 %:	Bush	⤵	Nader	⤵	Gore
20 %:	Nader	⤵	Bush	⤵	Gore
11 %:	Bush	⤵	Gore	⤵	Nader

If the last group of voters change their preferences by putting their favorite at the end, then they achieve a better result overall! (Compare with Example 7.8 on Slide 245.)

Can we find choice functions where such manipulations are not possible?

Strategy-Proofness

Earlier, in Definition 3.18 on Slide 49 we have introduced the notation using $-i$ to denote a profile without agent i 's entry. We are using this notation here as well.

Definition 7.19 (Strategy-Proofness)

A social choice function C^* is called **strategy-proof**, if there is no agent i and profile $\langle \vec{x}_{-i}, x'_i \rangle$ such that

$$C^*(O, \vec{x}) \succ_i C^*(O, \langle \vec{x}_{-i}, x'_i \rangle)$$

We compare the profile $\langle x_1, \dots, x_{i-1}, x'_i, \dots, x_A \rangle$, in which agent i **misrepresents** her true preference x_i , with the real profile $\langle x_1, \dots, x_{i-1}, x_i, \dots, x_A \rangle$.

Theorem of Gibbard and Satterthwaite

Lemma 7.20

Strategy-proofness implies strong monotonicity.

Proof.

We assume that a SCF is not strongly monotone and show that it is not strategy-proof. So we assume there is $o \neq o'$ and \succsim, \succsim' s.t.

- $C^*(O, \succsim) = o$, and $C^*(O, \succsim') = o'$, and
- for all $\mathbf{i} \in \mathbf{A}$, for all $p \in O \setminus \{o\}$: $p \prec_{\mathbf{i}} o$ implies $p \prec'_{\mathbf{i}} o$.

We are now modifying \succsim as follows. For $\mathbf{i} = 1, 2, \dots, \mathbf{n}$ we replace successively $\prec_{\mathbf{i}}$ by $\prec'_{\mathbf{i}}$, until the winner of the new action profile under C^* is no more o but somebody else (which must happen, because at the end it is o'). Let \mathbf{j} be this agent.

Because we can adapt our assumption to this new situation, we assume wlog that \succsim, \succsim' differ only at the entry \mathbf{j} . So wlog we can use the notation o' as in our assumption. □

Theorem of Gibbard and Satterthwaite (2)

Proof (cont.)

Case 1: “ $o' \succ_j o$ ”. If j 's true preferences are as in \succ_j , then it pays off for j to vote as in \succ_j (to ensure o is winning, and not o' .) So it is not strategy-proof.

Case 2: “not $o' \succ_j o$ ”. Then “not $o' \succ_j o$ ”, therefore $o \succ_j o'$. If j 's true preferences are as in \succ_j , then it pays off for j to vote as in \succ_j (to ensure o' is winning, and not o). So it is not strategy-proof.



Theorem of Gibbard and Satterthwaite (3)

Theorem 7.21 (Gibbard-Satterthwaite (1973/1975))

*If there are at least 3 candidates, then any SCF satisfying **surjectivity** and **strategy-proofness** must be **dictatorial**.*

Using Lemma 7.20 this is just a corollary to Theorem 7.15.

Note that we could replace **surjectivity** by the stronger **unanimity**.



7.4 Social Choice Correspondences

This section: ballots are strict total orders \prec . We also assume $V = O$ in order to simplify notation.

Definition 7.22 (Social choice correspondence)

A **social choice correspondence** is any function

$$W^* : U \rightarrow 2^O; \succ \mapsto v$$

where $v \neq \emptyset$.

A correspondence returns a **nonempty set** of winners.

The **Borda rule** is a typical example of a **correspondence**.

Impossibility results similar to Theorem 7.15 are rare. The most important one is due to **Duggan and Schwartz**: See Theorem 7.28 on Slide 291 .

Definition 7.23 (Positive responsiveness)

A SCC W^* satisfies

positive responsiveness: for any profile \succsim , $o \in O$:

if $o \in W^*(\succsim)$ then $W^*(\succsim')$ = $\{o\}$, provided that \succsim' is a different profile such that for all $o' \neq o \neq o''$ and i the following holds:

$$o' \prec_i o \text{ implies } o' \prec'_i o, \text{ and } o' \prec_i o'' \text{ iff } o' \prec'_i o''.$$

Intuitively, this property means that if o is among the winners and at least one voter raises o up, then o **should become the sole winner**. This is very intuitive in case there are only two candidates!

Anonymity and neutrality

Definition 7.24 (Anonymity and Neutrality)

\mathbf{W}^* satisfies **anonymity** if for all permutations π on $\{1, \dots, |\mathbf{A}|\}$:

$$\mathbf{W}^*(\langle \prec_1, \dots, \prec_{|\mathbf{A}|} \rangle) = \mathbf{W}^*(\langle \prec_{\pi(1)}, \dots, \prec_{\pi(|\mathbf{A}|)} \rangle)$$

\mathbf{W}^* satisfies **neutrality** if for all permutations π on O

$$\pi(\mathbf{W}^*(\vec{z})) = \mathbf{W}^*(\pi(\vec{z})),$$

where $\pi(\vec{z})$ is defined componentwise and $\pi(\prec_i)$ is defined in the obvious way: $o \prec_i o'$ iff: $\pi(o) \prec_i \pi(o')$.

Which properties satisfies the SCC that always declares

- all candidates as winners;
- the **two** top choices of plurality voting as winners?

Only two candidates: $|O| = 2$

Theorem 7.25 (May (1952))

Assume there are only two candidates, $|O| = 2$. A SCC W^* satisfies **anonymity**, **neutrality**, and **positive responsiveness** if and only if W^* is **simple majority voting**.

Assume there are only 2 candidates and 2 voters, and we are considering a social choice function C^* (not a correspondence). \rightsquigarrow **Then the choice function can not satisfy both anonymity and neutrality.**

Weak Monotonicity

There is a weaker version of the responsiveness condition, namely **weak monotonicity**, when we replace

“ $\mathbf{W}^*(\succsim')$ ” by $o \in \mathbf{W}^*(\succsim')$

- Does Theorem 7.25 hold for weak monotonicity instead of positive responsiveness?

Optimistic and pessimistic voters

Strategy-proof in Definitions 7.27 and ?? rules out **manipulability** by **untruthful** voting: by **misrepresenting their true preferences**, voters should not be able to get overall **better** results.

Social choice correspondences determine **sets of winners** (not just single winners).

⇒ We need to **rank** such sets.

Optimistic and pessimistic voters (2)

Definition 7.26 (Optimist and pessimists)

An agent i is an **optimist**, if she ranks X higher than Y , whenever $\text{top}(\succsim_i|_Y) \succsim_i \text{top}(\succsim_i|_X)$.

An agent i is a **pessimist**, if she ranks X higher than Y , whenever $\text{bot}(\succsim_i|_Y) \succsim_i \text{bot}(\succsim_i|_X)$.

By slightly abusing notation, we also write $Y \succsim_i X$.

As we are considering strict linear orders, top and bottom elements are always unique.

But even in the case of **weak orders** all maximal elements (and all minimal elements) are indifferent to each other, so both definitions are well-defined (see Slide 258).

Strategy-Proofness for correspondences

Definition 7.27 (Strategy-Proofness)

A social choice correspondence \mathbf{W}^* is called **strategy-proof**, if there is no agent i , profile \vec{v} and ranking \prec'_i such that for all $v \in V$

$$\mathbf{W}^*(\langle \vec{v}_{-i}, \prec_i \rangle) \prec_i \mathbf{W}^*(\langle \vec{v}_{-i}, \prec'_i \rangle)$$

We compare the profile $\langle \prec_1, \dots, \prec_{i-1}, \prec'_i, \dots, \prec_A \rangle$, in which agent i **misrepresents** her true preference \prec_i , with the real profile $\langle \prec_1, \dots, \prec_{i-1}, \prec_i, \dots, \prec_A \rangle$.

Note that we consider the relation

$\mathbf{W}^*(\langle \vec{v}_{-i}, \prec_i \rangle) \prec_i \mathbf{W}^*(\langle \vec{v}_{-i}, \prec'_i \rangle)$ only for optimistic or pessimistic voters i (see Definition 7.26).

Theorem 7.28 (Duggan/Schwartz (2000))

We assume that there are at least 3 candidates.
Then any SCC that is **nonimposed** (for each $o \in O$ there is \vec{z} s.t. $W^*(\vec{z}) = \{o\}$) and **strategy-proof** for both optimistic and pessimistic voters is **dictatorial**.

What is the importance of the **nonimposed** property?

The proof is based on several lemmas. It has striking similarity to the proof of Gibbard/Satterthwaite. Important concepts are **down monotonicity** and **dictating sets** (corresponding to decisive sets).

Definition 7.29 (Down-monotonicity)

Suppose we have a profile \vec{r} where $W^*(\vec{r})$ is a singleton and the following holds: If we modify \vec{r} by letting one agent move a losing alternative down one spot (obtaining profile \vec{r}'), then $W^*(\vec{r}) = W^*(\vec{r}')$.

Then we call W^* **down-monotone for singleton winners**.

Lemma 7.30

A SCC W^* satisfying **strategy-proofness for both optimistic and pessimistic voters** also satisfies **down-monotonicity for singleton winners**.

Definition 7.31 (Dictating Sets G, pG_o)

Let $G \subseteq A$ a group of agents and $p, o \in O$. We denote by $pG^{w*}o$ the fact that $W^*(\succ) \neq \{p\}$ for all profiles \succ in which all agents from G rank o above p .

A set $G \subseteq A$ is called **dictating for W^*** , if $pG^{w*}o$ holds for all pairs (p, o) .

Lemma 7.32

Suppose $G \subseteq A$, $p, o \in O$, and $pG^{w*}o$. Let $o \neq o' \neq p$ and $G = G_1 \dot{\cup} G_2$.

Then we have $o'G_1^{w*}o$ or $pG_2^{w*}o'$.

Lemma 7.33

If W^* is down-monotonic for singleton winners and nonimposed, then the **set of all agents is a dictating set**.

Lemma 7.34

We assume SCC W^* satisfies **strategy-proofness for both optimistic and pessimistic voters**.

- Let G be a dictating set. If it is the disjoint union of sets G_1 and G_2 , then one of these sets is dictating too.
- There is an agent whose maximal element is the unique winner, whenever $W^*(\succsim)$ is a singleton.



7.5 Social Welfare Functions

In this section agents ballots are **weak orders** \succsim . We also use $o \prec_i o'$ to denote “ $o \succsim o'$ and not $o' \succsim o$ ”: o' is **strictly greater** than o .

Definition 7.35 (Social welfare function)

A **social welfare function** is any function

$$f^* : U \rightarrow L; (\succsim_1, \dots, \succsim_{|A|}) \mapsto \succsim^*$$

For each $V \subseteq 2^O \setminus \{\emptyset\}$ the function f^* w.r.t. U induces a choice correspondence $C_{\langle \succsim_1, \dots, \succsim_{|A|} \rangle}$ as follows:

$$C_{\langle \succsim_1, \dots, \succsim_{|A|} \rangle} =_{\text{def}} \begin{cases} V & \longrightarrow & V \\ v & \mapsto & \text{top}(\succsim^*|_v) \end{cases}$$

Each tuple \succsim determines the election for all possible $v \in V$.

What are desirable properties for f^* ?

- **Weak Pareto Efficiency:**

for all $o, o' \in O$: $(\forall i \in \mathbf{A} : o \prec_i o')$ implies $o \prec^* o'$.

- **Independence of Irrelevant Alternatives (IIA):**

for all $o, o' \in O$:

- $(\forall i \in \mathbf{A} : o \prec_i o' \text{ iff } o \prec'_i o') \Rightarrow (o \prec^* o' \text{ iff } o \prec'^* o')$,
- $(\forall i \in \mathbf{A} : o \succsim_i o' \text{ iff } o \succsim'_i o') \Rightarrow (o \succsim^* o' \text{ iff } o \succsim'^* o')$
- $(\forall i \in \mathbf{A} : o \sim_i o' \text{ iff } o \sim'_i o') \Rightarrow (o \sim^* o' \text{ iff } o \sim'^* o')$

IIA expresses that the social ranking of two alternatives does only depend on the **relative** individual rankings of these alternatives.

Note that this implies in particular

$$(\forall i \in \mathbf{A} : \succsim_i|_v = \succsim'_i|_v) \Rightarrow$$

$$\forall o, o' \in v, \forall v' \in V \text{ s.t. } v \subseteq v' : (o \prec^*|_{v'} o' \text{ iff } o \prec'^*|_{v'} o')$$

- Shouldn't we also require in the pareto efficiency condition that

$$(\forall i \in \mathbf{A} : o \succsim_i o') \text{ implies } o \succ^* o'?$$

The answer is *no*, the stated condition, **weak pareto efficiency**, is perfectly sufficient for proving Arrows theorem. The stronger condition, called **pareto efficiency**, is not needed. Note that for strict linear orders, there is no such distinction.

Example 7.36 (Which champagne is the best?)

Suppose you go out for dinner and you want to start with a champagne.

- The waiter gives you the choice between a Blanc de Blancs or a Blanc de Noirs (both grands crus from respectable houses)
- You choose a Blanc de Blancs.
- Then the waiter returns and mentions that they also have a Rosé.
- *"Oh, in that case, I take a Blanc de Noirs."*

In this example, an irrelevant alternative (the **Rosé**) **does matter**.

Plurality Vote

The simple **plurality vote** protocol does not satisfy the **IIA**.

We consider 7 voters ($\mathbf{A} = \{1, 2, \dots, 7\}$) and $O = \{A, B, C, D\}$,
 $V = \{\{A, B, C, D\}, \{A, B, C\}\}$. The columns in the following table represent two different preference orderings of the voters (black and red).

	\prec_1 (\prec_1)	\prec_2 (\prec_2)	\prec_3 (\prec_3)	\prec_4 (\prec_4)	\prec_5 (\prec_5)	\prec_6 (\prec_6)	\prec_7 (\prec_7)
A	1 (2)	1 (2)	1 (1)	1 (1)	2 (2)	2 (2)	2 (2)
B	2 (3)	2 (3)	2 (2)	2 (2)	1 (1)	1 (1)	1 (1)
C	3 (4)	3 (4)	3 (3)	3 (3)	3 (3)	3 (3)	3 (3)
D	4 (1)	4 (1)	4 (4)	4 (4)	4 (4)	4 (4)	4 (4)

Let \prec^* be the solution generated by \prec and \prec^* the solution generated by the \prec .
 Then we have for $i = 1, \dots, 7$: $B \prec_i A$ iff $B \prec_i A$, but $B \prec^* A$ and $A \prec^* B$. The latter holds because on the whole set O , for $\prec^* A$ gets selected 4 times and B only 3 times, while for $\prec^* A$ gets selected only 2 times but B gets still selected 3 times. The former holds because we even have $\prec_i|_{\{A,B,C\}} = \prec_i|_{\{A,B,C\}}$.

The introduction of the **irrelevant** (concerning the relative ordering of A and B) **alternative** D changes everything: **the original majority of A is split and drops below one of the less preferred alternatives (B).**

Theorem 7.37 (Arrow (1951))

We assume that there are at least two voters and three candidates ($|O| \geq 3$). If the SWF f^ satisfies **Weak Pareto Efficiency** and **IIA**, then there always exists a dictator.*

Proof (of Arrows theorem).

The proof is the third proof given by John Geanakoplos (1996) and based on the following

Lemma 7.38 (Strict Neutrality)

We assume **Pareto Efficiency** and **IIA** and consider two pairs of alternatives a, b and α, β . Suppose each voter **strictly prefers** a to b **or** b to a , i.e. for all i : $a \succ_i b$ or $b \succ_i a$. Suppose further that each voter has the same preference for α, β as she has for a, b .

Then either $a \succ^* b$ and $\alpha \succ^* \beta$ or $b \succ^* a$ and $\beta \succ^* \alpha$.

A simple corollary is the following:

Corollary 7.39 (Extremal Lemma)

Let the social welfare function f^* satisfy **Pareto Efficiency** and **IIA**. Let $o \in O$ and suppose each voter i puts o either on the very top (unique top element wrt. \succsim_i) or to the very bottom (unique bottom element wrt. \succsim_i).

Then o is **either a unique bottom or a unique top element** of \succsim^* .

Proof (of the lemma).

We assume wlog that (a, b) is distinct from (α, β) and that $b \succ^* a$ (we have to show the preference is strict).

We construct a different profile $\langle \succ'_1, \dots, \succ'_{|A|} \rangle$ obtained as follows (for each i):

- If $a \neq \alpha$, we change \succ_i by moving α just strictly above a .
- If $b \neq \beta$, we change \succ_i by moving β just strictly below b .

This can be done **by maintaining the old preferences between α and β** (as preferences between a and b are strict).

By pareto efficiency, we have $a \succ_i^* \alpha$ (for $\alpha \neq a$)
and $\beta \succ_i^* b$ (for $\beta \neq b$).

By IIA, we have $b \succ_i^* a$. Using transitivity, we get
 $\beta \succ_i^* \alpha$.

By IIA again, we also get $\beta \succ_i^* \alpha$ (because $\alpha \succ_i \beta$ iff
 $\alpha \succ_i \beta$).

We now reverse the roles of (a, b) with (α, β) and
apply IIA again to get $b \succ_i^* a$. □

The proof of Arrows theorem is by considering two alternatives a, b and the profile where $a \prec_i b$ for all agents i . By pareto efficiency, $a \prec^* b$. **Note that this reasoning is true for all rankings with the same relative preference between a and b .**

We now consider a sequence of profiles $\langle \prec_1^l, \dots, \prec_{|A|}^l \rangle$ (from $l = 0, \dots, |A|$, starting with the one described above ($l = 0$)), where in step l , we let all agents numbered $\leq l$ change their profile by **moving a strictly above b** (leaving all other rankings untouched). In fact, only one agent, number l changes its ranking in step l .

We consider the rankings \prec^{l^*} obtained from $\vec{\prec}^l$.

There must be one step, let's call it d , where $a \prec^{d-1^*} b$ but $b \prec^{d^*} a$ (because of pareto efficiency and strict neutrality).

Again: this reasoning is true not just for one profile $\langle \prec_1^l, \dots, \prec_{|A|}^l \rangle$, **but for all such profiles** with the same relative ranking of a and b .

We claim that d is a dictator.

Take any pair of alternatives α, β and assume wlog $\beta \succ_d^d \alpha$
(otherwise the following argument works as well for $\alpha \succ_d^d \beta$).

d is a dictator, when we can show $\beta \succ_{d^*} \alpha$.

Take $c \notin \{\alpha, \beta\}$ (because $|O| \geq 3$) and consider the new profile $\langle \succ_{\mathbf{i}}^d, \dots, \succ_{|\mathbf{A}|}^d \rangle$ obtained as follows from $\langle \succ_{\mathbf{i}}^d, \dots, \succ_{|\mathbf{A}|}^d \rangle$:

- for $1 \leq \mathbf{i} \neq d$: we put c on top of each $\succ_{\mathbf{i}}^d$.
- for d : we put c inbetween α and β .
- for $d \neq \mathbf{i} \leq |\mathbf{A}|$: we put c to the bottom of each $\succ_{\mathbf{i}}^d$.

We are changing the profile $\langle \succ_{\mathbf{i}}^d, \dots, \succ_{|\mathbf{A}|}^d \rangle$ by moving c around in a very particular way.

We apply **strict neutrality** to the pair $\langle c, \beta \rangle$ and $\langle \alpha, b \rangle$. Because both pairs have the same relative ranking in $\langle \succ_{\mathbf{i}}^d, \dots, \succ_{|\mathbf{A}|}^d \rangle$, we have $\beta \succ_{d^*} c$.

Now we consider the new profile $\langle \succsim''_1, \dots, \succsim''_{|A|} \rangle$ obtained as follows from $\langle \succsim^d_1, \dots, \succsim^d_{|A|} \rangle$:

- for $1 \leq i \neq d$: we put c to the bottom of each \succsim^d_i .
- for d : we put c in between α and β .
- for $d \neq i \leq |A|$: we put c to the top of each \succsim^d_i .

We now apply **strict neutrality** to the pair (α, c) and (a, b) . Because both pairs have the same relative ranking in $\langle \succsim''_1, \dots, \succsim''_{|A|} \rangle$, we have $c \succ^{d^*} \alpha$.

By transitivity: $\beta \succ^{d^*} \alpha$. □

Ways out (of Arrow's theorem):

- 1 Choice function is not always defined.
- 2 Independence of alternatives is dropped.



7.6 References



Geanakoplos, J. (1996, April).

Three brief proofs of arrow's impossibility theorem.

Cowles Foundation Discussion Papers 1123R3, Cowles Foundation, Yale University.

available at

<http://ideas.repec.org/p/cwl/cwldpp/1123r3.html>.



Sandholm, T. W. (1997).

Chapter 5: Distributed Rational Decision Making.

In G. Weiss (Ed.), *Multiagent Systems*. Cambridge, MA: MIT Press.