

Multi-Agent Systems course - Summary

Marc van Zee

February 2012

Contents

Preface	3
1 Introduction	4
2 Noncooperative Game Theory: Games in Normal Form	7
2.1 Introduction	7
2.2 Games in Normal Form	7
2.2.1 Examples	8
2.2.2 Definition of games in normal form	8
2.2.3 More examples of normal-form games	9
2.2.4 Strategies in normal-form games	10
2.2.5 Expected utility	11
2.3 Analyzing games: from optimality to equilibrium	12
2.3.1 Pareto efficiency	12
2.3.2 Best response and Nash equilibrium	13
2.3.3 Finding Nash equilibria	14
2.4 Further solution concepts for normal-form games	15
2.4.1 Maxmin strategies	15
2.4.2 Minmax strategies	15
2.4.3 Removal of dominated strategies	16
2.5 Preferences and Utility	17
3 The Extensive Form: Games with Sequential Actions	20
3.1 Perfect information extensive-form games	20
3.1.1 From extensive-form to normal-form	22
3.1.2 Subgame-perfect equilibrium	22
3.1.3 Computing equilibria: backward induction	24
3.2 Imperfect-information extensive-form game	26
4 Communication	29
4.1 "Doing by talking" I: Cheap talk	29
4.2 "Talking by doing": signaling games	31
4.3 "Doing by Talking" II: Speech-act theory	33
4.3.1 Speech acts	33
4.3.2 A game-theoretical view of speech acts	33
4.3.3 Agent communication languages	35

5	Aggregating Preferences: Social Choice	36
5.1	Introduction	36
5.2	A formal model	37
5.3	Voting	38
5.3.1	Voting methods	38
5.3.2	Voting paradoxes	38
5.4	Existence of social function	39
5.5	Restrictions on preferences	41
5.5.1	Single-peaked preferences	41
5.5.2	Dichotomous preferences	42
6	Protocols for Strategic Agents: Mechanism Design	43
6.1	Social choice functions and strategic game forms	43
6.2	Incentive compatibility of social choice functions	45
6.3	Implementation in dominant strategies	45
7	Protocols for Multiagent Resource Allocation: Auctions	47
7.1	Single-good auctions	47
7.1.1	Second-price auctions	48
7.1.2	First-price auctions	49
7.1.3	Revenue equivalence	49
7.2	Multiunit auctions	50
7.2.1	Single-unit demand	51
7.3	Combinatorial auctions	52
7.3.1	Simple combinatorial auction mechanism	53
7.3.2	Expressing a bid: bidding languages	54

Preface

This is a summary of the material covered in the course Multi-agent Systems, given in 2011/2012 by dr. Dastani at the University of Utrecht. The official website of the course is <http://www.cs.uu.nl/docs/vakken/mas/>. This document is supposed to cover all the slides of the course and (parts of) Chapter 3, 5, 8, 9, 10 and 11 of the book *Multiagent Systems* by Shoham and Leyton-Brown. Since large parts are copied from this book and the slides, this summary in no way attempts to provide new information. The book *Multi-agent systems* can be bought or downloaded from the website <http://www.masfoundations.org/>.

The correspondence between the book and this summary is as follows: Chapter 2 in this summary will discuss games in normal form, which corresponds to chapter 3 of the book. Chapter 3 discusses games in extensive form, which is chapter 5 of the book. In chapter 4 we will consider communication; chapter 8 of the book. Subsequently, chapter 5 (social choice theory) and chapter 6 (mechanism design) correspond to chapter 9 and 10 of the book. Chapter 10 is relatively the most deviating chapter from the book. Here we will mainly discuss how to actually implement a social choice function, which is largely left implicitly in the book. Finally, chapter 7 considers auctions (chapter 11 in the book).

This summary does not contain all material that is covered in the book by Shoham and Leyton-Brown. By taking the lectures of dr. Dastani as a guide, I tried to make an estimate of what would be important for the course and what not, but it is in no way meant to be complete. Many technical details are missing, and nearly all proofs are omitted.

For questions etc. feel free to contact me.

- Marc van Zee (marcvanee@gmail.com)

Chapter 1

Introduction

The motivation for studying multiagent systems often comes from an interest in artificial (software or hardware) agents. The umbrella of multiagent systems is very broad; it ranges from artificial intelligence to economics, philosophy and linguistics. These topics are on its own worth several books. The goal of the book *Multi-agent systems* is to gather the most important elements from these subjects and in this way give a balanced and accurate introduction to this broad field. In this summary, we will not consider all of these subjects. In general, it is important to keep the following keywords in mind: *coordination*, *competition* and *communication*.

Working definition of a Multi-Agent System (MAS): A MAS consists of a set of autonomous entities, called agents, which interact with each other and their surrounding environment to achieve their (joint) objectives. This can be looked upon from three perspectives:

1. *Computing perspective:* From this perspective MAS are an advance in computing science because of 1) the increase in computational power, 2) the interconnection between the different entities, 3) the increase in intelligence because of the increase in complexity of the system, 4) delegation of control between the agents and 5) human-oriented interaction with high-level concepts and metaphors such as goals and beliefs.
2. *Software engineering perspective:* Multi-agent programming can be seen as the next step after object oriented programming:
 - Structural analysis methodology: divide and conquer (e.g. C, Basic)
 - Object-oriented methodology: design patterns (e.g. Java, C++)
 - Agent-oriented methodology: organization and society patterns (2APL/2OPL, Jason, JADE)
3. *Artificial intelligence perspective:* while Artificial Intelligence is used for the development of MAS, they are not the same. Essential to MAS are social intelligence and emergent behavior, while AI considers broader topics such as planning, learning an vision.

Characteristics of MAS

- Consists of interacting autonomous agents.
- Designed to achieve a *global* goal.
- Agents are able to *cooperate*, *coordinate*, and *negotiate*.
- Specified using high-level abstract concepts (roles, permissions, interactions, ...)

Examples of MAS include a transportation system, auctions (which will come back in Chapter 11) or a marketplace.

As said before a MAS consists of autonomous agents. We now will define what these agents are, and then specify what the three above-mentioned abilities - cooperate, coordinate and negotiate - mean. **Autonomous agents** are reactive, pro-active, social (interaction and communication) and rational (behave to maximize achievement). An agent is an abstract notion which can be any autonomous entity, ranging from a thermostat to a robot or even a human. Capabilities of agents:




Figure 1.1: *Schematic overview of an agent*

- *Communication*: agents communicate using messages. Besides the message, there is also a need for some agreement between the agents to what communication protocol will be used. See Chapter 8 for more information on communication.
- *Cooperation*: agents can cooperate by sharing/distributing tasks or by distributing results of their actions.
- *Negotiation*: agents negotiate to reach agreement, examples include auctions (Chapter 11) or argumentation.

Besides these capabilities, agents also need to be embedded in some system that make them able to perform their tasks. This system is called an *organization*.

Organization: aims to arrange and manage the relation between agents to ensure qualities and outcomes of the overall system. Examples include normative systems, electronics institutions or market places.

A crucial aspects of a MAS is coordination; agents activities need to be synchronized in some way. This can be established endogenous (from within the agents) or exogenous

(facilitated by the organization).

What are the challenges for the designer of a MAS?

- Analyze, specify, design and build individual agents that can act autonomously in order to successfully carry out a task...
- ...Especially when the agents cannot be assumed to share the same interests/goals.

While the concept of a MAS is very powerful and provokes the imagination, building one in practice is very challenging. When multiple agents participate in interaction, it is computationally infeasible to predict all the outcomes. Therefore, as usually in science, to understand the complexity of this phenomenon we turn to a simplified situation and try to understand this first. This is where game theory comes in...

Chapter 2

Noncooperative Game Theory: Games in Normal Form

Note: while the corresponding chapter in the book actually starts with a section on preferences, I have decided to start by explaining games in Normal Form first, because I believe it will give a better intuition on what follows.

2.1 Introduction

Game theory: the mathematical study of interaction between independent, self-interested agents.

Noncooperative Game Theory: focuses on the modeling of an individual, *self-interested* agents (including his beliefs, preferences, and possible actions).

Cooperative Game Theory: focuses on *groups* of agents. This will not be discussed in this summary.

Noncooperative game theory involves *self-interested agents*, which means that every agent has his own, individual description. This is usually modeled using *utility theory*. This is explained in more detail in Section 3.5. For now we should just realize that agents have a utility function whose, when playing a game, expected value they want to maximize. It reflects the *preferences* of an agent for the different possible actions they can perform.

2.2 Games in Normal Form

On first sight, game theory might look simple; agents simply choose the actions that maximize their expected utility. This is not the case when considering two or more agents with actions that can affect each other's utility. In this situation, choosing an action might be more beneficial to another agent than it is to the agent that performs the action. To study such situations in a formal and structured setting, we turn to game theory.

2.2.1 Examples

Let us consider a classic of game theory, the prisoner's dilemma. Its story runs as follows:

Two suspects are taken into custody and separated. The district attorney is certain that they are guilty of a specific crime, but he does not have adequate evidence to convict them at a trial. He points out to each prisoner that each has two alternatives: to confess to the crime the police are sure they have done, or not to confess. If they will both do not confess, then the district attorney states he will book them on some very minor trumped up charge such as petty larceny and illegal possession of a weapon, and they will both receive minor punishment; if they both confess they will be prosecuted, but he will recommend less than the most severe sentence; but if one confesses and the other does not, then the confessor will receive lenient treatment for turning state's evidence whereas the latter will get "the book" slapped on him. (Luce and Raiffa, 1957, p. 95)

		Player 2	
		Not Confess	Confess
Player 1	Not confess	2,2	0,3
	Confess	3,0	1,1

Figure 2.1: *The Prisoner's Dilemma*

A natural way to represent such a dilemma, or game, in game theory is via a two-dimensional matrix (or grid). In general, each row corresponds to a possible action for player 1, each column to a possible action for player 2, and each cell corresponds to one possible outcome. The utility that each player receives for an outcome is written in the cell, with player 1's utility listed first. The Prisoners' Dilemma is shown in Figure 2.1. As can be seen, it is best for both agents not to confess, with the danger of the other agent to confess and receiving the highest utility possible. While the player do not gain a lot when they both confess, this intuitively seems to be safe strategy. Given these options, should Player 1 choose to confess or to defect? Does it depend on what it colleague does? Will the game change if the players are allowed to communicate? Game theory answers many of these questions.

2.2.2 Definition of games in normal form

The normal, or strategic, form is the most familiar representation of a game in game theory. The game in Figure 2.1, for example, is in normal form. A game written in this way represents each player's utility for every combination of actions for all the players. In this chapter, we will only consider **two player** games, so we can simply represent every game in a two-dimensional matrix, but one could imagine that a three player game would require a three-dimensional matrix. Notice that when representing a game in this way, we look at a so-called *stage game*, which means that we only look at one round; we

do not consider the notion of time. This means that we assume that all players make a move simultaneously. Alternative, an *extensive form game* (Chapter 5) involves an element of time. We will now formally introduce the normal, or strategy, game form.

Definition 1 (Strategic game form). *A strategic game form is a quadruple (N, A, O, g) , where:*

- N is a finite set of n players;
- $A = A_1 \times \dots \times A_n$, where A_i is a finite set of strategies or actions available to player i . Each vector $a = (a_1, \dots, a_n) \in A$ is called a strategy profile;
- O is a set of outcomes;
- $g : A \mapsto O$ is an outcome function.

The game form does not consider the preferences or *utility functions* for the players. If we add these, we end up with a *strategic game*.

Definition 2 (Strategic game). *A strategic game is a quintuple (N, A, O, g, u) , where:*

- (N, A, O, g) is a game form;
- $u = (u_1, \dots, u_n)$, where $u_i : A \mapsto \mathbb{R}$ is a utility function for player i .

Sometimes a normal-form game is denoted as (N, A, u) instead of (N, A, O, g, u) by assuming that $A = O$, which makes the outcome function g unnecessary.

Notice that A is the cartesian product of all the sets of actions that each player has available. Consider, for example, the Prisoner's Dilemma in Figure 2.1. This game contains two sets of actions (one for each player), namely $A_1 = \{Not\ Confess, Confess\}$ and $A_2 = \{Not\ Confess, Confess\}$. Taking the cartesian product of these sets gives us the action profiles:

$$A_1 \times A_2 = \{(Not\ Confess, Not\ Confess), (Not\ Confess, Confess), (Confess, Not\ Confess), (Confess, Confess)\}.$$

As may be clear now, the action profiles, or strategy profile, are depicted in the cells of the matrix.

In the next section we will look at several more examples of normal-form games and introduce different classes of games.

2.2.3 More examples of normal-form games

Games in which, for every action profile, all players have the same payoff are called **common-payoff games**.

Definition 3 (Common-payoff game). *A common-payoff game is a game in which for all action profiles $a \in A_1 \times \dots \times A_n$ and any pair of agents i, j , it is the case that $u_i(a) = u_j(a)$.*

Such games are also known as (*pure*) *coordination games*. Players have no conflicting interests, which means they in fact have to cooperate to obtain an optimal result. Figure 2.2 shows such a coordination game.

In **zero-sum games**, players receive exactly the opposite payoff of their opponent. This can be more generally described as a constant-sum game where the sum of the payoffs for the two players in each of the action profiles is always the same.

	Left	Right
Left	1,1	0,0
Right	0,0	1,1

Figure 2.2: *Coordination game*

Definition 4 (Constant-sum game). *A two-player normal-form game is constant sum if there exists a constant c such that for each strategy profile $a \in A_1 \times A_2$ it is the case that $u_1(a) + u_2(a) = c$.*

Notice that we will only consider *zero-sum* games, while they are actually a subset of the constant-sum games. Figure 2.3 contains a classic example of such a game, called *Matching Pennies*. This game represents two players that toss a coin. If the sides of the pennies match (both Heads or both Tails), player 1 wins. Otherwise, player 2 wins. Another famous game is *Rock, Paper, Scissors*, depicted in Figure 2.4. The intuition behind these games is that the win of the one player equals the loss of the other.

	Heads	Tails
Heads	1,-1	-1,1
Tails	-1,1	1,-1

Figure 2.3: *Matching Pennies game*.

	Rock	Paper	Scissors
Rock	0,0	-1,1	1,-1
Paper	1,-1	0,0	-1,1
Scissors	-1,1	1,-1	0,0

Figure 2.4: *Rock, Paper, Scissors game*.

2.2.4 Strategies in normal-form games

We need a formal way of describing what action(s) players will take in the future, so called *strategy profiles*. A *strategy profile* is a set of strategies for each player which fully specifies the actions that he will choose in a game. It must include one and only one strategy for every player. We distinguish between two kinds of strategy profiles:

- *Pure strategy profile*: each player selects a single action. For example: in Figure 2.3, (Heads, Heads) is one of the four pure strategy profiles.
- *Mixed strategy profile*: each player plays according to some probability distribution over its available actions. In our example: playing Heads 50% of the time is a mixed strategy. We can formally define the mixed strategy of one player used the notion of a *mixed strategy*.

Definition 5 (Mixed strategy). *Let (N, A, u) be a normal-form game, and for any set X let $\Pi(X)$ be the set of all probability distributions over X . Then the set of mixed strategies for player i is $S_i = \Pi(A_i)$.*

So, since A_i represents the set of all actions for player i , $\Pi(A_i)$ is actually a function that takes this set and returns all possible probability distributions for these actions. When we want to denote a single mixed strategy for a player, we actually choose one of these distributions and denote this with s_i . Therefore $s_i \in S_i$.

Definition 6 (Mixed-strategy profile). *The set of mixed-strategy profiles is simply the Cartesian product of the individual mixed-strategy sets, $S_1 \times \dots \times S_n$.*

$s_i(a_i)$ represents the probability that an actions a_i will be played under mixed strategy s_i . The subset of actions that are assigned a positive probability by the mixed strategy s_i is called the *support* of s_i .

Definition 7 (Support). *The support of a mixed strategy s_i for a player i is the set of pure strategies $\{a_i | s_i(a_i) > 0\}$.*

Note that a pure strategy is actually also a mixed strategy, in which the support is a single action with a probability of 1. A *fully mixed strategy* assigns to every action a nonzero probability.

2.2.5 Expected utility

The *expected utility* of a player is the payoff that the players expect to receive when playing according to some strategy.

Definition 8 (Expected utility of a mixed strategy). *Given a normal-form game (N, A, u) , the expected utility u_i for player i of the mixed-strategy profile $s = (s_1, \dots, s_n)$ is defined as*

$$u_i(s) = \sum_{a \in A} u_i(a) \prod_{j=1}^n s_j(a_j).$$

Let consider an example to understand all the previously describe definitions. Recall the game in Figure 2.2 (the Coordination game). A possible *mixed-strategy* for player 1 is $s_1 = (0.4, 0.6)$, meaning that he will play *Left* 40% of the time, and *Right* for 60%. Notice that to play according to this mixed strategies, the game has to be played for several rounds. Likewise assume that player 2 plays according to $s_2 = (0.25, 0.75)$. Notice also that the probabilities for one player always have to sum to 1 (see Figure 2.5).

	Left (0.25)	Right (0.75)
Left (0.4)	1,1	0,0
Right (0.6)	0,0	1,1

Figure 2.5: *Coordination game*

Now, lets apply definition 8 by looking at the *expected utility* for player 1. We can simply multiply the probability that an action profile will be selected with the payoff the player receives for it. So,

$$\begin{aligned}
 u_1(s) &= Pr(a_1 \equiv Left) \times Pr(a_2 \equiv Left) \times 1 + Pr(a_1 \equiv Left) \times Pr(a_2 \equiv Right) \times 0 + \\
 &\quad Pr(a_1 \equiv Right) \times Pr(a_2 \equiv Left) \times 0 + Pr(a_1 \equiv Right) \times Pr(a_2 \equiv Right) \times 1 \\
 &= 0.4 \times 0.25 \times 1 + 0.4 \times 0.75 \times 0 + 0.6 \times 0.25 \times 0 + 0.6 \times 0.75 \times 1 \\
 &= 0.1 + 0 + 0.15 + 0 \\
 &= 0.25
 \end{aligned}$$

2.3 Analyzing games: from optimality to equilibrium

An *optimal strategy* is a strategy that maximizes the agent's expected payoff. Below we will discuss two notion of optimality in strategies: pareto efficiency and dominance.

2.3.1 Pareto efficiency

Definition 9 (Pareto efficiency). *An outcome $o \in O$ is (weakly) Pareto efficient if there is no outcome that is strictly better for all players, i.e. if there is no $o' \in O$ such that $\forall i \in N : o' >_i o$.*

Since we usually consider games in which the outcomes of the game equal the strategy, or action, profiles, we will further refine this notion in the context of strategy profiles.

Definition 10 (Pareto efficient strategy profile). *A mixed strategy profile $s \in \Delta(A)$ is (weakly) Pareto efficient if there is no mixed strategy profile that is strictly better for all players, i.e., if there is no $s' \in \Delta(A)$ such that $\forall i \in N : u_i(s') > u_i(s)$.*

What is/are the Pareto efficient strategy/ies in the games in Figure 2.6 and 2.7? Answer in this footnote¹.

	L	R
T	2,2	2,4
D	3,4	1,1

Figure 2.6: Pareto game #1.

	L	R
T	2,2	0,3
D	4,-1	1,1

Figure 2.7: Pareto game #2.

Definition 11 (Dominance). *A pure strategy a_i for player i (strongly) dominates another strategy a'_i if for any strategies of the opponents, a_i leads to more preferable outcome than a'_i , i.e. if $\forall b \in A : (b_1, \dots, a_i, \dots, b_n) >_i (b_1, \dots, a'_i, \dots, b_n)$.*

For example, in the game in Figure 2.7 the dominant pure strategy for player 1 is D . This is because playing D will always yield a higher payoff than playing T , irrelevant of

¹Game #1: Only (D,L) because in all other strategy profiles both players can improve by playing (D,L); Game #2: all except (D,R).

what the opponent does. If a strategy profile is the product of dominant strategies, we say that it is a *dominant strategy profile*. The strategy profile (*Confess, Confess*) in the Prisoner's Dilemma of Figure 2.1 is thus a strongly dominant strategy profile.

Definition 12 (Dominance for mixed strategies). *A mixed strategy s_i for player i (strongly) dominates another strategy s'_i if for any mixed strategies of the opponents, s_i has a greater expected utility than s'_i , i.e., if $\forall t_{j \neq i} \in \Delta(A_j) : u_i(t_1, \dots, s_i, \dots, t_n) > u_i(t_1, \dots, s'_i, \dots, t_n)$.*

A mixed strategy s_i of player i that (strongly) dominates all other mixed strategies of i is called a *strongly dominant strategy* for player i . A mixed strategy profile (s_1, \dots, s_n) is called a (strongly) *dominant mixed strategy equilibrium* if s_i is a (strongly) dominant strategy for player i for every $i = 1, \dots, n$.

2.3.2 Best response and Nash equilibrium

Best reponse: the best reaction of a player to a current strategy profile, resulting in the highest payoff possible, given all the strategies of the other players. For example, in the game of Figure 2.7, the best response to L for player 1 is D , because $u_1((D, L)) > u_1((T, L))$.

We use s_{-i} to denote a strategy profile without i 's strategy. So in Figure 2.7, s_{-1} can be L or R .

Definition 13 (Best response). *Player i 's best response to the strategy profile s_{-i} is a mixed strategy $s_i^* \in S_i$ such that $u_i(s_i^*, s_{-i}) \geq u_i(s_i, s_{-i})$ for all strategies $s_i \in S_i$.*

Normally an agent does not know the strategies of the other players, so best responding is not a direct solution. Still, it is used in the most important notion in game theory: the infamous Nash equilibrium.

Definition 14 (Nash equilibrium). *A strategy profile $s = (s_1, \dots, s_n)$ is a Nash equilibrium if, for all agents i , s_i is a best response to s_{-i} .*

In this strategy profile, no player has a reason to deviate, because it is a best response strategy for all players. There are two kinds of Nash equilibria, namely pure and mixed:

- *Pure Nash equilibrium:* can be read from the game matrix directly, by checking for each strategy profile whether a player will not gain when deviating. If this is the case of all the players, we have found one.
- *Mixed Nash equilibrium:* this is the result of players playing a mixed strategy, and cannot be read from the matrix directly. Therefore, it should be computed. We turn to this in the next section.

Theorem 1 (Nash, 1950). *Every strategic game with a finite number of pure strategies has a Nash equilibrium in mixed strategies.*

2.3.3 Finding Nash equilibria

Consider the game in Figure 2.8 and 2.9, can you find the pure Nash equilibria? Answers in footnote².

	L	R
T	2,1	0,0
D	0,0	1,2

Figure 2.8: *Nash equilibrium game #1.*

	A	B	C
A	3,1	1,1	0,0
B	1,1	1,2	5,0
C	0,1	4,0	0,0

Figure 2.9: *Nash equilibrium game #2.*

So how do we find a mixed Nash equilibrium? The trick is to choose the mixed strategy of player 1 in such a way that player 2 will become indifferent between its actions; meaning that no matter what player 2 chooses, they will all result in the same payoff. We can do exactly the same for player 2; we choose his strategy in such a way that player 1 will get the same payoff, no matter what he plays. We have now found a mixed strategy profile in which both players make sure that the other player will always get the same payoff and thus does not have a reason to deviate; every strategy is a best response!

Let try to find the mixed Nash equilibrium for the case of the game in Figure 2.8, which are not pure. Lets assume that player 1 plays T with probability p and D with probability $1 - p$. Now, we want to choose p is such a way that the utility for player 2 for playing L equals that of R :

$$\begin{aligned}
 u_2(L) &= u_2(R) \\
 1 * p + 0 * (1 - p) &= 0 * p + 2 * (1 - p) \\
 p &= \frac{2}{3}
 \end{aligned}$$

We do the same for player 2, denoting the probability that player 2 plays L with q :

$$\begin{aligned}
 u_1(T) &= u_1(D) \\
 2 * q + 0 * (1 - q) &= 0 * q + 1 * (1 - q) \\
 q &= \frac{1}{3}
 \end{aligned}$$

This means that the mixed strategy profile $((\frac{2}{3}, \frac{1}{3}), (\frac{1}{3}, \frac{2}{3}))$ is a mixed-strategy equilibrium. We will now try to find the Nash equilibrium of the Matching Pennies game

²game #1: (T,L) and (D,R); game #2: (A,A). An effective technique to find them is to start with any strategy profile and see if it is improvable for a player. For example, in game #2, start with (B,B), player 1 can improve by playing C which leads to (C,B), there player 2 improves by playing A resulting in (C,A), then player 1 plays A which leads to the Nash equilibrium (A,A) where no player can improve by deviating

(Figure 2.3).

$$\begin{aligned}u_2(\text{Left}) &= u_2(\text{Right}) \\1 * p + 0 * (1 - p) &= 0 * p + 1 * (1 - p) \\p &= \frac{1}{2}\end{aligned}$$

Similarly we can find the mixed strategy of player 2, resulting in the mixed-strategy equilibrium $((\frac{1}{2}, \frac{1}{2}), (\frac{1}{2}, \frac{1}{2}))$.

2.4 Further solution concepts for normal-form games

We reason about multiplayer games using *solution concepts*, principles according to which we identify interesting subsets of the outcomes of a game. While the most important solution concept is the Nash equilibrium, there are also a large number of others, and we will discuss some of them here.

2.4.1 Maxmin strategies

Player i 's *maxmin strategy* is a strategy that maximizes i 's worst-case payoff, in the situation where all the other players $-i$ happen to play the strategies which cause the greatest harm to i . The *maxmin value* (or *security level*) of the game for player i is the minimum amount of payoff guaranteed by a maxmin strategy. So why would an agent want to play a maxmin strategy?

- It could be a conservative agent maximizing worst-case payoff...
- ...Or a paranoid agent who believes everyone is out to get him

It is anyway a very safe strategy (hence the term security level). By convention, the maxmin value for player 1 is called the *value of the game*.

Definition 15 (Maxmin strategy). *The maxmin strategy for player i is $\arg \max_{s_i} \min_{s_{-i}} u_i(s_1, s_2)$.*

Definition 16 (Maxmin value). *The maxmin value for player i is $\max_{s_i} \min_{s_{-i}} u_i(s_1, s_2)$.*

Can you find the pure maxmin strategies and values of both players in the game of Figure 2.9?³.

2.4.2 Minmax strategies

Player i 's minmax strategy in a 2-player game is a strategy that minimizes the other player's best-case payoff. The minmax value of the 2-player game for player i is that maximum amount of payoff that $-i$ could achieve under i 's minmax strategy. Why would an agent want to play a minmax strategy?

- If it wants to punish the other agent as much as possible.

Definition 17 (Minmax strategy). *The minmax strategy for player i is $\arg \min_s \max_{s_{-i}} u_{-i}(s_1, s_2)$.*

³for Player 1: first we determine the worst-case payoff of every action, which results in A: 0, B: 1, C: 0. Then we maximize this, which leads to strategy B and value 1. For player 2 the worst-case payoffs are A: 1, B:0, C:0 so the maxmin strategy is A and the value is 1 as well.

Definition 18 (Minimax value). *The minmax value for player i is $\min_s \max_{s_{-i}} u_i(s_1, s_2)$.*

Can you find the pure minmax strategies and values of both players in the game of Figure 2.9?⁴

An important theorem regarding these solution concepts is the minimax theorem.

Theorem 2 (Minimax theorem (von Neumann, 1928)). *In any finite, two-player, zero-sum game, in any Nash equilibrium each player receives a payoff that is equal to both his maxmin value and his minmax value.*

2.4.3 Removal of dominated strategies

Games with dominant strategies play an important role in game theory, but they are rare in naturally occurring games. More common are dominated strategies.

Definition 19 (Dominated strategy). *A strategy s_i is (strictly) dominated for an agent i if some other strategy s'_i (strictly) dominates s_i .*

The procedure of iterated elimination of dominated strategies is as follows:

- Eliminate one after another actions of a player that are (weakly or strongly) dominated, until this is no longer possible.
- If only one payoff profile remains, we say the game is *dominance solvable*.

	D	E	F
A	3,1	0,1	0,0
B	1,1	1,1	5,0
C	0,1	4,1	0,0

	D	E
A	3,1	0,1
B	1,1	1,1
C	0,1	4,1

	D	E
A	3,1	0,1
C	0,1	4,1

Figure 2.10: *Elimination (1).*

(2)


(3)

The strategy profiles that survive iterated elimination of weakly dominated actions may depend on the order of elimination. This is not the case for iterated elimination of strongly dominated actions. We will now illustrate this process for the two games in Figure 2.10 and 2.11.

In Figure 2.10 (1), the strategy F is dominated by all other strategies, so we can eliminate it. In (2), there seem to be not dominated strategies anymore. This is indeed the case for the pure strategies, but when we also consider mixed strategies we see that B for the row player is dominated by the mixed strategy that selects either A or C with equal probability. And so we are left with the maximally reduced game in (3).

In Figure 2.11, we first remove strategy C for the row player because it is dominated by A . Then we remove F for the column player because it is dominated by E , resulting

⁴for Player 1: first we determine the maximum payoff of player 2 for every action, which results in A: 1, B: 2, C: 1. Then we minimize this, which leads to strategy A or C and value 1. We now also see that the minmax value is not necessarily unique. For player 1 the maximum payoffs to consider are A: 3, B:4, C:5 so the minmax strategy is A and the value is 3.



in (2). Finally we remove both D for the column player and A for the row player, resulting in the single strategy profile (B, E) . This means that the game is *dominance solvable*.

Note that this is a much weaker solution concept than Nash equilibrium—the set of strategy profiles will include all the Nash equilibria, but it will include many other mixed strategies as well.

Let us conclude by stating several characteristics of Nash equilibria.

- Nash equilibrium is perhaps the most important solution concept for non-cooperative games, for which numerous refinements have been proposed.
- Any combination of dominant strategies is a Nash equilibrium.
- Nash equilibria are not generally Pareto efficient.
- Existence in (pure) strategies is not in general guaranteed.
- Nash equilibria are not in general unique (equilibria selection, focal points).
- Nash equilibria are not generally interchangeable.
- Payoffs in different Nash equilibria may vary.
- Every game possesses at least one (mixed) Nash Equilibrium.

2.5 Preferences and Utility

Because the idea of utility seems so straightforward and logical, one might wonder why anyone would argue with this formal model for reasoning about an agent's happiness in different situations. But when we look at it closer, one might wonder why it is actually a one-dimensional preference. Is this expressive enough for all the different preferences an agent may have over the different alternatives? Perhaps we need more dimensions, say one for each alternative.

To show that a quantitative notion of utility (meaning some numerical value, as shown for example in Figure 2.1) is actually sufficient, we start by abstracting away from it by considering *preferences* of an agent. First we assume that agents possess a qualitative preference over their actions, meaning that all actions are comparable in some way. This is reasonably fair to assume, since humans are often able to express such preferences (for example: "*I rather go fishing than to stay at home, but I prefer to go to hunting the most*"). Then we define several characteristics for these preferences, which allow us to prove that for every preference relation, there exists some utility function that is equally expressive. These utility functions then return a quantitative value, which is easier to use when reasoning about preference.

Let O be a finite set of outcomes. For any $o_1, o_2 \in O$, we define the following preference orderings for the agent:

- $\mathbf{o}_1 \sim \mathbf{o}_2$: o_1 and o_2 have equal preference;
- $\mathbf{o}_1 \succeq \mathbf{o}_2$: o_1 is weakly preferred over o_2 ;
- $\mathbf{o}_1 \succ \mathbf{o}_2$: o_1 is strictly preferred over o_2 (so $o_1 \succeq o_2$ and $\neg(o_1 \sim o_2)$)

Now we will define six axioms that will together imply the theorem we are looking for. The first axiom simply states that all outcomes are comparable with each other; the \succeq relation induces an ordering over them.

Axiom 1 (Completeness). $\forall o_1, o_2 : o_1 \succ o_2$ or $o_2 \succ o_1$

The second axiom means that outcomes are transitive. The absence of this axiom would lead to irrational behavior of our agent. For example, an agent would then be allowed to prefer staying in bed to going to work and prefer going to work to a funeral, while at the same time he would prefer a funeral to staying in bed. Now who would want that?

Axiom 2 (Transitivity). *If $o_1 \succeq o_2$ and $o_2 \succeq o_3$, then $o_1 \succeq o_3$*

For the third axiom we need the notion of a *lottery*, which assigns a probability to each outcome, written as $[p_1 : o_1, \dots, p_k : o_k]$, where each $o \in O$, each $p_i \geq 0$ and $\sum_{i=1}^k p_i = 1$. For example, one could see tossing an unbiased coin as the following lottery: $[0.5 : heads, 0.5 : tail]$. Notice that this is also a bit counterintuitive, because before we looked at an outcome as a *choice* (for example, in Figure 2.1 the player clearly has a choice between Confess or Not Confess, there was no probability involved). The trick is that we can use the characteristics of probabilities to specify the utility that an agent assigns to an outcome later. Axiom 3 states that outcomes with equal preference can be substituted in a lottery, given that the rest of the lottery remains unchanged.

Axiom 3 (Substitutability). *If $o_1 \sim o_2$, then for all sequences of one or more outcomes o_3, \dots, o_k and sets of probabilities p, p_3, \dots, p_k for which $p + \sum_{i=3}^k p_i = 1$, $[p : o_1, p_3 : o_3, \dots, p_k : o_k] \sim [p : o_2, p_3 : o_3, \dots, p_k : o_k]$.*

The next axiom states that an agent is indifferent between lotteries that induce the same probabilities. For example, having a lottery within a lottery, for example: $l_1 = [0.3 : o_1, 0.7 : [0.8 : o_2, 0.2 : o_1]]$, is equal to having one lottery that contains all the outcomes at once: $l_2 = [0.44 : o_1, 0.56 : o_2]$. We denote the probability of an outcome o_1 to be selected in a lottery l with $P_l(o_i)$. So in the previous example we have $P_l(o_1) = P_{l'}(o_1) = 0.44$.

Axiom 4 (Decomposability). *If $\forall o_i \in O : P_{l_1}(o_i) = P_{l_2}(o_i)$ then $l_1 \sim l_2$.*

The next axiom states that agents prefer more of a good thing. If an agent prefers o_1 to o_2 and considers two lotteries over these outcomes, he prefers the lottery that assigns the larger probability to o_1 .

Axiom 5 (Monotonicity). *If $o_1 \succ o_2$ and $p > q$ then $[p : o_1, 1-p : o_2] \succ [q : o_1, 1-q : o_2]$.*

The last axiom of Continuity assumes that there is a "tipping point" between being better than and worse than a given middle option.

Axiom 6 (Continuity). *If $o_1 \succ o_2$ and $o_2 \succ o_3$, then $\exists_{p \in [0,1]}: o_2 \sim [p : o_1, 1 - p : o_3]$.*

From these six axioms we can infer a single-dimensional *utility function* that the agent wants to maximize.

Theorem 3 (von Neumann and Morgenstern, 1944). *If a preference relation \succeq satisfied the axioms completeness, transitivity, substitutability, decomposability, monotonicity, and continuity, then there exists a function $u : O \mapsto [0, 1]$ with the properties that:*

1. $u(o_1) \geq u(o_2)$ iff $o_1 \succeq o_2$; and
2. $u([p_1 : o_1, \dots, p_k : o_k]) = \sum_{i=1}^k p_i u(o_i)$.

In this theorem, the utilities are defined for the range $[0, 1]$, but it follows logically that if $u(o)$ is a utility function for a given agent then $u'(o) = au(o) + b$ is also a utility function for the same agent, as long as a and b are constants and a is positive. In other words; we can use any numerical for the utility of the agents.

Chapter 3

The Extensive Form: Games with Sequential Actions

3.1 Perfect information extensive-form games

An *extensive-form game* represents turn-taking in games; it does not always assume that players act simultaneously. It can be transformed to a normal-form representation of Chapter 3, so the game-theoretic properties such as Nash equilibria are preserved. It is exponentially smaller to represent than its induced normal form and there are other solution concepts such as subgame-perfect equilibria, which we will discuss below.

An extensive-form game is represented as a graph-theoretic tree (also called a *game tree*, e.g. Figure 3.1), where each node represents a choice of one of the players and each edge a possible action. The root is the starting point and the leafs are possible outcomes for which players have preferences (utility).

Definition 20 (Perfect-information game). *A (finite) perfect-information game (in extensive form) is a tuple $G = (N, A, H, Z, \chi, \rho, \sigma, u)$, where:*

- N is a set of n players;
- A is a (single) set of actions;
- H is a set of nonterminal choice nodes;
- Z is a set of terminal nodes, disjoint from H ;
- $\chi : H \mapsto 2^A$ is the action function, which assigns to each choice node a set of possible actions;
- $\rho : H \mapsto N$ is the player function, which assigns to each nonterminal node a player $i \in N$ who chooses an action at that node;
- $\sigma : H \times A \mapsto H \cup Z$ is the successor function, which maps a choice node and an action to a new choice node or terminal node such that for all $h_1, h_2 \in H$ and $a_1, a_2 \in A$, if $\sigma(h_1, a_1) = \sigma(h_2, a_2)$ then $h_1 = h_2$ and $a_1 = a_2$;
- $u = (u_1, \dots, u_n)$, where $u_i : Z \mapsto \mathbb{R}$ is a real-valued utility function for player i on the terminal nodes Z .

The sets H and Z represent the nodes of the graph, and the function σ the arcs, which are mappings between nodes parameterized with an action. Each player is assigned a set of nodes (the ρ function), and at each node he has a set of possible actions (the χ function). Since this is a *tree*, each node has a single parent and we thus can always determine the history of the game.

The following two notions can be considered important terminology:

- *subtree*: for each node h , its subtree is the tree rooted at h .
- *descendants* (or *children*): all the nodes in the subtree of h .

Just like in a norm-form game, the extensive-form game also has a notion of *pure strategy*.

Definition 21 (Pure strategy in extensive-form game). *A pure strategy in an extensive-form game $G = (N, A, H, Z, \chi, \rho, \sigma, u)$ for player i is a complete specification of what action to take at every node belonging to that player; it consists of the Cartesian product $\prod_{h \in H, \rho(h)=i} \chi(h)$.*




Figure 3.1: The Sharing game.

For illustrative purposes, imagine the following game.

A brother and a sister are in the process of sharing two indivisible presents. First, the brother may choose how he wants to divide the presents. Then the sister has the possibility to accept the proposal (Yes) or to decline it (No).

This game is visualized as an extensive game in Figure 3.1. The first choice of the brother is represented as the root node (player 1) with three outgoing arcs (representing actions) which define the way the presents can be divided by the brother. Enumerating the strategies (or possible actions) of the brother is fairly straightforward: $S_1 = \{2-0, 1-1, 0-2\}$. The strategies of the sister are somewhat less intuitive. Remember that we need to specify all possible combinations of actions *for each node*. Since there are three choice points for player 2 and two actions for each choice point, we end up with $2^3 = 8$ strategies:

$$S_2 = \{(yes, yes, yes), (yes, yes, no), (yes, no, yes), (yes, no, no), (no, yes, yes), (no, yes, no), (no, no, yes), (no, no, no)\}$$

3.1.1 From extensive-form to normal-form

It is fairly easy to transform this game to a normal-form game. For each combination of strategies in S_1 and S_2 we denote the resulting utility (Figure 3.2). The pure Nash equilibria can be read from the table directly and are denoted in bold. We see now that the extensive form is indeed much more compact than its induced normal form, and that it can result in a certain redundancy within the normal form; in the extensive-form game there are 6 outcomes, while there are 24 in its induced normal form. Transformation

	(y,y,y)	(y,y,n)	(y,n,y)	(y,n,n)	(n,y,y)	(n,y,n)	(n,n,y)	(n,n,n)
(2-0)	(2,0)	(2,0)	(2,0)	(2,0)	(0,0)	(0,0)	(0,0)	(0,0)
(1-1)	(1,1)	(1,1)	(0,0)	(0,0)	(1,1)	(1,1)	(0,0)	(0,0)
(0-2)	(0,2)	(0,0)	(0,2)	(0,0)	(0,2)	(0,0)	(0,2)	(0,0)

Figure 3.2: The induced normal-form game of Figure 3.1.

from extensive form to normal form is always possible, but it can result in an exponential blowup of the game representation.

Theorem 4. *Every (finite) perfect-information game in extensive form has a pure-strategy Nash equilibrium*

But, there does *not* always exist an extensive form game for a normal form. For example: the Prisoner’s Dilemma in Figure 2.1 cannot be transformed to an extensive form. Intuitively, this is because perfect-information extensive-form games cannot model simultaneity. In Section 3.2 we will consider a class of games in which this is possible.

In more complex games, the induced normal form may contain different strategies that result in the same action. To understand this, consider the extensive-form representation of Figure 3.3. As soon as player 1 chooses A as his first action, it is no longer relevant whether he chooses G or H . Still, if we look at the induced normal-form representation in Figure 3.4, we see that both (A, G) and (A, H) are included.

Also, the equilibria in the normal-form game may feel unsatisfying when considering them in the extensive-form. Take for example the equilibrium $\{(B, H), (C, E)\}$. Notice that if player 1 would play (B, G) instead, player 2 would respond with (C, F) . So we could interpret the reason for player 1 playing H instead of G as a *threat*; there would be no other reason since G would actually result in a higher payoff than H for both players. Still, if player 2 were to play (C, F) anyway, would player 1 really following through his treat and play H , or would he pick G ?

3.1.2 Subgame-perfect equilibrium

Because of the previous intuition that some equilibria in the norm-form game are somewhat unsatisfying in the extensive-form game, we have another solution concept for extensive-form games: the *subgame-perfect equilibrium*. Before we define this, we need to understand the notion of a *subgame*.




Figure 3.3: A perfect-information game in extensive form.

	(C,E)	(C,F)	(D,E)	(D,F)
(A,G)	(3,8)	(3,8)	(8,3)	(8,3)
(A,H)	(3,8)	(3,8)	(8,3)	(8,3)
(B,G)	(5,5)	(2,10)	(5,5)	(2,10)
(B,H)	(5,5)	(1,0)	(5,5)	(1,0)

Figure 3.4: The induced normal-form game of Figure 3.3.

Definition 22 (Subgame). *The subgame of an extensive game G with root h is the restriction of G to the descendants of h . The set of subgames of G include all subgames of G rooted at some node in G .*

Definition 23 (Subgame-perfect equilibrium (SPE)). *The subgame perfect equilibrium of extensive game G are all strategy profiles s such that for any subgame g' of G , the restriction of s to G' is a Nash equilibrium of G' .*

So the subgame perfect equilibria are the strategy profiles that are Nash equilibria in all subgames. Notice that a subgame of a game G is its own subgame, so every SPE is a Nash equilibrium. This is not always the case the other way around: not every NE is a SPE. Still, every perfect-information game in extensive form has at least one subgame-perfect equilibrium. Can you find the SPE of Figure 3.1¹ and Figure 3.3²?

¹Player two will never play $(*, n, *)$ or $(*, *, n)$, so the only Nash equilibria that are also SPE are $((2, 0), (y, y, n))$ and $((1 - 1), (n, y, y))$.

²The only SPE is (AG, CF) , because playing H is not subgame perfect for player 1 (consider the subgame rooted at the second choice node of player 1, here G is subgame perfect). Therefore all NE with H can be dropped.

3.1.3 Computing equilibria: backward induction

Backward induction is the most straightforward way to find NE using subgame-perfect equilibria. The idea is to identify the equilibria in the "bottom-most" subgames and work your way up. It is computationally simple (linear in size of the game tree). The algorithm goes as follows (see next page for a sketch of the algorithm):

Algorithm 1 Finding the value of a sample (subgame-perfect) Nash equilibrium

```
1: function BackwardInduction (node  $h$ ) :  $u(h)$ 
2: if  $h \in Z$  then
3:   return  $u(h)$ 
4: end if
5:  $best\_util \leftarrow -\infty$ 
6: for all  $a \in \chi(h)$  do
7:    $util\_at\_child \leftarrow \text{BackwardInduction}(\sigma(h, a))$ 
8:   if  $util\_at\_child_{\rho(h)} > best\_util_{\rho(h)}$  then
9:      $best\_util \leftarrow util\_at\_child$ 
10:  end if
11: end for
12: return  $best\_util$ 
13: end function
```

First we choose the node h of which we want to compute the value of a subgame-perfect Nash equilibrium; we choose the root node. Line 2 can be ignored, because we are not considering a leaf node. We initialize $best_util$ with $-\infty$ (Line 5) and iterate over all possible actions (Line 6). For each action we determine the utility of the node that this action leads to (Line 7). Then we take the maximum of this utility *for the player who's turn it is* (Line 8 and 9). Lets walk through this algorithm using Figure 3.3 as an example again. In the root node there are two actions, A and B . For action A we have again two action C and D . These are the leaf nodes and thus return the utilities $(3, 8)$ and $(8, 3)$. Since it is player 2 his turn to choose between C and D , he chooses utility 8 as his $best_util$. As might be clear, this technique leads to a *labeling* of the nodes, as shown in Figure 3.5. In the root node we now have the value of the subgame-perfect equilibrium. The way to find this subgame-equilibrium is very trivial, it are simply the choices that the players made: $(A, G), (C, F)$.

There are several problems with the backward induction algorithm:

- We often cannot enumerate the entire tree in advance, and we need to make assumptions to simplify the tree. An example of a technique for this is alpha-beta pruning in zero-sum games, which is not explained here.
- Only one SPE is found, not all. This can be done, but we do not discussed that here.
- There is a conceptual drawback. Consider the *centipede game* in Figure 3.6. The only SPE is for all players to always choose D . This is a problem because of two reasons:
 - It seems to be more intuitive (proven by experiments) to keep playing A until we reach the end of the chain and achieve a high payoff




Figure 3.5: Labeling of nodes in an extensive-form game using the backward induction algorithm.

- Imagine you are the second player, and you reach your first choice point (so player 1 has chosen A). What should you do? According to the SPE you should play D , but according to the same analysis it is not even possible to reach this point. You have reached a state which your analysis has given a probability zero: a *mesasure-zero event*. This is a fundamental problem in game theory, and it related to the topic of *belief revision*.




Figure 3.6: A perfect-information game in extensive form.

3.2 Imperfect-information extensive-form game

An *imperfect-information game* is an extensive-form game in which each player's choice nodes are partitioned into information sets; if two choice nodes are in the same information set then the agent cannot distinguish between them.

Definition 24 (Imperfect-information game). *An imperfect-information game (in extensive form) is a tuple $(N, A, H, Z, \chi, \rho, \sigma, \mu, I)$, where:*

- $(N, A, H, Z, \chi, \rho, \sigma, \mu)$ is a perfect-information extensive-form game;
- $I = (I_1, \dots, I_n)$, where $I_i = (I_{i,1}, \dots, I_{i,k_i})$ is a set of equivalence classes on (i.e., a partition of) $\{h \in H : \rho(h) = i\}$ with the property that $\chi(h) = \chi(h')$ and $\rho(h) = \rho(h')$ whenever there exists a j for which $h \in I_{i,j}$ and $h' \in I_{i,j}$.

Consider the imperfect-information extensive-form game of Figure 3.7. Player 1 has two information sets: the set including the top choice node, and the set including the bottom choice nodes. The two bottom choice nodes in the second information set have the same set of possible actions. We can interpret this as player 1 that does not know whether player 2 chose A or B when he makes his choice between l and r .




Figure 3.7: An imperfect-information game.

A pure strategy for an agent in an imperfect-information game selects one of the available actions in each information set of that agent.

Definition 25 (Pure strategies). *Let $G = (N, A, H, Z, \chi, \rho, \sigma, u, I)$ be an imperfect-information extensive-form game. Then the pure strategies of player i consist of the Cartesian product $\prod I_{i,j} \in I_i \chi(I_{i,j})$*

Using this type of game it turns out to be possible to model the Prisoners' Dilemma, as shown in Figure 3.8. In general, any normal-form game can be trivially transformed into an equivalent imperfect-information game.

Transforming an imperfect-information extensive game to a normal form game is done by enumerating the pure strategies of each agent, but now it selects one of the available




Figure 3.8: The Prisoner's Dilemma game in extensive form.

actions in each information set of that agent. We can distinguish between two kinds of strategies in the imperfect-information game:

- *mixed strategies*; this is simply the set of mixed strategies in its image normal-form game.
- *behavioral strategies*: independent randomization at each information set.




Figure 3.9: A game with imperfect recall.

	U	D
L	1,0	1,0
R	5,1	2,2

Figure 3.10: Its induced normal form.

To understand the difference between these strategies, consider the game in Figure 3.9. First we turn this game into its induced normal form. Player 1 only has one information set containing both its choice nodes. Player 2 has one information set with its only choice node. This leads to the normal form game in Figure 3.10. A mixed strategy for player 1 is to play L with probability p and R with probability $1 - p$. Because R is a strictly dominant strategy for player 1 and D is a strict best response, (R, D) is the unique Nash equilibrium. Notice that the payoff of 100 for player 1 has been completely left out, because it is unreachable for a mixed strategy.

If we on the other hand consider the behavioral strategy of player 1, he will randomize at each time he finds himself in the information set. Since D is the dominant strategy for player 2, player 1 computes the best response to D as follows. Using the behavioral strategy $(p, 1 - p)$, his expected payoff is

$$\begin{aligned} u_1(s) &= p \times p \times 1 + p \times (1 - p) \times 100 + (1 - p) \times 0 \times 5 + (1 - p) \times 1 \times 2 \\ &= 1p^2 + 100p(1 - p) + 2(1 - p) \end{aligned}$$

whose maximum is obtained at $p = 98/198$. Thus we end up with a different equilibrium in behavioral strategies.

The games for which the expressive power of mixed and behavioral strategies coincide are called games of *perfect recall*. Every perfect-information game is a game of perfect recall.

Chapter 4

Communication

As said in Chapter 2, game theory studies the interaction between independent, self-interested agents. Therefore it makes sense that agents need to communicate. Communication is one of the defining characteristics of a multiagent system. It is characterized by:

- *syntax*: the form of the communication (words, sentences, dialogues and conversations).
- *semantics*: the meaning, relation between the communication form and the meaning it carries.
- *pragmatics*: the context of the communication and how it is influenced.

4.1 "Doing by talking" I: Cheap talk

In cheap talk, players can communicate before taking actions. Communication is *costless*, it does *not need to be truthful* so it does *not imply commitment*. Cheap talk can be seen as a two-stage game: first the players communicate, and then they decide action.

Consider the Prisoners' Dilemma once more, reproduced in Figure 4.1. The unique Nash equilibrium is (*Confess, Confess*), but it is not Pareto optimal: playing (*Not Confess, Not Confess*) is better for both players. Will the outcome of this game change if players are allowed to communicate? The answer is *no*, because *Confess* remains the safest option for the player, if he cannot trust the other player. Players that choose to communicate can not necessarily be trusted. Therefore the talk of the agents is *cheap*; not credible and therefore useless.

	Not Confess	Confess
Not Confess	2,2	0,3
Confess	3,0	1,1

Figure 4.1: *The Prisoner's Dilemma game.*

	L	R
U	1,1	0,0
D	0,0	1,1

Figure 4.2: *The Coordination game.*

Consider now in contrast the Coordination game reproduced in Figure 4.2. In this case it is in the interest of the players to tell the truth before playing. For example, if player 1 declares "I will play U", it is not in his interest to lie about this. This utterance is both *self-committing* and *self-revealing*:

- **self-committing:** Once uttered, and assuming it is believed by the other player, the declared action is the optimal one.
- **self-revealing:** Assuming it is uttered with the expectation that it will be believed, it is uttered only when it was the intention to act that way.

Let us analyze two games using these notions. Consider the game of Stag hunt in Figure 4.3, where the row player declares "I will play Stag". Is this utterance self-committing or self-revealing?

- *Self-committing:* yes, because if row player thinks column player believes his utterance (meaning that the column player will play Stag to maximize its outcome, since players are assumed to be rational), then the declared action Stag is optimal for the row player, since any other (mixed) strategy will result in a lower payoff.
- *Self-revealing:* no, because the row player would like the column player to believe the utterance, because there is a strategy for which the row player is better off by not playing Stag when announcing to do so. In general, to find out whether such a strategy exists we should assume that both players currently play an equilibrium strategy. This is because if you communicate a strategy which is not your side of a Nash Equilibrium, the utterance is already not self-committing (your best reply to the opponents best reply will never be the same as the declared strategy), so the matter of self-revealingness is not relevant in this situation. Assume that both players play the mixed equilibrium strategy $(\frac{7}{8}, \frac{1}{8})$. The expected utility in this equilibrium is 7.875 for both players. Now the row player announces: "I will play Stag". If the column player believes this utterance, he will best respond by playing Stag. Now if row player, instead of doing what he said, would remain playing the equilibrium strategy, the expected utility of the column player would remain 7.875, while that of the row player would increase to 8.875 (the reader should verify this). Therefore the row player would like the column player to believe the utterance, because it would increase its payoff while the payoff of the column player does not increase.

In general, self-revelation plus self-commitment is very credible; self-commitment alone is not credible. Now consider the game depicted in Figure 4.4. What is the status of the utterance "I will play D" by player 1? Answer in the footnote¹. What about "I will play L" by the column player?²

In the context of communication we define two notions of an equilibrium:

- *Babbling equilibrium:* an equilibrium in which one party send a meaningless signal and the other party ignores it (perhaps because the sending player is judged incredible). Every cheap talk game has a babbling equilibrium.

¹*self-committing:* no, because if column player believes the utterance, he will play *R*. Then it would be optimal for player 1 to play *T* instead of *D*. *self-revealing:* in general, an utterance cannot be not self-committing and self-revealing at the same time, but we can also derive it if we assume that the players play the equilibrium strategy $s = (s_1, s_2)$ where $s_1 = (\frac{1}{4}, \frac{3}{4})$ and $s_2 = (\frac{1}{2}, \frac{1}{2})$. If player 2 were to believe the utterance, player 1 would benefit from it while player 2 his outcome would remain the same.

²both not self-revealing and self-committing using similar arguments are in the previous example.

	Stag	Hare
Stag	9,9	0,8
Hare	8,0	7,7

Figure 4.3: *Stag hunt game*.

	L	R
T	0,4	3,1
D	1,2	2,3

Figure 4.4: *Fictive game*.

- *Revealing equilibrium*: an equilibrium that is not a babbling equilibrium, which means that the signals sent by one party are meaningful. This can be because the party is judged perfectly honest or perfectly dishonest (which results in the other party to always believe the opposite of what is signaled).

4.2 "Talking by doing": signaling games

In the games we discussed so far, communication was in the context of games with perfect information. The only information that can be revealed in these games is the intention for the players to act in a certain way. We will now consider games of incomplete information, which allows for the opportunity to reveal one's own private information prior to acting. We consider a class of imperfect-information games called *signaling games*. In these games, the game to play is decided by chance (often called *Nature*).

Definition 26 (Signaling games). *A signaling game is a two-player game in which Nature selects a game to be played according to a commonly known distribution, player 1 is informed of that choice and chooses an action, and player 2 then chooses an action without knowing Nature's choice, but knowing player 1's choice.*

Consider the following example.

The Acme Corporation wants to hire Sally into one of two positions: a demanding and an undemanding position. Sally may have high or low ability. Sally prefers the demanding position if she has high ability (because of salary and intellectual challenge) and she prefers the undemanding positions if she instead has low ability (because it will be more manageable). Acme too prefers that Sally be in the demanding position if she has high ability, and that she be in the undemanding position if she is of low ability.

The perfect information version of this game is shown in Figure 4.5.

Only Sally knows what her true ability level is, but before they play the game, Sally can send Acme a signal about her ability level, meaning she can make the utterance "I have high ability" or "I have low ability". We model this as a signaling game by letting Nature decide which game to play, i.e., Nature decides if Sally has a high or a low ability (with equal probability). Then Sally, who knows her ability (she knows which game is being played), sends a signal to Acme about her choice (her ability level). Then, Acme decides what action to take. The situation is modeled by the two games in Figures 4.6 and 4.7.

What signal should Sally send? It is obvious that she should tell the truth; if she would lie and Acme believes her, she would receive a lower payoff than if she had told

		Job Acme gives Sally	
		Demanding	Undemanding
Sally's ability	High	2,1	0,0
	Low	0,0	1,3

Figure 4.5: *Perfect information version of the Job Hunt game.*

the truth. Therefore she should signal High ability when the left game is selected. Acne knows that Sally tells the truth (the messages are credible), so will respond with the demanding job. The message is therefore *self-signaling*; assume Sally will be believed, she will send the message only if it is true.

Signal high ability	3,1
Signal low ability	0,0

Figure 4.6: *High-ability game.*

Signal high ability	0,0
Signal low ability	2,1

Figure 4.7: *Low-ability game.*

Consider another example depicted in Figures 4.8 and 4.9. Again, the row player knows which game is selected by Nature (with equal probability). Row player chooses his message (U or D) and Column player (who does not know which game is being chosen by Nature) will choose his action (L or R).

	L	R
U	4,-4	1,-1
D	3,-3	0,0

Figure 4.8: *Incomplete information game #1.*

	L	R
U	1,-1	3,-3
D	2,-2	5,-5

Figure 4.9: *Incomplete information game #2.*

In this game, it is actually an advantage for the Row player to exploit the privileged information. Assume that Row player selects/signals U . Now Column player believes that game #1 is selected, since it is the dominant strategy for Row player in this game. Therefore Column player will select its dominant strategy in this game, which is R . Similarly, if Row player signals D , Column player will respond with L . From this we can compute the expected utility for Row player if he were always telling the truth: when game #1 is selected, (U, R) will be played, and game #2 will result in (D, L) . Each game is selected with equal probability.

$$E(u_1) = \frac{1}{2} \times 1 + \frac{1}{2} \times 2 = 1.5$$

A better strategy for Row player is actually to always play D . We calculate the expected payoff for player 1 when playing this strategy, assuming that player 2 plays L with probability p and R with probability $(1 - p)$ as follows:

$$E(u_1) = \frac{1}{2}(3p + 0(1 - p)) + \frac{1}{2}(2p + 5(1 - p)) = 2.5$$

So although player 1 has privileged information, it may not always be to his advantage to exploit it (share it). Thus, in some cases player 1 can receive a higher payoff by ignoring this information.

4.3 "Doing by Talking" II: Speech-act theory

4.3.1 Speech acts

Speech-act theory claims that some communication can be viewed as actions, intended to achieve some goal, instead of simply sharing information. It distinguishes between three different kinds of speech acts (levels of analyzing an utterance):

- *locutionary act*: the act of saying something, the emission of a signal carrying a meaning; the actual message. e.g. "there is a car coming your way". Locutions establish a proposition, which may be true or false.
- *illocutionary act*: the performance of a conventional force on the receiver through the utterance. e.g. "warn", "request", "inform".
- *perlocutionary act*: the intention of the speaker; the effect of the illocutionary act on the hearer. e.g. making sure the hearer avoids a car.

The philosopher P. Grice's *cooperative principle* states that humans seem to undertake the act of conversation cooperatively. Humans generally seek to understand and to be understood. It is in both parties' best interest to communicate clearly and efficiently. *Gricean maxims* are four basic rules that humans use to achieve the cooperative principle:

- *quantity*: provide exactly the amount of information required in the conversation.
- *quality*: provide information that is true or believed to be true.
- *relation*: provide information that is relevant to the conversation.
- *manner*: provide information clearly and briefly.

The *implicature* contains the meaning that can be derived from the actual communication between humans; which is generally a lot. The *conversation implicature* is the implied meaning (implicature) that relies on the fact that the hearer assumes that the speaker is following the Gricean maxims. This may help to avoid unwanted implicatures. This might result in more efficient communication between artificial agents.

4.3.2 A game-theoretical view of speech acts

We will try to model the process of two agents -the speaker and the hearer- involved in communication as a game in the sense of game theory. We will illustrate this via the phenomenon of *disambiguation* in language. We will analyze the following sentence-level ambiguity:

Every ten minutes a person gets mugged in New York City.

This sentence could have two meanings: a different person is mugged every ten minutes in New York City; or there is a very unfortunate individual walking around in New York City that gets mugged every ten minutes. We could model this as a game with player A (the speaker) and B (the hearer) that consists of the following steps:

1. There exist two situations (or games)
 - s**: Muggings of different people take place every ten minutes in NYC.
 - t**: The same person is repeatedly mugged every ten minutes in NYC.
2. Nature selects randomly between s and t .
3. Nature's choice is revealed to A but not to B .
4. A decides between uttering one of three possible sentence:
 - p**: "Every ten minutes a person gets mugged in New York City."
 - q**: "Every ten minutes some person or another gets mugged in New York City."
 - r**: "There is a person who gets mugged every ten minutes in New York City."
5. B hears A , and must decide whether s or t obtain.

We now simplify the game by decreasing the number of available choices, and add probabilities for the choice points. Finally we add payoffs which makes sense, resulting in the extensive game depicted in Figure 4.10. We then can analyze this game using the techniques discussed in Section 3.1, allowing us to find the following two equilibria:

1. A 's strategy: say q in s and r in t . B 's strategy: when hearing p , select between the s and t interpretations with equal probabilities.
2. A 's strategy: say p in s and r in t . B 's strategy: when hearing p , select the s interpretation.

Notice that this example is very concise and purely meant as an illustration of how communication can be structured between agents using game theory.

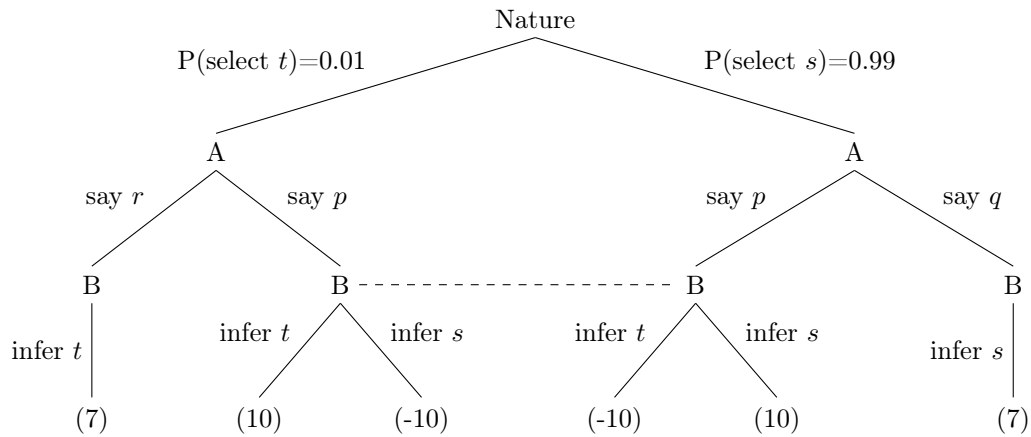


Figure 4.10: Communication as a signaling game.

4.3.3 Agent communication languages

Several proposals have been made for artificial languages to serve as the medium for communication. We will describe three shortly.

- **KQML:** *Knowledge Query and Manipulation Language*. It defines the outer communication language and provides a set of speech acts (ask-if, perform, tell, reply). An example of a KQML message:

```
(tell
  :sender      Agent1
  :receiver   Agent2
  :language   KIF
  :ontology   Blocks-World
  :content    (AND (Block A) (Block B) (On A B)))
```

- **KIF:** *Knowledge Interchange Format*. It defines the message content.

```
A to B:  (ask-if (> (size chip1) (size chip2)))
B to A:  (reply true)
B to A:  (inform (= (size chip1) 20))
B to A:  (inform (= (size chip1) 18))
```

- **FIPA:** *Foundation for Intelligent Physical Agents*.

```
(inform
  :sender agent1
  :receiver agent5
  :content (price good200 150)
  :language s1
  :ontology hpl-auction
)
```

Chapter 5

Aggregating Preferences: Social Choice

In the previous chapters we looked at game theory from the agent perspective: what choices should an agent make as to optimize its expected utility? We now look at the "designer perspective": what rules should be the authority (the "designer") invent in order to enforce agents to act in a certain way? In this chapter we will deviate from game theory, but in the next chapter on Mechanism design we will apply the concepts of this chapter to game theory.

5.1 Introduction

A simple example of a situation in which the designer perspective matters is voting. How should a central authority pool the preferences of different agents so as to best reflect the wishes of the population as a whole? Consider the following example:

You are babysitting three children - Wil, Liam, and Vic - and need to decide on an activity for them. You can choose among going to the video arcade (a), playing basketball (b), and driving around in a car (c). Each kid has a different preference over these activities, which is represented as a strict total ordering over the activities and which he reveals to you truthfully. $a \succ b$ means that a is strictly preferred over b . The children hold the following preferences:

Will: $a \succ b \succ c$

Liam: $b \succ c \succ a$

Vic: $c \succ b \succ a$

The question is now: what would be the fairest choice? A possibility is to apply *plurality voting*; pick the activity with the largest number of votes. This problem faces several problems:

- It needs a tie-breaking rule that selects between outcomes of equal preference (in the case of our example all the activities are equally often chosen as first choice).

A tie-breaking rule of our example could be to select the activities alphabetically; this would make a the winner..

- It does not necessarily meet the *Condorcet condition*. The Condorcet condition states that if there exists a candidate x such that if for all other candidates y at least half the voters prefer x to y , then x must be chosen. The Condorcet rule would choose b , since two of the three children prefer b to a , and likewise prefer b to c .

From this example it might seem that the Condorcet winner gives us a satisfying outcome. But this is not always the case, as shown in the next example.

Will: $a \succ b \succ c$

Liam: $b \succ c \succ a$

Vic: $c \succ a \succ b$

Here the Condorcet condition does not tell us what to do, neither does plurality voting give us satisfying result. Social choice is not a straightforward manner, and offers many solutions. We will consider more voting methods in Section 9.3.1., but first we establish a formal model of social choice.

5.2 A formal model

Let $N = \{1, 2, \dots, n\}$ denote a set of agents (in our first example this were the children), and let O denote a finite set of outcomes/candidates (the activities of our example). Let L_{-} denote the set of all possible orderings. We define two kinds of social functions: social choice functions and social welfare functions. *Social choice functions* select one of the alternatives given an ordering. For example, plurality can be seen as a social choice function.

Definition 27 (Social choice function). *A social choice function is a function $C : L_{-}^n \mapsto O$.*

Social choice correspondences select a set of candidates instead of a single one.

Definition 28 (Social choice correspondence). *A social choice correspondence is a function $C : L_{-}^n \mapsto 2^O$.*

Let $\#(o_i \succ o_j)$ denote the number of agents who prefer outcome o_i to outcome o_j .

Definition 29 (Condorcet winner). *An outcome $o \in O$ is a Condorcet winner if $\forall o' \in O : \#(o \succ o') \geq \#(o' \succ o)$.*

We already know that the Condorcet condition is not guaranteed to give an outcome (recall our example). The following condition does guarantee an outcome.

Definition 30. *The Smith set is the smallest set $S \subseteq O$ with $\forall o \in S, \forall o' \notin S, \#(o \succ o') \geq \#(o' \succ o)$.*

That is, every outcome *in* the Smith set is preferred by at least half of the agents to every outcome *outside* the set. This set always exists and moreover, when the Condorcet

winner exists, then the Smith set is a singleton consisting of the Condorcet winner. The Smith set is sometimes also referred to as the *weak condorcet winners*.

Social welfare functions produce a total ordering on the set of alternative, and are thus more informative than social choice functions.

Definition 31 (Social welfare function). *A social welfare function is a function $W : L_-^n \mapsto L_-$.*

Definition 32 (Social welfare correspondence). *A social welfare correspondence is a function $W : L_-^n \mapsto 2^{L_-}$.*

The use of these functions will become clear when discussing Arrow's famous impossibility theorem in Section 9.4.

5.3 Voting

5.3.1 Voting methods

We will now survey some important voting methods and discuss their properties:

- *Plurality voting*: each voter makes a single vote (**non-ranking voting**) and the candidate with most votes wins.
- *Cumulative voting*: each voter makes k votes (multiple votes on one candidate allowed) and the candidate with most votes wins.
- *Approval voting*: each voter makes a single vote for as many candidates as he wishes and the candidate with most votes wins.
- *Plurality with elimination*: Each voter makes a single vote for the most preferred candidate. The candidate with the fewest votes is eliminated and the voting repeats with the remaining candidates until one candidate remains. This allows voters to express their full preference orderings (called *ranking voting*).
- *Borda voting*: Each voter submits a full ordering on the candidates. If there are n candidates, it contributes $n - 1$ points to the highest ranked candidate, $n - 2$ to the second-highest, and so on; no points are given to the lowest ranked candidate. The winners are those whose total sum of points for all the voters is maximal. The Borda rule does not always provide a social choice function, but a social choice correspondence.
- *Pairwise elimination*: Voters are given a schedule for the order in which pairs of candidates will be compared. Given two candidates, the candidate that each voter prefers is determined. The candidate who is preferred by a minority of votes is eliminated, and the next pair of non-eliminated candidates in the schedule is considered. This continues until only one candidate remains.
- *Majority voting*: The outcome that is preferred by more than half of the voters is the winner.

5.3.2 Voting paradoxes

One might wonder why there are so many different voting methods. This is because there does not seem to be one voting method that is appropriate for all circumstances. Voting schemes that seem reasonable can often fail in surprising ways, as is shown in the following examples.

- **Condorcet condition.** Consider 1000 agents with the following preferences:

499 agents: $a \succ b \succ c$

3 agents: $b \succ c \succ a$

498 agents: $c \succ b \succ a$

Who is the Condorcet winner in this case? Answer in footnote.¹ Does this match with plurality voting?² And with plurality voting with elimination?³ And what about Borda voting?⁴

- **Sensitivity to a losing candidate.** Consider 100 agents with the following preferences:

35 agents: $a \succ c \succ b$

33 agents: $b \succ a \succ c$

32 agents: $c \succ b \succ a$

Who are the winners if we vote according to plurality, Borda, Codorcet?⁵ What would plurality pick if candidate c did not exist?⁶ What can you thus say about candidate c ?⁷

- **Sensitivity to the agenda setter.** In some cases, the agenda setter of pairwise elimination can select whichever outcome he wants by selecting the appropriate elimination order. Reconsider the preferences in the previous example. Pairing first a and b will eliminate a . Then c is chosen in the pairing of b and c . Alternatively, starting with a and c eliminates c . Then b is chosen in the pairing between a and b . Finally, if we start with b and c , b is eliminated and ultimately a is chosen. So the agenda setter is able to let a , b , or c win, depending on the order of elimination.

5.4 Existence of social function

We will now try to understand the paradoxes above by describing some results formally. First we define three notions of so-called *fairness*. Finally we will state *Arrow's impossibility theorem*, showing that, when defining a social welfare function, it is impossible to achieve all these notions of fairness at once! Assume that W is a social welfare function.

Definition 33 (Pareto efficiency (PE)). W is Pareto efficient if for any $o_1, o_2 \in O, \forall i : o_1 \succ_i o_2$ implies that $o_1 \succ_W o_2$.

The intuition behind this definition is that if alternative o_1 is unanimously preferred to alternative o_2 , o_1 should be ranked higher than o_2 in the social ordering.

¹ b is the Condorcet winner, since 501 out of 1000 people prefer b to a , and 502 prefer b to c .

²No, since a wins by plurality voting.

³Also not, since plurality voting with elimination will first eliminate b . After this, 499 agents prefer a over c , while 501 prefer c over a . Therefore c is selected as the winner.

⁴Borda voting does select b as a winner with 1003 points. c follows with 999 and a is last with 998.

⁵plurality: a , Borda: a , Condorcet: a

⁶candidate b .

⁷it stands no chance of being selected, but it acts as a "spoiler", changing the selected outcome.

Definition 34 (Independence of irrelevant alternatives (IIA)). W is independent of irrelevant alternatives if, for any $o_1, o_2 \in O$ and any two preference profiles $[\succ'], [\succ''] \in L^n, \forall i : (o_1 \succ'_i o_2 \text{ if and only if } o_1 \succ''_i o_2)$ implies that $(o_1 \succ_{W([\succ'])} o_2 \text{ if and only if } o_1 \succ_{W([\succ''])} o_2)$.

This really complicated definition states that the social preference as returned by the social welfare function of two alternatives only depends on the *relative ordering* of these two alternatives in the individual preference relations. So as long as the relative ordering between two candidates does not change, changing the preference of other candidates should never change the social preference of these two candidates.

Definition 35 (Nondictatorship). W does not have a dictator if $\neg \exists i : \forall o_1, o_2 (o_1 \succ_i o_2 \Rightarrow o_1 \succ_W o_2)$.

Meaning there does not exist a single agent whose preferences always determine the social ordering. Now for the most important theorem in social choice theory:

Theorem 5 (Arrow, 1951). If $|O| \geq 3$, any social welfare function W that is Pareto efficient and independent of irrelevant alternatives is dictatorial.

This theorem tells us that we cannot hope to find a voting scheme that satisfies all of the notions of fairness that we find desirable. Therefore there is no hope for general social welfare functions (the identification of a social ordering over *all* outcomes), but maybe this problem is too hard. What about social welfare functions (i.e., only one outcome)? We first define the notions of fairness in the scope of social choice functions.

Definition 36 (Weak Pareto efficiency). A social choice function C is weakly Pareto efficient if, for any preference profile $[\succ] \in L^n$, if there exist a pair of outcomes o_1 and o_2 such that $\forall i \in N, o_1 \succ_i o_2$, then $C([\succ]) \neq o_2$.

In other words, if all agents prefer o_1 to o_2 , the social choice rule cannot choose o_2 .

Definition 37 (Monotonicity). C is monotonic if, for any $o \in O$ and any preference profile $[\succ] \in L^n$ with $C([\succ]) = o$, then for any other preference profile $[\succ']$ with the property that $\forall i \in N, \forall o' \in O, o \succ'_i o'$ if $o \succ_i o'$, it must be that $C([\succ']) = o$.

Intuition: improving one alternative should not influence the relative ordering of other alternatives. So an outcome o must remain the winner whenever the support for it is increased relative to a preference profile under which o was already winning. This definition is similar to the IIA condition (Definition 31).

Definition 38 (Nondictatorship). C is non-dictatorial if there does not exist an agent j such that C always selects the top choice in j 's preference ordering.

The sad conclusion for social welfare functions unfortunately pertains in social choice functions.

Theorem 6 (Muller-Satterthwaite, 1977). If $|O| \geq 3$, any social choice function C that is weakly Pareto efficient and monotonic is dictatorial.

Consider for example plurality voting. Clearly, it satisfies weak Pareto efficiency and is not dictatorial. According to the previous theorem, this means that plurality voting

has to be nonmonotonic. To see why, consider the following scenario with seven voters.

3 agents: $a \succ b \succ c$
 2 agents: $b \succ c \succ a$
 2 agents: $c \succ b \succ a$

Under these preferences, plurality chooses candidate a . Now imagine that the ground of last two agents increase their support for a in the following way:

3 agents: $a \succ b \succ c$
 2 agents: $b \succ c \succ a$
 2 agents: $b \succ a \succ c$

Since the support for a is increased in the individual preference ordering, it must remain the winner to satisfy monotonicity. As we can see, plurality selects b as the winner, hence plurality voting is nonmonotonic.

5.5 Restrictions on preferences

5.5.1 Single-peaked preferences

Definition 39 (Single-peaked preferences). *Given a predetermined linear ordering of the alternative set O , a preference relation \succ , is single-peaked if there exists a point $p \in O$ (the peak) such that for all $o_1, o_2 \in O$ such that $p \geq o_1 > o_2$ or $o_2 > o_1 \geq p$ then $o_1 \succ_i o_2$.*

To find out whether a single-peaked preference profile exists for all the players, we need to order the preferences in such a way that each preference has one single peak. Consider for example the preferences in Figure 5.1. The horizontal axis contains the candidates, while the vertical axis shows the level of preference that each voter has for a candidate. Each voter has its own color. On the left side, the preferences are ordered in a way that reveals no single-peaked preference profile. But once we change the order of the preferences as shown on the right, a single-peaked preference profile reveals.




Figure 5.1: *Single-peaked preferences*

Definition 40 (Median voter rule). *Given a profile of single-peaked preferences \succ_1, \dots, \succ_n with peaks p_1, \dots, p_n (with respect to a predetermined linear ordering on alternatives), the median voter rule selects the median among some superset of the peaks.*

Theorem 7 (Black, 1948). *The median voter rule selects a weak Condorcet winner.*

When we apply the median voter rule to Figure 5.1, we first determine the center point of the peaks of the preferences. There are three peaks: a , b and c . b is exactly in the middle of these three, so the median voter rule selects b , which is then also the weak condorcet winner. If the center of the peaks is not exactly on a candidate, we choose the candidate one which the most single-peaked preferences can be found.

5.5.2 Dichotomous preferences

The intuition of dichotomous preferences is that they only distinguish between "good" and "bad" alternatives.

Theorem 8 (Inada, 1964). *If all agents have dichotomous preferences, a weak Condorcet winner is guaranteed to exist.*

In case we restrict the preferences to *single-peaked* or *dichotomous*, the existence of a weak Condorcet winner is guaranteed.

Chapter 6

Protocols for Strategic Agents: Mechanism Design

Mechanism design is a strategic version of social choice theory, which adds the assumption that agents try to maximize their individual payoff, which means that various agents *declare* their preferences, which they may do truthfully or not. It is sometimes called "inverse game theory".

To get a feeling for what mechanism design constitutes, consider the following example.

agent 1:	$b \succ a \succ c$
agent 2:	$b \succ a \succ c$
agent 3:	$a \succ c \succ b$
agent 4:	$c \succ a \succ b$

Given plurality voting, candidate b is selected. If agent 4 knows the preferences of the other agents, he can manipulate this outcome in his advantage by lying about his true preference and instead announcing the following preference: $a \succ b \succ c$. Assuming an alphabetical tie-breaking rule, he has enforced his more preferred candidate a to be selected.

6.1 Social choice functions and strategic game forms

We will now see how we can implement a social choice function as described in the previous chapter (for example, a voting method) into a strategic game form. This will allow us to apply the solution concepts such as a Nash equilibrium (Chapter 3) and in this way derive properties of the social choice function. In general, the desired outcome is for the social choice function to be *truthfully implementable*, which means that players have no incentive to lie. For this, we create a strategic game in which the strategies of the players are their preferences. The true preferences of the agents are assumed to be unknown to the designer of the game. Therefore we need to consider every possible combination of preferences, and design a social choice function that will give us the desired outcome in every situation.

Consider the following example: two women both claim a new born child is theirs. Living in the middle-ages, there is no possibility to examine DNA material, which requires

King Solomon to take more drastic measure:

He sent for a sword, and when it was brought, he said: "Cut the living child in two and give each woman half of it". The real mother, her heart full of love for her son, said to the king, "Please, Your Majesty, don't kill the child! Give it to her!" But the other woman said, "Don't give it to either of us; go on and cut it in two". Then Solomon said, "Don't kill the child! Give it to the first woman, she is the real mother." (1 Kings 3: 16-28)

We will model this situation as a social dilemma. First, realize that there are three outcomes:

- a:** The first woman gets the baby.
- b:** The second woman get the baby.
- c:** The baby is bisected.

For each woman, we can model two possible types (the good mother and the bad mother).

$$\begin{array}{ll} \succ_1: a \succ b \succ c & \succ_2: b \succ a \succ c \\ \succ'_1: a \succ c \succ b & \succ'_2: b \succ c \succ a \end{array}$$

As we witnessed in the story, the actual preference profile is (\succ_1, \succ'_2) , since the first mother prefers to be separated from her child over the child itself being separated, while the second mother prefers it the other way around.

If both women tell the truth, this results in outcome a (since King Solomon decided that the first woman should have the baby). If the first woman lies and the second woman tells the truth, or if the second woman lies and the first tells the truth, the baby is bisected (outcome c). Finally, if both women lie, the second woman gets the baby. The strategic game in normal-form is depicted in Figure 6.1.

	\succ_2	\succ'_2
\succ_1	c	a
\succ'_1	b	c

Figure 6.1: *Solomon's Verdict game.*

Comparing this strategic game with the initial preferences of the women, we can conclude that telling the truth is not a dominant strategy for the second woman. Consider the case where the first woman tells the truth. If the second woman were to lie, the resulting preference profile is (\succ_1, \succ_2) , resulting in outcome c to be selected. Since woman two with preference profile \succ'_2 prefers c over a , it is beneficial for her to lie. If the first woman lies, the second woman prefers to lie as well, since she prefers b over c .

This means that the social choice function that Solomon came up with is not *dominant strategy incentive compatible*; players can enforce a better outcome when lying.

6.2 Incentive compatibility of social choice functions

Definition 41 (Dominant strategy incentive compatible). *A social choice function f is dominant strategy incentive compatible with relation to a set A of type profiles, if $\forall \succeq^* \in A, \forall \succeq, \succeq' \in A, \forall i \in N : f(\succeq_1, \dots, \succeq_i^*, \dots, \succeq_n) \geq f(\succeq_1, \dots, \succeq'_i, \dots, \succeq_n)$.*

The intuition is: a social choice function f is truthfully implementable if for no player there are situations in which telling the truth can hurt. This is an important property, because otherwise the mechanism does not guarantee that the goals of the designer, which are defined relative to the preferences of the individuals, are reached.

A *direct mechanism* is one in which the only action available to each agent is to declare its preferences. It turns out that for implementation of a social choice function in dominant strategy equilibria, we can restrict ourselves to direct mechanisms. This is called the *revelation principle*.

Theorem 9 (Revelation principle). *If there exists any mechanism that implements a social choice function C in dominant strategies then there exists a direct mechanism that implements C in dominant strategies and is truthful.*

We now ask what social choice functions can be implemented in dominant strategies. For this we first need to define what it means for a social choice function to be *dictatorial*:

Definition 42 (Dictatorial social choice function). *A social choice function f is dictatorial if there is a player i such that $\forall \succeq, o \in O : f(\succeq) \succeq_i o$.*

Intuition: for every preference profile \succeq , the outcome of the social choice function, i.e. $f(\succeq)$, is among the dictator's most preferred outcomes.

The following theorem is somewhat disappointing:

Theorem 10 (Gibbard, 1973 and Satterthwaite, 1975). *If $|O| \geq 3$, every incentive compatible social choice function onto O is dictatorial.*

Intuition: for all non-dictatorial social choice functions there are type profiles in which it is profitable for some player to lie about his true preferences. In other words, if a social choice function is incentive compatible, then it is dictatorial.

6.3 Implementation in dominant strategies

Now that we have seen what it means for a social choice function to be truthfully (dominant strategy) implementable, we are going to see how we can actually establish this result. We can generalize the notion of truthful implementable to S implementable for any other solution concept S . This results in the next definition.

Definition 43 (S -implementable). *A social choice correspondence $\phi : A \mapsto 2^O$ is S -implementable if there is some game form $G = (N, A, O, g)$ such that for all preference profiles \succeq : $\phi(\succeq) = \{g(s) \in A : s \in S(G, \succeq)\}$*

Intuition: The game form G implements ϕ in S if for each \succeq the sets of outcomes selected by S and ϕ coincide. The only type of implementability that we will discuss is implementation in *dominant strategies*, this means that the social choice function always selects a dominant strategy equilibrium. This equilibrium is very robust because

it assumes very little about the agents. On the other hand, it does not allow for much flexibility.

We will now consider an example of a voting game with two alternatives. Two players have to vote, and both have the ability to vote for (action a) or to vote against (action b). Thus, the players can choose from the following preferences:

$$\begin{array}{l} \succ^a: \quad a \succ b \\ \succ^b: \quad b \succ a \end{array}$$

We want to "design" the social choice function in such a way that the outcome of the social choice function is the dominant strategy in the strategic game. Assume that we come up with the following social choice function:

$$f(\succ_1, \succ_2) = \begin{cases} a & \text{if } \succ_1 = \succ^a \text{ and } \succ_2 = \succ^a \\ b & \text{otherwise} \end{cases}$$

We are now going to see whether this social choice function is implementable in dominant strategies. Therefore we have to consider every possible combination of strategies that the players can have, and verify that the strategy profile is a dominant strategy equilibrium. First we create the strategic game that is the result of this social choice

	\succ^a	\succ^b
\succ^a	a	b
\succ^b	b	b

Figure 6.2: *Social choice function implemented in a strategic game*

function, shown in Figure 6.2. Now we are going to consider every possible combination of preferences. This leads to four different preference profiles (resulting in the four different outcomes), which all turn out to be a dominant strategy equilibrium. We will briefly discuss why; in the following enumerating the first element of the pair represents the preference/strategy of the row player and the second element that one of the column player.

- (\succ^a, \succ^a) : the outcome is a . Consider the case where Row player lies, this results in outcome b which he prefers less than the actual outcome. A similar results holds for the column player.
- (\succ^a, \succ^b) : the outcome is b . For Row player, lying results in b as well, which makes him indifferent. The column player would actually be worse of by lying since it would result in outcome a .
- (\succ^b, \succ^a) : the outcome is b . due to symmetry, exactly the same results hold here of the column player and the row player as in the previous preference profile.
- (\succ^b, \succ^b) : the outcome is b . Lying would never change the outcome, so this strategy profile is also not dominated.

Chapter 7

Protocols for Multiagent Resource Allocation: Auctions

Auctions provide a solution to the problem of allocating resources among agents. They are used mainly because they provide a framework for understanding resource allocation problems. There are several types of auctions for both one and several goods.

Definition 44 (Auction). *An auction is any protocol that allows agents to indicate their interest in one or more resources and that uses these indications of interest to determine both an allocation of resources and a set of payments by the agents.*

7.1 Single-good auctions

Auctions are simply mechanisms (see Chapter 10) for allocating goods. A *single-sided auction* is the most simple auction with one good for sale (one seller) and multiple buyers. Each buyer has a valuation for this good and wishes to buy for the lowest price. The most familiar types of auctions are:

- **English auctions:** the auctioneer sets a starting price and buyers have the possibility to propose higher prices (usually with a minimum increment). The rules when the auction closes differ, but it is mostly related to some time limit. The highest bidder must buy for the price he proposed.
- **Japanese auctions:** the auctioneer again sets a starting price, but now the auctioneer keeps announcing increased prices where the buyers should choose whether they are "in" or "out". Once an agent is "out" he cannot become "in" anymore. The auction ends when only one buyer is "in", and this buyer should buy for the current price.
- **Dutch auction:** the auctioneer announces a high price and keeps decreasing it (like a clock) until one of the buyers signals the auctioneer that he is willing to buy for the current price. The auction stops and the buyer buys for the current price.
- **Sealed-bid auctions:** each agent submits a secret, "sealed" bid. The agent with the highest bid must purchase the good, but the price he has to pay depends on the

kind of auction: in *first-price auction* he pay the price he offered, in *second-price auction* (or *Vickrey auction*) he pays the amount equal to the next highest bid.

There are many more different kind of auctions, where at least the following rules should be defined:

1. *Bidding rules*: who can bid, when, what is the content?
2. *Clearing rules* when does the auction finish, who receives the goods for what price?
3. *Information rules*: are the bids private, what else is public/private?

7.1.1 Second-price auctions

We will now see whether the above described auction types are truthful or *incentive compatible* (see Chapter 10). First we will state a theorem addressing *second-price auctions*: auctions in which the winning buyer has to pay the price of the second-winning buyer. It turns out that this kind of auction is incentive compatible. There is one assumptions though: bidders have *independent private values* (IPV). This means that the valuations are drawn independently from the same distribution where the agent has no information about the valuations of the others.

Theorem 11. *In a second-price auction where bidders have independent private values, truth telling is a dominant strategy*

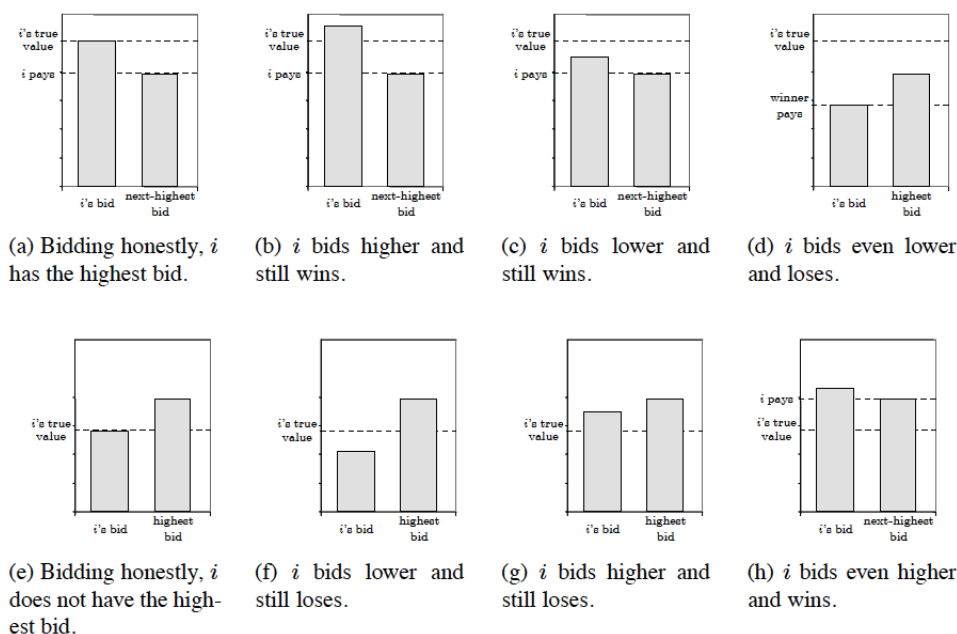


Figure 7.1: Schematic overview of an agent

Proof. We consider a player i that bids for a good and see whether he would have gained by deviating from its true value. If he never gains by deviating, truth-telling is dominant strategy. We can distinguish two possibilities in which the auctions terminates, namely that i wins or does not win. We are going to consider both cases and

see what happens when when i lies about his preferences. See Figure 7.1: in (a) player i tells the truth and has the highest bid. He has to pay the price of the next-highest bid, which is lower. In (b), (c) and (d) its shown what happens when i deviates. As can be seen, in all possible cases the price that i has to pay does not decrease. The same argument goes for (e), where i tells the truth and does not win. In this case (f), (g) and (h) show what happens if he deviates; they give similar results.

It can be proven that both Japanese auction and English auction are second-price auctions; they are incentive compatible too. The Japanese action is closely related to the second-price auction if we describe it as follows: each player writes down a number at which he will drop out. Now the winning agent will pay the second-highest price, because the winning agent pays the price when the second-winning agent drops out. There are small differences, but as long as the IPV assumption is satisfied, Japanese auctions - as well as English auctions - are dominant-strategy truthful.

7.1.2 First-price auctions

The Dutch auction is clearly a first-price auction. Each agent must select an amount without knowing about the other agents' selections; the agent with the highest amount wins the auction and must purchase for that amount. For the second-price auction it was clear that truth-telling is dominant. For first-price auctions this is not so easy, but it can be proven that in the case of two agents, there is an equilibrium in which each player bids half of his true valuation.

Proposition 1. *In a first-price auction with two bidders whose valuations are drawn independently and uniformly at random from the interval $[0, 1]$, $(\frac{1}{2}v_1, \frac{1}{2}v_2)$ is a Bayes-Nash equilibrium strategy profile.*

Although this proposition is quite narrow, it is already sufficient to show that first-price auctions are not incentive compatible. This means that in a first-price auction participants are better off by not telling the truth, i.e., truth telling is not rewarding. Somewhat more generally, we have the following theorem.

Theorem 12. *In a first-price sealed-bid auction with n agents the unique equilibrium is given by the strategy profile $(\frac{n-1}{n}v_1, \dots, \frac{n-1}{n}v_n)$.*

In other words, the unique equilibrium of the auction occurs when each player bids $\frac{n-1}{n}$ of his valuation.

7.1.3 Revenue equivalence

So what kind of auction should the auctioneer choose in order to enforce a fair outcome? The following theorem shows that this actually does not matter so much as one might think:

Theorem 13 (Revenue equivalence theorem). *Under the IPV assumption for a single good at auction, and the assumption that bidders are risk-neutral, any efficient auction mechanism results in any bidder with valuation v_i making the same expected payment.*

Thus all of the auctions we described so far must yield the same expected revenue to the seller, as long as the assumptions are satisfied. This theorem contains several terms that we are not introduced with yet, so we will briefly discuss them.

A *risk-neutral* attitude can be seen as a risk attitude in which a bidder is indifferent between a sure payoff and a gamble for a payoff, if both result in the same outcome. For example, if a risk-neutral bidder is offered a choice between 50 euro or a 50% change on 100 euro, he will be indifferent between these choices. A person with a *risk seeking* attitude prefers the gamble and a person with a *risk averse* attitude prefers the sure payoff.

An *efficient* mechanism selects the choice that maximizes the sum of agent's utilities, without considering the monetary payments that agents are required to make. So efficiency is defined in terms of agents' true valuations, not their declared valuations.

7.2 Multiunit auctions

In a *multiunit auction* there is still only one *kind* of good available, but there are now multiple identical (indivisible) copies of that good. We will review the same example auctions as in 11.1, but now extended to the multiunit case.

- **Sealed-bid auctions:** there are three main differences:
 - The payment rules are different, since there are now multiple winners who have to pay different prices. There are two common solutions for the payment rules:
 - * Pay-your-bid scheme (the *discriminatory pricing rule*): each of the winners pay its own bid. This can be seen as a generalization of the first-price auction.
 - * Uniform-pricing rule: all winners pay the same amount, which can be the highest losing bid or the lowest winning bid.
 - Bidders have to offer a bid for every number of units (meaning that the price may differ between buying one, two or three items). The bidder can place an *all-or-nothing* bid, meaning he will only accept the number of units that he bids for, or a *divisible* bid, where he will accept any smaller number at the same price-per-unit.
 - Tie-breaking (choosing between bidders with equal bids) can be difficult when considering all-or-nothing bids. Consider for example an auction for 10 units with the following all-or-nothing bids: 5 units for €20/unit, 3 units for €15/unit, 5 units for €15/unit, and 1 unit for €15/unit. It would seem that the first bid should be satisfied, but which of the other three? There are different tie-breaking rules one could think of: by quantity, by time or a combination of these.
- **English auctions:** faces the same problems there were addressed above, but bidders don't have to specify more than one number of units per price offer. Still it may be difficult to decide what bid to accept next, when bids are all-or-nothing. Consider the following example. An auction with 10 units for sale with two bids: one for 5 units for €1/unit, and one for 5 units at €4/unit. What is a next acceptable bid? A bid of 3 units at €2/unit is, but a bid for 7 units at €2/unit is not. This problem is solved if the bids are divisible.
- **Japanese auctions:** similar to English auctions; an agent calls a price and the number of units he is willing to buy. Usually the number decreases as the price

increases. Different implementations offer different solutions when demand exceeds supply: goods go unsold, etc.

- **Dutch auctions:** buyers call out the amount of units they want to buy for the current price. If there are units left, the auction continues. This may be from the current price, or it can be reset to a set percentage above the current price, or even to the original high price.

7.2.1 Single-unit demand

In the case of single-unit auctions we were able to establish useful results on the outcome of the mechanism using first-price and second-price auctions. Will this also be possible for the case of a multiunit auction? We will first look at the most basic case in which buyers only demand one unit (single-unit demand). The auction mechanism here is to sell the units to the k highest bidders for the price of the highest losing bid. Thus, instead of a second-price auction we have a $k + 1^{\text{st}}$ price auction.

There is something remarkable about this auction: selling more items will not necessarily result in a higher revenue for the seller. For example, consider the valuations in Figure 7.2. If the seller would only offer a single unit, he would receive €20 (the highest losing bid), while if he would sell three times he would receive $€8 \times 3 = €24$, and if he would sell four units he would receive no money at all! Why is this?

Bidder	Bid amount
1	€25
2	€20
3	€15
4	€8

Figure 7.2: Example valuations in a single-unit demand multiunit auction

The answer is logical: a dominant-strategy efficient mechanism can only select losing bids as the price per unit, so this bid will decrease as the number of bidders increases, eventually to zero when demand equals supply. How to solve this? The seller could set a minimum price, or select the lowest winning bid as the price to pay per unit. The payoff is that the mechanism will no longer be truthful. Still, in the single-unit demand case it does not matter much what auction the seller chooses, stated in the following theorem, which is an extension of Theorem 6.

Theorem 14 (Revenue equivalence theorem, multiunit version). *Under the IPV assumption for a single unit of k identical goods at auction, and the assumption that bidders are risk-neutral, any efficient auction mechanism results in any bidder with valuation v_i making the same expected payment.*

Note that these results only hold as long as players are risk-neutral and the IPV condition is satisfied. The fact that auction houses like eBay actually use non-dominant-strategy mechanisms (since they don't want items to be given away for free, the buyers often pay the price of the lowest winning bid), suggest that these assumptions are not very realistic in practice.

7.3 Combinatorial auctions

We will now consider an even broader auction setting, in which multiple different goods are for sale, which means that goods are no longer interchangeable. Examples of such auctions are: energy auctions, auctions for paths (routing, shipping rights, bandwidth) in a network.

So now we not only have a set of bidders $N = \{1, \dots, n\}$, but also a set of goods $G = \{1, \dots, m\}$. Let $v = (v_1, \dots, v_n)$ denote the true *valuation functions* of the different bidders, where for each $i \in N$, $v_i : 2^G \mapsto \mathbb{R}$, which means that the valuation of a bidder only depends on the set of goods he wins (he is indifferent about the allocations and payments of the other agents).

We are interested in *nonadditive valuation functions*, which means that the value of a bundle of goods is unequal to the sum of the values for single goods. We define two important kinds of nonadditivity: substitutability and complementarity.

Definition 45 (Substitutability). *Bidder i 's valuation v_i exhibits substitutability if there exist two sets of goods $G_1, G_2 \subseteq G$, such that $G_1 \cap G_2 = \emptyset$ and $v(G_1 \cup G_2) < v(G_1) + v(G_2)$. When this condition holds, we say that the valuation function v_i is subadditive.*

We can distinguish two ways in which two items can substitute each other:

- *partial substitutes*: $v(G_1 \cup G_2) < v(G_1) + v(G_2)$
- *strict substitutes*: $v(G_1 \cup G_2) = v(G_1) = v(G_2)$.

Examples for partial substitutes are two TVs of different brands or a CD player and an MP3 player. Two tickets to the same concert can be considered strict substitutes. Notice that strict substitution is simply a stronger version of a partial substitution.

Definition 46 (Complementarity). *Bidder i 's valuation v_i exhibits complementarity if there exist two sets of goods $G_1, G_2 \subseteq G$, such that $G_1 \cap G_2 = \emptyset$ and $v(G_1 \cup G_2) > v(G_1) + v(G_2)$. When this condition holds, we say that the valuation function v_i is superadditive.*

An example of complementary goods are a left shoe and a right shoe.

So how should the seller sell its good? There are several alternatives.

- Sell goods individually. This might not be fair because of the exposure problem. The *exposure problem* means that bidders offer a high price for a set of goods, but succeed only in winning a subset of these goods and therefore pay too much. This problem is especially common when the valuations for goods are strongly complementary, because then bidders might be willing to pay more for bundles and they would pay if the goods were sold separately.
- Run different auctions for the different goods, but connect them in certain ways. For example, one could run several auctions in parallel, synchronizing on rounds so that the buyers have a good idea of what is going on in the other auctions. This approach can be made more effective through the establishment of constraints on bidding that span all the auctions. For example, bidders may only increase their aggregate (total) bid between rounds with a certain amount. Therefore they will have an incentive to participate in early rounds as to bid the desired amount in a later round.

- A *combinatorial auction*: let the auctioneer sell all good in a single auction where bidders have the possibility to bid on a bundle of goods. This eliminates the exposure problem (bids satisfy the all-or-nothing criterium).

7.3.1 Simple combinatorial auction mechanism

The simplest combinational auction is the one in which the auctioneer computes the allocation that maximizes the social welfare.

Definition 47 (Social welfare). Social welfare means an allocation of goods that maximizes the sum of the valuations that the players have for it: $\chi = \max_{x \in X} \sum_{i \in N} v_i(x)$.

This mechanism charges the winners their bids, so it can be seen as a generalization of the first-price sealed-bid auction and therefore is not incentive compatible. Consider the example in Figure 7.3.

Bidder 1	Bidder 2	Bidder 3
$v_1(x, y) = 100$	$v_2(x) = 75$	$v_3(y) = 40$
$v_1(x) = v_1(y) = 0$	$v_2(x, y) = v_2(y) = 0$	$v_3(x, y) = v_3(x) = 0$

Figure 7.3: Example valuations in a combinatorial auction setting.

The goods to be distributed are x and y (one of each). If we apply the previously describe auction to this example, the outcome that maximizes the social welfare is to allocate item x to buyer 2 and item y to buyer 3. This results in a social welfare of $75 + 40 = 115$, which is higher than the other option, namely to allocate both goods to buyer 1. From this example it is clear that the mechanism is not incentive compatible; bidder 3 would for example be better off by bidding 38 for y .

Another option is to apply a generalized version of the second-price auction, which is a bit more tricky in this case. Consider the above example where bidder 2 and 3 receive the goods. To see how much bidder 2 has to pay, we need to see how much the highest losing bid pays. First note that without bidder 2, bidder 1 would have bought both items with a valuation of 100. But since bidder 3 buys item y for a value of 40 with bidders 2, we need to subtract 40 from 100, resulting in a price of 60. Via similar reasoning we conclude that bidder 3 has to pay a price of $100 - 75 = 25$. It can be verified that this mechanism is indeed truthful; bidders cannot gain by deviating from their real value.

Still this mechanism, and the second-price mechanism in general, has shortcomings. There are two main reasons why these actions are rarely seen in practice:

- **shill bid**: the seller can observe the bid of the truthtelling bidder and submit a fake bid just below, so that the winning bidder will have to pay more than the actual first losing bid.
- Both the seller and other buyers find out the true valuation and exploit this in a future transaction.

Other shortcomings are:

- vulnerability to collusion among bidders.
- one bidder masquerades as several different ones.
- computational.

The last point is arguably the biggest potential hurdle. In combinatorial auctions, determining the winners is a very challenging computational problem.

Definition 48 (Winner determination problem (WDP)). *The winner determination problem (WDP) for a combinatorial auction, given the agents' declared valuations \hat{v} , is to find the social-welfare-maximizing allocation of goods to agents.*

This problem turns out to be NP-complete, which means that it is not likely that a polynomial-time algorithm exists for the problem.

7.3.2 Expressing a bid: bidding languages

As we saw in our previous simple example, bidders had to specify a valuation for every subset of the goods at auction. Since this number is exponential in the number of goods, this is infeasible. Therefore we must find a way for bidders to express their bids in a more succinct manner. We will describe different languages that express different classes of bids. These languages should at least have the following properties: the language should be...

- ...*expressive* enough to represent all possible valuation functions;
- ...*concise*, when possible, values that are equal should not be stored double;
- ...*natural* for humans to both understand and create;
- ...*tractable* so the mechanism can determine the allocation and payments using only polynomial-time computation.

Moreover, we assume that the valuation functions have the following properties:

- *Free disposal*: goods have nonnegative value, so if $S \subseteq T$ then $v_i(S) \leq v_i(T)$, meaning that a buyer can never value a subset of a set of goods more than the set itself.
- *Nothing-for-nothing*: $v_i(\emptyset) = 0$; a bidder who gets no goods also gets no utility.

We will now discuss four languages for expressing bids that fulfill these requirements:

- **Atomic bid**: the most basic bid request. A pair (S, p) stating that an agent is willing to pay p for the subset of goods S . Notice there is an implicit AND operator between the different goods in the bundle. So an atomic bid of $(\{TV, DVD \text{ player}\}, 100)$ means a bid on the TV *and* the DVD player for €100.
- **OR bid**: atomic bids cannot express disjunctive bids such as €10 for a DVD Player or €20 for an MP3 Player. It is only possible to use the OR bid if the goods in the bid exhibit no substitutability. Recall that this means that the valuation for the bundle of goods equals the sum of the valuations for the separate goods.

Theorem 15. *OR bids can express all valuation function that exhibit no substitutability, and only these.*

So the atomic bid and the OR bid both offer us the possibility to express expressions such as $(x \vee y) \wedge z$. Still we cannot express bids such as: "I want the TV and the DVD Player for €100, or the TV and the MP3 Player for €75, but not both". This is because we are dealing with a XOR bid.

- **XOR bid:** is an exclusive OR of atomic bids $(S_1, p_1) \oplus (S_2, p_2) \oplus \dots \oplus (S_k, p_k)$ that indicates that the agent is willing to accept exactly one of the atomic bids.

XOR bids have unlimited representational power, since it is possible to construct a bid for an arbitrary valuation using a XOR for every possible bid.

Theorem 16. *XOR bids can represent all possible valuation functions.*

It may already be clear to the reader that although XOR bids can represent everything, this does not necessarily have to be efficient.

Consider the following example: we want to do a bid where we state that we want to buy a TV for €250 or a table for €50 or a vase for €10. If we assume that these goods are additive (which is reasonable), then this can simply be expressed using the OR bid as follows: $(\{TV\}, 250) \vee (\{table\}, 50) \vee (\{vase\}, 10)$. If we need to use the XOR bid, we will explicitly have to express all possible outcomes: $(\{TV\}, 250) \oplus (\{TV, table\}, 300) \oplus (\{TV, vase\}, 260) \oplus (\{TV, table, vase\}, 310) \oplus (\{table\}, 50) \oplus (\{table, vase\}, 60) \oplus (\{vase\}, 10)$. This example supports the next theorem.

Theorem 17. *Additive valuations can be represented by OR bids in linear space, but require exponential space if represented by XOR bids.*

- **OR-of-XOR bid:** a set of XOR bids, such that the bidder is willing to obtain any number of these bids. This does not expressivity (we already could express everything using only XOR bids), but it may add compactness, which is interesting from a computational perspective. It is also possible to create XOR-of-OR bids, or allow arbitrary nesting.