Collective Decision Making

Sylvain Bouveret, Jérôme Lang, and Michel Lemaître

Abstract This chapter introduces two prominent models of collective (multiagent) decision making, namely the basic ordinal model, and the utilitarian (numerical/quantitative) model. These models are then illustrated on three major collective decision making problems: voting, fair division and auctions. For each of these three problems we give a formal definition and we discuss the main links with computer science and artificial intelligence.

1 Introduction

1.1 Collective Decision Making Problems

This chapter focuses on collective decision making (CDM) problems, in which a group of people (agents) has to make a collective decision cooperatively. The chosen decision, to be selected among a set of eligible decisions, will engage each agent. Most procedures presented in this chapter are centralized procedures.

Typical CDM problems examples are: political elections; private everyday votes (for example, friends choosing a restaurant); fair allocation (for example, dividing goods in a divorce, allocating courses to students in a university...); a jury seeking for a consensus in a court.

Sylvain Bouveret
Univ. Grenoble-Alpes, CNRS, LIG, F-38000 Grenoble, France,
e-mail: `sylvain.bouveret@imag.fr`

Jérôme Lang
LAMSADE-CNRS, PSL Research University, Université Paris Dauphine, France,
e-mail: `lang@lamsade.dauphine.fr`

Michel Lemaître
Formerly ONERA, Toulouse, France,
e-mail: `michel.lemaitre.31@gmail.com`

The study of CDM problems dates back to Antiquity. The name “social choice” refers nowadays to the formal study of such collective problems. Nicolas de Condorcet (1743–1794) was one of the first to formalize some CDM problems. His contribution to the field of voting systems (de Condorcet, 1735) is widely recognized. Other prominent contributors are Kenneth Arrow (born in 1921), celebrated for his famous impossibility theorem (see page 4), and Amartya Sen (born in 1933), well known for his work on social inequalities (Sen, 1970).

Classical social choice theory has never been very much concerned about algorithmic issues. This is where computer science and more specifically artificial intelligence and operations research come into play: a recent research field has emerged, named *computational social choice*, bringing computer science and social choice together. Two research directions have appeared: the first one (from social choice theory to computer science) aims at exploiting social choice theory concepts and procedures in order to solve problems arising in computer science applications (for example, aggregation procedures for web page ranking and information retrieval, using voting for pattern recognition and classification, or computational resource allocation). The other direction (from computer science to social choice theory) aims at using notions and methods coming from computer science (representation languages, complexity, algorithmics, interaction protocols...) in order to solve complex group decision making problems. This last direction is by far the most important.

Formally, a CDM problem consists of three elements: a set of *agents* $\mathcal{N} = \{1,...,n\}$; a set of *eligible decisions* or *alternatives* $\mathcal{X}$; an expression of individual preferences (or sometimes *beliefs* — we will go back to this later) of each agent over the alternatives. The expected result is, as the case may be, the choice of a “socially optimal” alternative, the choice of a set of alternatives, or a ranking of the alternatives.

Three of the most important sub-fields of social choice are:

- *voting*: agents (or *voters*) express their preferences over alternatives (in this case *candidates*) and must agree on the choice of a candidate (or a subset of candidates).
- *fair allocation*: a common resource has to be divided amongst agents expressing their preferences about the possible shares they can possibly receive.
- *judgement aggregation*: agents express their beliefs over the real word and must come up to a common conclusion.

The first two examples above concern *preference* aggregation (the most frequent case in social choice), whereas in the last case, *belief* aggregation is at stake. Belief aggregation is addressed specifically in Chapter 14 of this volume and will not be discussed in this one. In the following we are only concerned with preference aggregation, focusing successively on voting, fair allocation, and finally on combinatorial auctions, which are a special case of resource allocation.

There are many models dedicated to the problem of aggregating individual preferences into a collective one. We will now present the two prominent models on which most works on CDM are based.

1.2 The Basic Model: Ordinal Preferences

Let $\mathscr{P}$ be the set of total preorders[1] on $\mathscr{X}$. In the ordinal model, the preferences of an agent i are represented by the *individual preorder* $\succeq_i \in \mathscr{P}$.

Let $G : \mathscr{P}^n \to \mathscr{P}$ be a *collective preorder aggregation function*. The *collective preorder* $\succeq_{col} = G(\langle \succeq_1, \succeq_2, \ldots, \succeq_n \rangle)$ (or *social welfare ordering*) represents the collective preference which results from the aggregation by G of the individual preference profile $\langle \succeq_1, \succeq_2, \ldots, \succeq_n \rangle$. A collectively preferred alternative is an alternative maximizing the collective preorder $\succeq_{col}$.

Let us give a simple example of a preorder aggregation procedure: let $N(a)$ be the number of times for which an alternative a is (one of) the most preferred in the individual preorders. Now define $a \succeq_{col} b \equiv N(a) \geq N(b)$. This is obviously a preorder. On the other hand, the aggregation which prefers alternative a to alternative b when a majority of agents prefer a to b is not a preorder, because it may generate a cyclic collective preference, in which case the preference relation is not transitive — this is the celebrated Condorcet paradox, see Section 2.

In this context, the centralized CDM problem consists in defining an aggregation function G having "good" properties. What are these "good" properties for CDM? We now introduce the main ones.

1.2.1 The Pareto-Efficiency Property and the Unanimity Principle

Informally, an efficient alternative is an alternative which satisfies all agents "as well as possible". The simplest and mostly used expression of efficiency is the *Pareto-efficiency* property, based on *Pareto-dominance*. Let $\langle \succeq_1, \succeq_2, \ldots, \succeq_n \rangle$ be a preference profile. We say that alternative a Pareto-dominates alternative b when $a \succeq_i b$ for all agents, with $a \succ_i b$ for at least one agent ($\succ_i$ designates the strict part of $\succeq_i$, that is $a \succ_i b \equiv [a \succeq_i b$ and not $b \succeq_i a]$). A Pareto-efficient (or Pareto-optimal) alternative is a non-dominated one. It is such that we cannot switch to another alternative increasing strictly the satisfaction of an agent, without strictly decreasing the satisfaction of another agent. We say that an aggregation function G satisfies the Pareto-efficiency property if the alternatives collectively preferred are Pareto-efficient.

The *unanimity principle* simply requires that the aggregation function G satisfies the Pareto-efficiency property.

1.2.2 The Independence of Irrelevant Alternatives (IIA) Property

This natural property asks that for each pair of alternatives a and b, the strict collective preference between a and b ($a \succ_{col} b$ or $b \succ_{col} a$) only depends on the way

[1] A preorder $\succeq$ is a binary relation that is reflexive and transitive. In a total preorder (or weak order), no pair of alternatives is incomparable: $x \succeq y$ or $y \succeq x, \forall x, y \in \mathscr{X}$.

each agent strictly compares a and b ($a \succ_i b$ or $b \succ_i a$) — the other alternatives are irrelevant.

1.2.3 Arrow's Theorem

Most results in classical social choice theory consist of *impossibility* or *possibility theorems* of the following form: *there is no collective decision procedure satisfying a set of natural and desirable conditions* $R_1, \dots, R_p$, or *the set of collective decision procedures satisfying the set of natural and desirable conditions* $R_1, \dots, R_p$ *is exactly the set of procedures of form* F. A celebrated example is Arrow's theorem (Arrow, 1951). Consider strict preference profiles $\langle \succ_1, \dots, \succ_i, \dots, \succ_n \rangle$ on $\mathscr{X}$ (total strict orders). Let $\mathscr{S}$ be the set of all possible strict profiles. Arrow's theorem states that if there are at least 3 alternatives, then any aggregation function G defined on $\mathscr{S}^n$ satisfying the unanimity principle and the IIA property is *dictatorial*, meaning that there is an agent i such that for any profile P, $G(P) = \succ_i$.[2]

1.3 The Utilitarian Model, or the Model of Quantitative Preferences

The utilitarian model (or numerical or quantitative preference model) represents the preferences of agent i by an *individual utility function* $u_i : \mathscr{X} \to \mathbb{R}$. To each alternative a corresponds a vector $\langle u_1(a), u_2(a), \dots, u_n(a) \rangle$ (utilities of a for each agent) called the *utility profile* of a.

In order to compare the (quantitative) satisfaction of two agents for a given alternative, utilities must be defined on a common scale. But this is not always possible: *interpersonal comparison of utility* is a critical question in CDM. Actually, agents may use their own non commensurable utility scales. However in the following we will assume (unless explicitly stated) that agents' utilities are expressed on a common utility scale.[3]

Let $g : \mathbb{R}^n \to \mathbb{R}$ be a *collective utility aggregation function*, and let $u : \mathscr{X} \to \mathbb{R}$ be the function defined this way: $u(a) = g(\langle u_1(a), u_2(a), \dots, u_n(a) \rangle)$, for any alternative a. The function u, called *collective utility function* or *social utility function*, represents the collective preference obtained through aggregation by g of individual utility functions u_i. A collectively preferred alternative is an alternative maximizing this function u.

Given individual and collective utility functions, we can easily recover the ordinal model by defining individual and collective preorders on $\mathscr{X}$ as follows: for any

---

[2] Arrow's theorem also holds (in a weaker form) when the individual preferences are preorders (that is, with possible indifference between alternatives).

[3] A straightforward way to obtain a common utility scale, often used in fair allocation problems, is to normalize the individual utility of each agent relatively to the utility she would get if she was given all the resource (Kalai-Smorodinsky normalization).

agent i, $a \succeq_i b \equiv u_i(a) \geq u_i(b)$, and $a \succeq_{col} b \equiv u(a) \geq u(b)$. The individual utility function u_i (respectively the collective utility function u) is said to *represent* the individual preorder $\succeq_i$ (respectively the collective preorder $\succeq_{col}$). In this way, each purely ordinal property (such as the Pareto-efficiency property) can be expressed in the utilitarian model.

The main two collective utility aggregation functions are sum $u(a) = \sum_{i \in \mathcal{N}} u_i(a)$, and minimum: $u(a) = \min_{i \in \mathcal{N}} u_i(a)$. These two functions respectively correspond to the main two agendas of the utilitarian model, namely *classical utilitarianism* and *egalitarianism*. Classical utilitarianism (sum) seeks to produce collective utility, irrespective of the agent from which this utility comes from (hence ignoring any equity concern). On the other hand, egalitarianism (min) seeks to maximize and equalize at the same time individual utilities: it selects an alternative which maximizes the satisfaction of the least satisfied agent, hence conveying a very strong equity flavor.

Classical utilitarianism and egalitarianism are two opposite and extreme attitudes towards CDM.[4] In classical utilitarianism, an agent is a “collective utility producer”. The marginal collective utility produced by an agent does not depend on her present degree of individual utility. Hence, the collective preference maximization could indeed lead to lower the satisfaction of the least satisfied agents, if more satisfied agents “produce” more utility. Agents must show a high degree of solidarity: some of them could be sacrificed on the altar of the collective utility maximization. Conversely, in egalitarianism, even a large utility increment of an agent already satisfied does not compensate for a tiny loss of utility of the least satisfied agent.

These two variations of utilitarianism are linked with two different approaches in philosophy and economics: Rawls (1971) and Sen (1970) for egalitarianism, and Harsanyi (1955) for classical utilitarianism are often advocated.

The utilitarian model often refers to an efficiency definition which is a refinement of Pareto-efficiency, namely *sum-efficiency*. A sum-efficient alternative is an alternative maximizing the sum of individual utilities, that is, alternatives maximizing the sum of individual utilities. A sum-efficient alternative is Pareto-efficient, but the converse is not true in general.

In the utilitarian model, maximizing the collective utility function yields Pareto-efficient (or Pareto-optimal) decisions if and only if the aggregation function g is strictly increasing. This is the case for sum, but not for min. The leximin total preorder is a refinement of the total preorder induced on $\mathcal{X}$ by min, of which the maximization always results in a Pareto-efficient alternative.

---

[4] For example, egalitarianism prefers the utility profile $\langle 10,10,10 \rangle$ to the profile $\langle 9,100,100 \rangle$. Classical utilitarianism prefers $\langle 1,100,100 \rangle$ to $\langle 66,67,67 \rangle$, and even to $\langle 2,99,99 \rangle$. See Section 3 for two families of trade-offs between utilitarianism and egalitarianism.

1.4 Centralized vs. Distributed CDM

Another essential dichotomy exists — orthogonal to the ordinal versus quantitative/numerical preferences one — namely the way agents interact in the decision making. In a *centralized* CDM resolution, a central authority (arbitrator, chairperson, Home Office, auctioneer, ...) gathers in a first phase the preferences of the agents (or at least a part of them, being informative enough for the decision to be made), and then decides on the optimal alternative and communicates it to the agents. The interaction phase between the central authority and the agents intended for collecting preferences is generally called the *elicitation* phase (the interested reader can see Chapter 7 of this volume. for more details). In fully *distributed* CDM resolution, there is no central authority, agents interact freely, and negotiate to reach a common consensus. There are also intermediate CDM frameworks, that lie between centralized resolution and distributed negotiation.

This chapter concerns mostly centralized CDM resolution, because of its importance, and because distributed decision (in particular negotiation) is considered in detail in Chapter 20 of this volume.

1.5 Discussion

In practice, the nature of the CDM problem at stake dictates the choice of a particular model of preferences (ordinal or numerical/quantitative) and type of resolution (centralized or distributed). By way of illustration, voting theory generally assumes ordinal preferences and centralized resolution; fair allocation with money (as with auctions), numerical preferences and centralized resolution; some fair allocation problems — such as *cake-cutting*, see Section 3.5 — ordinal preferences and distributed resolution.

Some difficulties may impair the use of a centralized CDM model:

- agents might be unable to reveal their preferences or simply refuse to do so, hence complicating the elicitation phase;
- in the real world, the agents often have intricate preferences that are difficult to translate into preorders or utilities: agents are often sensitive to several criteria (see Chapter 7 of this volume), and are often not indifferent to other agents preferences, as well as to some social norms.

However, to be accepted by the agents at stake, the CDM model and resolution should be based on clear concepts, easy to explain and use.

Finally, we should be aware of the limitations of the standard CDM models presented above. Their interest can be mostly found for CDM problems having a technical aspect — typically, the routine allocation of numerous physical resources — for which direct negotiations are hardly possible, and for which a kind of automatic processing is required. These models can serve as well to build some technically relevant solutions, initiating a negotiation process.

The following sections are devoted to three specific CDM problems. The first concerns voting, making use of the ordinal model of preferences. The next one is dedicated to fair allocation problems, for which equity is a strong concern. We present in the last section the auction problem, a specific allocation problem in which agents interact in a limited way, and which is solved by the mean of maximizing the profit of a particular agent (the auctioneer).

2 Voting

The use of voting procedures for collective decision making is not only of tremendous importance in large-scale political contexts, but it is also more and more applied in low-stake contexts such as social networks, workplaces or other local communities (and perhaps also societies of autonomous agents), which explains why it has acquired so much importance in the last fifteen years in the Artificial Intelligence literature. Since there are now hundreds of papers on voting in the mainstream AI conferences (such as AAAI or IJCAI) and journals (such as AIJ or JAIR) as well as in more specialised conferences (such as AAMAS), we cannot report on every research stream and we will only give a brief overview of the main topics (measured in number of papers). For a finer overview we advise to consult Chapters 2 to 10 of the Handbook of Computational Social Choice (Brandt et al, 2016b).

2.1 Introduction to Voting Theory

A common assumption in voting theory is that the agents (which will be called 'voters' in this section) have ordinal preferences, and furthermore that these preferences are *linear orders* (or *rankings*) over the set of alternatives. (There are exceptions to this, such as in approval voting, which will be discussed further.)

Let $N = \{1,...,n\}$ be a set of voters, and $\mathscr{X} = \{x_1,\ldots,x_m\}$ be a set of *alternatives*, or *candidates*. A *profile* is a collection of n votes, where each vote is a linear order over $\mathscr{X}$:

$$P = \langle V_1,\ldots,V_n\rangle = \langle \succ_1,\ldots,\succ_n\rangle$$

where V_i (also denoted $\succ_i$) is the vote expressed by voter i.

A *resolute voting rule* F is a function that maps each profile P to a candidate $F(P)$ in $\mathscr{X}$, who is the socially preferred candidate.

An *irresolute voting rule* F is a function that maps each profile P to a nonempty subset of $\mathscr{X}$: $F(P)$ is the set of socially preferred candidates, called *cowinners*; the candidate that will be chosen in the end will be one of the candidates of $F(P)$, obtained by means of a *tie-breaking mechanism* whose specification is outside the definition of F.[5]

When there are only two candidates a and b, the arguably most reasonable irresolute rule is the *majority rule*:

$$\text{maj}(V_1,\ldots,V_n) = \begin{cases} \{a\} & \text{if a strict majority of voters prefers } a \text{ to } b \\ \{b\} & \text{if a strict majority of voters prefers } b \text{ to } a \\ \{a,b\} & \text{otherwise} \end{cases}$$

May's theorem (1952) gives an axiomatic characterisation of the majority rule.

Things become more complicated when the number of candidates is at least 3. We now give an incomplete (but representative) list of voting rules. Unless stated otherwise, we define only their irresolute version; again, a resolute version can be obtained by composition with a tie-breaking mechanism.

A *positional scoring rule* is defined by a vector $\mathbf{s} = \langle s_1,\ldots,s_m\rangle$ of m integers, with $s_1 \geq \ldots \geq s_m$ and $s_1 > s_m$: each time voter i ranks candidate x in position j, x gets a score $score_i(x) = s_j$; the cowinners for the scoring rule $F_{\mathbf{s}}$ are the candidates maximizing $s(x) = \sum_{i=1}^n score_i(x)$. Here are three important examples of positional scoring rules:

- *plurality*: $s_1 = 1$, $s_2 = \ldots = s_m = 0$ (the cowinners are the candidates ranked first most often);
- *veto* (or *antiplurality*): $s_1 = s_2 = \ldots = s_{m-1} = 1$, $s_m = 0$ (the cowinners are the candidates ranked last least often);
- *Borda*: $s_1 = m-1$, $s_2 = m-2$, $\ldots s_m = 0$.

Consider the profile P composed of one vote $c \succ a \succ b \succ d$, two votes $a \succ b \succ d \succ c$ and two votes $d \succ b \succ c \succ a$: the cowinners for plurality are a and d; for Borda and veto, it is b.

Another important family of voting rules is that of the rules *based on the majority graph*. Given two candidates x and y, and a profile P, let $N_P(x,y)$ be the number of voters who prefer x to y in P. The *majority graph* M_P associated with P is the directed graph whose vertices are the candidates, and which contains an edge from x to y if and only if $N_P(x,y) > \frac{n}{2}$. A voting rule is *based on the majority graph* if the cowinners can be computed from M_P.

A candidate x is *Condorcet winner* for P if for any $y \neq x$, we have $N_P(x,y) > \frac{n}{2}$, that is, if it beats every other candidate in a pairwise duel by a majority of votes. Clearly, x is Condorcet winner for P if M_P contains an edge from x to every other candidate. Of course, when there exists a Condorcet winner, it is unique. However,

---

[5] The reason why we sometimes need irresolute rules is the possibility of a tie: suppose for instance that we have two candidates a and b, $n = 2q$ voters, and a profile P containing q votes $a \succ b$ and q votes $b \succ a$. For an irresolute voting rule, we simply let $F(P) = \{a,b\}$, and the final winner will be chosen by the tie-breaking mechanism. For a resolute rule, however, we have to specify the tie-breaking mechanism as part of the rule. For this, a choice must be made: either we give up *neutrality* and use a predefined priority relation over candidates; or we give up *anonymity*, and use predefined a priority relation over voters or sets of voters.

for some profiles, there is no Condorcet winner (see the example below). A voting rule is *Condorcet-consistent* if it elects the Condorcet winner when there exists one.

The majority graph M_P associated with the previous profile P is given in Figure 2.1. Each candidate being dominated by another candidate, there is no Condorcet

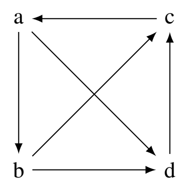

Fig. 1 Majority graph M_P

winner for P. If the first voter, instead of voting $c \succ a \succ b \succ d$, had voted $c \succ b \succ a \succ d$, then the edge $a \to b$ would have been replaced by an edge $b \to a$ and b would have been a Condorcet winner.

Here are three examples of rules based on the majority graph:

- *Copeland*: the *Copeland score* of a candidate x with respect to a profile P is the number of candidates that x beats in the majority graph M_P, plus half the number of candidates for which there is a pairwise tie with x (there is a pairwise tie between x and y if M_P contains neither an edge from x to y nor one from y to x). The Copeland (co)winners are the candidates with largest Copeland score. For example, for the profile P above, the Copeland winners are a and b.
- *Slater*: a *Slater order* for P is a linear order on $\mathscr{X}$ minimizing the number of edges disagreeing with M_P. A Slater winner is a candidate ranked first in some Slater order. For the profile P above, the unique Slater order is $a \succ b \succ d \succ c$, with only one disagreement with M_P (about (a,c)), and the Slater winner is a.
- *Banks*: for $S \subseteq \mathscr{X}$, let $M_P^{\downarrow S}$ be the restriction of M_P to S. $M_P^{\downarrow S}$ is a *maximal acyclic subtournament* of M_P if $M_P^{\downarrow S}$ is acyclic and for each S' such that $S \subset S' \subseteq \mathscr{X}$, $M_P^{\downarrow S'}$ is not acyclic. Then x is a Banks winner if x is non-dominated in a maximal acyclic subtournament of M_P. For the profile P above, the maximal acyclic subtournaments of M_P are obtained for $\{a,b,d\}$, $\{a,c\}$ and $\{b,c,d\}$, and the Banks winners are a, b and c.

Clearly, these three rules are Condorcet-consistent. For a survey of rules based on the majority graph, with a focus on computation, see Brandt et al (2016a).

The *weighted majority graph*, or *pairwise comparison matrix* W_P is defined by: for each $x, y \in \mathscr{X}$, $x \neq y$, $W_P(x,y)$ is the number of voters who prefer x to y minus the number of voters who prefer y to x. A voting rule is *based on the weighted majority graph* if the (co)winners can be computed from the weighted majority graph (or pairwise comparison matrix) W_P.

Consider the profile Q:

4 voters: $a \succ b \succ c \succ d$
2 voters: $b \succ c \succ d \succ a$
3 voters: $c \succ d \succ a \succ b$

The weighted majority graph for Q is

	a	b	c	d
a	–	5	−1	−1
b	−5	–	3	3
c	1	−3	–	9
d	1	−3	−9	–

Here is a rule based on the weighted majority graph: the *Simpson* (or *maximin*) rule outputs the candidates maximizing $\min_{y \neq x} W_P(x,y)$. The maximin winner for Q is a, with $\min_{y \neq a} W_Q(a,y) = -1$. Clearly, $\min_{y \neq x} W_P(x,y) > 0$ if and only if x is a Condorcet winner for P, and there cannot be two candidates with this property, therefore, the maximin rule is Condorcet-consistent.

Another rule based on the weighted majority graph is the *Kemeny* rule, defined as follows: the *Kemeny score* $K(V,P)$ of a linear order V with respect to profile P is defined by $K(V,P) = \sum_{(x,y) \in \mathscr{X}, x \neq y} W_P(x,y)$. A *Kemeny consensus* for P is a linear order V^* maximizing $K(V^*,P)$, and a *Kemeny winner* is a candidate ranked first in a Kemeny consensus. The Kemeny rule is Condorcet-consistent as well. The Kemeny rule can be used a voting rule, but perhaps even more so as a social welfare function (outputting the set of Kemeny consensuses). Also, it is easily adaptable to truncated votes, which explains why it is used for the aggregation of rankings of web pages given by different search engines (Dwork et al, 2001). On profile Q, the Kemeny consensus is $a \succ b \succ c \succ d$, with Kemeny score 18.

Some Condorcet-consistent rules are not based on the weighted (and a fortiori, unweighted) majority graph. Here is an example: the Dodgson rule[6] is defined as follows: for each $x \in \mathscr{X}$, $D(x)$ is the smallest number of elementary changes needed for making x a Condorcet winner, where an elementary change consists in swapping two adjacent candidates in a vote.

In order c to become a Condorcet winner for Q, it has to move one position up in two out of the first 6 votes; as for a, it needs to move two positions up in one of the last 5 votes; b and d need respectively 3 and 7 elementary changes in order to become Condorcet winners: therefore, a and c are the Dodgson cowinners for Q.

Here are now two rules that proceeds by *successive rounds*. First, *single transferable vote* (STV) proceeds in $n-1$ rounds, as follows:

1. let y be the candidate ranked first by the smallest number of voters (using a tie-breaking mechanism if necessary);[7]
2. eliminate y; the votes where y was ranked first are 'transferred' to the voter's preferred candidate among those who remain;

[6] Charles Dodgson was better known under the name of Lewis Carroll.

[7] Another way of handling ties consists in considering all tie-breaking possibilities and gather the corresponding winning candidates; the resulting rule is called the *parallel universe* version of STV (Conitzer et al, 2009).

3. iterate the process until there remains only one candidate.

Consider the profile R containing 3 votes $a \succ d \succ b \succ c$, 4 votes $b \succ d \succ a \succ c$, 3 votes $c \succ d \succ a \succ b$ and 2 votes $d \succ c \succ b \succ a$. At the first round, d is eliminated; the votes of the two voters who preferred d are transferred to their second choice, that is, c. At the second round, we have the reduced following profile: 3 votes $a \succ b \succ c$, 4 votes $b \succ a \succ c$, 3 votes $c \succ a \succ b$ and 2 votes $c \succ b \succ a$: a is eliminated. At the last round, only b and c remain; 7 voters out of 12 prefer b to c, and the winner is b.

When there are only three candidates, STV coincides with *plurality with runoff*, which is defined more generally as follows: the first round selects the two candidates with the largest plurality scores (again, using tie-breaking if necessary), and the winner of the second round is selected according to majority.[8]

Social choice theorists have studied some desirable properties of voting rules. *Condorcet-consistency* is one of them; note that no positional scoring rule is Condorcet-consistent (Moulin, 1988), and that STV and plurality with runoff are not Condorcet-consistent either. We give three other important properties, which for the sake of brevity we define for resolute rules only:

- *monotonicity*: when x is the winner for profile P, it remains the winner for a profile obtained from P by moving x up in some vote, the rest being unchanged;
- *participation*: when x is the winner for P, the winner for a profile obtained from P by adding one more vote is either x, or a candidate which the new voter prefers to x;
- *reinforcement*: when x is elected separately by two profiles, it is also elected by their union.
- *clone-proofness*: if a candidate x is cloned into a set of clones $\{x^1, \ldots, x^p\}$, and assuming that these clones of x will be ranked contiguously (in an arbitrary order) in each vote, and that the rest of the vote is equal to the vote before x was cloned, then the winner after cloning x will be (a) the same winner as before cloning x, if this winner was not x, and (b) one of the clones of x, if the winner was x.

For instance, positional scoring rules satisfy monotonicity, participation and reinforcement, but not clone-proofness; Copeland and maximin satisfy monotonicity, but not participation, reinforcement, nor clone-proofness; more generally, as soon as there are at least 4 candidates, Condorcet-consistency is incompatible with participation and with reinforcement. STV fails monotonicity, participation, reinforcement and Condorcet-consistency, but satisfies clone-proofness. Plurality with runoff fails to satisfy all these properties! For a survey on voting rules and their properties, see Brams and Fishburn (2004) and Zwicker (2016).

In *approval voting* (see Laslier and Sanver, 2010 for an extensive survey), the input is different: each voter specifies an *approval ballot*, which is subset of candidates she approves; the cowinners are the candidates approved by the largest number of voters. After adapting the properties we listed above to approval ballots, we obtain

[8] Plurality with runoff is used for political elections in many countries, such as France. STV – arguably better than plurality with runoff – is used for political elections in some countries such as Ireland and Australia.

that monotonicity, participation, reinforcement and clone-proofness are satisfied by approval voting, and that Condorcet-consistency is not.

2.2 Computing Voting Rules

Many voting rules are computable in polynomial time. This is the case for positional scoring rules and plurality with runoff, that are computable in time $O(nm)$, and for Copeland, Simpson, and STV[9], computable in time $O(nm^2)$.

But for some other rules, winner determination is hard. The first article that shows that a voting rule is computationally hard is (Bartholdi et al, 1989b), which shows that Dodgson and Kemeny rules are NP-hard. The exact complexity of the problem was determined by Hemaspaandra, Hemaspaandra and Rothe (1997): deciding whether x is a Dodgson winner is Θ_2^{P}-complete, that is, needs a logarithmic number of calls to NP-oracles.

The Kemeny, Slater and Banks rules are also hard to compute. Deciding if x is a Kemeny winner is Θ_2^{P}-complete (Rothe et al, 2003).[10] Deciding if x is a Banks winner is NP-complete (Woeginger, 2003); however, it is possible to find an arbitrary Banks winner in polynomial time by a greedy algorithm (Hudry, 2004b). Note that when comparing Banks to Kemeny, an important difference is that for Banks, since winner determination is “only” in NP, we can always find a succinct certificate for verifying that x is a winner (such a certificate is the subset S such that $M_P^{\downarrow S}$ is maximal acyclic), while for Kemeny, certificates are exponentially large (unless the polynomial hierarchy collapses). Finally, winner determination for the Slater rule is NP-hard, even under the restriction that ties between candidates do not occur (Ailon et al, 2005; Alon, 2006; Conitzer, 2006); but it is not known whether the problem is in NP or not.

These hardness results do not mean that we should give up using these rules, especially when they have good properties. Here are three ways of dealing with hardness.

First, *practical computation*: sometimes using a translation into a well-known setting with good solvers, such as integer linear programming; and sometimes using specific heuristics.

Second, *approximation*: interestingly, a polynomial approximation algorithm of a voting rule defines a new voting rule — sometimes already known under an other name, sometimes not. For example, let us consider the *Tideman rule*, defined as follows: if x, y are two candidates, let $\mathit{Deficit}(x,y) = \max(0, 1 + \lfloor \frac{N(y,x)-N(x,y)}{2} \rfloor)$ ($\mathit{Deficit}(x,y)$ is the number of votes which x needs in order to win against y, if that is possible) and the *Tideman score* is defined by $T(x) = \sum_{y \neq x} \mathit{Deficit}(x,y)$. The Tideman winner is the candidate minimizing the Tideman score. This rule is computable in $O(nm^2)$, and is a good approximation of the Dodgson rule, in the following sense:

---

9 For its version where ties are broken as soon as they appear.

10 For the more general problem of the computation of median orders, see Hudry (2004a).

under the assumption that the profiles are uniformly distributed (also called *impartial culture assumption*), the probability that a Tideman winner is a Dodgson winner converges asymptotically towards 1 when the number of voters tends to infinity (McCabe-Dansted et al, 2008). Also, sometimes it is possible to design an approximation of a voting rule that not only is easier to compute than the original rule, but also satisfies more desirable properties! For instance, while the Dodgson rule does not satisfy monotonicity, monotonic polynomial approximation of it have been designed by Caragiannis et al (2009).

Third, *fixed-parameter tractable algorithms*: sometimes a rule is hard but becomes polynomial-time computable when the number of candidates is fixed; this is the case for instance for the Kemeny rule, for which winner determination can be made by inspecting each of the $m!$ orders and computing their scores in polynomial time.

These following three handbook chapters review the complexity, the approximation, and the practical computation of these rules that are hard to compute: Brandt et al (2016a) for rules based on the majority graph, such as Banks or Slater; Fischer et al (2016) for rules based on the weighted majority graph, such as Kemeny; and Caragiannis et al (2016a) for other rules, such as Dodgson.

2.3 Voting on Combinatorial Domains

In many contexts, a decision has to be taken over several variables that may be intercorrelated. Two typical examples:

- *multiple referenda*: variables correspond to binary issues. For example, the inhabitants of a town may have to decide whether the town should build a swimming pool or not, and whether it should build a tennis court or not.
- *committee elections*: for example, a president, a vice-president and a secretary have to be elected, and some candidates (not necessarily the same ones) run for these positions. Sometimes there are no specific positions and the aim is just to elect k people.

In these situations, the space of candidates is a *combinatorial domain*: it consists in a Cartesian product $\mathscr{X} = D_1 \times \cdots \times D_m$, where D_i is a finite domain of values for variable X_i.

When the preferences of a voter on the values of a variable do not depend on the values of other variables, there are said to be *separable*. When all the voters have separable preferences, the vote can be decomposed into several independent voting processes, each bearing on a variable: for instance, there will be a vote about the swimming pool, and independently, a vote about the tennis court. Problems arise when some voters have nonseparable preferences. Consider the following example: there are two binary variables P (build a swimming pool), T (build a tennis court), and five voters whose preferences are

voters 1 and 2: $P\overline{T} \succ \overline{P}T \succ \overline{PT} \succ PT$
voters 3 and 4: $\overline{P}T \succ P\overline{T} \succ \overline{PT} \succ PT$
voter 5: $PT \succ P\overline{T} \succ \overline{P}T \succ \overline{PT}$

A first problem is concerned with the way the voters can express their preferences on $\{S,\overline{S}\}$ and on $\{T,\overline{T}\}$. This is not a problem for voter 5, whose preferences are separable. On the other hand, for voters 1 to 4, this is problematic. Take for example voter 2. If she votes for the swimming pool, she can favour, according to the votes of other voters, $S\overline{T}$ (her best candidate) or ST (her worst candidate); and if she votes against the swimming pool, she can favour one of her two intermediate candidates. In both cases, she can feel regret once the final outcome of the vote is known. A second problem is that the outcome of the vote can be extremely bad. If the voters majoritarily vote 'optimistically', the outcome will be ST, which is the worst alternative for all voters but one. Such paradoxes have been studied under the name *multiple election paradoxes* (Brams et al, 1998; Lacy and Niou, 2000).

When there is no guarantee that voters have separable preferences, the decomposition into independent voting processes is thus a bad idea, and other solutions must be found. There is no perfect solution; some possibilities:

1. ask voters to specify their preference relation *explicitly* on the set of all alternatives.
2. restrict the possible combinations of values for which one can vote.
3. ask voters to report a small part of their preference relation (*e.g.*, their top alternative), and apply a voting rule that needs only this information (*e.g.*, plurality).
4. ask voters to report their preferred alternative(s) and complete their preferences using a *distance* between alternatives.
5. use a *compact preference representation language* in which the voters' preferences will be represented succinctly.
6. *sequential voting*: vote about the variables one ofter the other, and communicate the outcome for a variable to the agents before they vote on the next variable.

One has to keep in mind that there are $\Pi_{1 \le i \le m} |D_i|$ alternatives. Therefore, as soon as there are more than three or four variables, Solution 1 is unrealistic.

Solution 2 is somewhat arbitrary: who decides which combinations are allowed? Moreover, in order this method to be realistic, the number of possible combinations has to be limited to a small number: voters thus express their preferences on a very small part of the set of alternatives.

Solution 3 is likely to give completely insignificant results as soon as the number of variables is significantly larger than the number of voters ($2^m \gg n$). For example, consider 5 voters and 6 binary variables, that is, 2^6 candidates, and choose plurality as the voting rule; one can expect the votes to be completely scattered, for example 001010: 1 vote; 010111: 1 vote; 011000: 1 vote; 101001: 1 vote; 111000: 1 vote; all other candidates: 0 vote. This solution is then completely pointless.

Solution 4 presupposes the existence of a natural and objective (voter-independent) distance between alternatives. It is used, among others, for defining the *minimax* committee election rule (Brams et al, 2007), and other rules, as well as in belief merging (see Chapter 14 of this volume). This solution is communicationwise cheap; it is however more costly in terms of computation, and it requires a significant domain restriction.

Solution 5 comes down to aggregate preferences specified in a compact representation language (see Chapter 7 of this volume), such as CP-nets. It is potentially highly costly in terms of computation.

Finally, Solution 6 is an interesting trade-off: it is relatively cheap in communication and computation, and it is applicable to nonseparable preferences. However, in order it to work well, the following domain restriction has to be made (Lang and Xia, 2009): there must exist a linear order on the variables $X_1 > \ldots > X_p$, common to all voters, such that the preferences of each voter on X_i are independent of the values of $X_{i+1}, \ldots, X_p$: for example, for the choice of a collective menu, `MainDish` $>$ `FirstCourse` $>$ `Wine` looks reasonable enough.

More details on voting over combinatorial domains can be found in (Lang and Xia, 2016).

2.4 Computational Barriers to Strategic Behaviour

A key problem in voting theory is that in some circumstances, some voters have an incentive to report insincere preferences in order to give more chances of winning to a candidate they prefer to the one who would be elected normally. Such a behaviour is called *manipulation*.

Consider for example plurality with runoff applied to the following profile: 8 votes $a \succ b \succ c$, 4 votes $c \succ b \succ a$ and 5 votes $b \succ a \succ c$. At the first round, c is eliminated, and at the second round, b is elected. Suppose now that 2 of the 8 first voters (those whose preference is $a \succ b \succ c$) decide to vote $c \succ b \succ a$ (all other votes being unchanged). The new profile is then composed of 2 votes $c \succ a \succ b$, 6 votes $a \succ b \succ c$, 4 votes $c \succ b \succ a$ and 5 votes $b \succ a \succ c$. At the first round, b is eliminated, and at the second round, a is elected. Since the actual preferences of these two voters are $a \succ b \succ c$, they are better off, since a is now the winner.

This example is not an isolated case. Indeed, the Gibbard-Satterthwaite theorem (Gibbard, 1973; Satterthwaite, 1975) shows that when there are at least three candidates, any voting rule which is nondictatorial and surjective (that is, for each candidate x, there is a profile for which x wins) is manipulable: for some profiles, some voters will have an incentive to report insincere preferences.

Although it is not possible to find a reasonable rule which is not manipulable, a way of limiting the impact of manipulation consists in making sure that a manipulation, whenever there is one, *is hard to compute*; this has lead computer scientists to study the *computational resistance to manipulation*. In practice, one considers that for a given voting rule, if finding a manipulation is NP-hard, then one can assume that voters – whose rationality is limited – will give up the idea of looking for one. Let us state the problem more formally by defining the following problem called COALITIONAL CONSTRUCTIVE MANIPULATION: for a voting rule F, given a distinguished candidate $x \in \mathscr{X}$, and the votes $\succ_1, \ldots, \succ_k$ of voters $1, \ldots, k$, is there a vote $\succ_i$ for each of the voters $i = k+1, \ldots, n$ such that x is elected by application of F on the profile $\langle \succ_1, \ldots, \succ_k, \succ_{k+1}, \ldots, \succ_n \rangle$?

The first articles on this topic have been written by Bartholdi et al (1989a) and Bartholdi and Orlin (1991). Then this question came back in the early 2000s, with Conitzer and Sandholm (2002a). Since then, more than thirty papers on the problem of complexity of (several variants of) manipulation have been written. They are surveyed by Conitzer and Walsh (2016). (See next section for another interpretation.)

Let us start by an example illustrating the *constructive manipulation of the Borda rule by a single voter*. Consider the following profile: $P = \langle a \succ b \succ d \succ c \succ e,\ b \succ a \succ e \succ d \succ c,\ c \succ e \succ a \succ b \succ d,\ d \succ c \succ b \succ a \succ e \rangle$. The current Borda scores (from these 4 votes) are a: 10, b: 10, c: 8, d: 7 and e: 5. Obviously, the last voter can make a or b win. Can she make c win? Yes, by ranking c first, then ranking in second position the least threatening candidate (e), then the least threatening after e (d), then a, then b (or vice versa). The final scores are then a: 11; b: 10; c: 12; d: 9; e: 8. Can she make d win? The same algorithm leads to rank d first, then e, then c, then, without loss of generaltiy, a, then b. The final scores are then a: 11; b: 10; c: 10; d: 11; e: 8: the existence of a constructive manipulation for d here depends on the tie-breaking priority order (there exists a constructive manipulation for d if and only if d has a higher priority than a or than b). On the other hand, there exists no constructive manipulation for e. The greedy algorithm we have applied (rank first the candidate that we want to be the winner, then the others by increasing order of their current Borda score, possibly taking tie-breaking priority into account) gives a successful manipulation if and only if there is one: the manipulation of the Borda rule by a single voter is therefore polynomial.

What about the same problem for two voters or more? Consider a profile for which the current Borda scores are a: 12; b: 10; c: 9; d: 9; e: 4; f: 1, with tie-breaking priority $a > b > c > d > e > f$. Generalizing the previous greedy algorithm does not work: suppose that the last two voters want e to win; after they rank it first, e has 14 points, and after they both rank f second, f has 9 points. They can continue with ranking d third for one of them and fifth for the other one (d then has 13 points). Then there are two ways of going further, depending on whether c is ranked once third and once fifth, or twice fourth; one can check that in the first case, is will not be possible to make e win, but that in the second case it will. This example suggests that computing a manipulation of the Borda rule by two voters or more is hard; its NP-hardness was long conjectured, and was proven independently by Betzler et al (2011) and Davies et al (2011).

Such complexity studies were done for numerous voting rules, in several contexts (constructive or destructive manipulation, by a single voter or a coalition of voters, by weighted or unweighted voters, with a retsriction to single-peakedness profiles oir not, etc.). We refer to Conitzer and Walsh (2016) for detailed results.

Some other works have also considered the issue of the *average* complexity of manipulation, starting from the constatation that an NP-hardness result talks about the worst case and does not guarantee that computing a manipulation will be *usually* hard. The results in this direction tend to show that there does not exist any rule that is *often enough* hard to manipulate (Procaccia and Rosenschein, 2007).

Beyond manipulation by coalitions of voters, there exists other types of strategic behaviour, such as "procedural control": some voting rules can be strategically controlled by the central authority ('chair') in charge of the election. The first article on this topic (Bartholdi et al, 1992) defines several types of control: by addition, suppression or partitioning of candidates or voters. For example, for control by addition of candidates, the chair can add a certain number of candidates in the hope of diluting the support to the candidates that can beat her favorite candidate. For each type of control and each voting rule F, three possibilities exist:

- F is *insensitive to control*: the chair can never act so as to change the winner.
- F is *resistent to control*: F is not insensitive to control but the control problem is computationally hard.
- F is *vulnerable to control*: F is not insensible to control and the control problem is computationally easy.

For example, the plurality rule is computationally resistant to control by addition or suppression of candidates, but computationally vulnerable to control by suppression of voters (Bartholdi et al, 1992).

Oher types of strategic behaviour, related to control, have been considered more recently, such as *bribery*, *control of sequential voting on a combinatorial domain* or *manipulation by cloning candidates*.

For a synthesis on computational barriers to strategic behaviour, see the work by Conitzer and Walsh (2016) for manipulation and by Faliszewski and Rothe (2016) for control and bribery.

2.5 Incomplete Knowledge and Communication

We consider here questions such as: given an *incomplete* description of the votes, is the outcome already determined? If not, what are the candidates who can still win and what are the relevant pieces of information to ask to the voters? How can we do that in order to minimize the amount of communication exchanged between the voters and the central authority?

For example, let us consider the following *partial profile*, with 4 candidates and 9 votes, out of which only 8 have been expressed:

4 voters: $c \succ d \succ a \succ b$
2 voters: $a \succ b \succ d \succ c$
2 voters: $b \succ a \succ c \succ d$
1 voter: $? \succ ? \succ ? \succ ?$

If the voting rule is plurality, then it is not difficult to see that the outcome is already determined independently of the last vote (the winner is c), while for Borda, the partial scores (computed from the 8 votes expressed) are a: 14; b: 10 ; c: 14; d: 10; thus, only a and c can win, and in order to determine the winner one needs only to know whether the last voter prefers a to c or vice versa. This problem is known under the name *vote elicitation* (Conitzer and Sandholm, 2002b; Walsh, 2008).

More generally, in order to model situations where the central authority has an *incomplete knowledge* of the voters' preferences, one considers that each voter has expressed a *partial order* over candidates, and a *partial profile* is a n-uple of partial orders $P = \langle P_1, \ldots, P_n \rangle$. A *completion* of P is a (complete) profile $T = \langle T_1, \ldots, T_n \rangle$, where each T_i is a linear order extending P_i. Then one defines the *possible and necessary winners* (Konczak and Lang, 2005) with respect to a voting rule and a partial profile:

- c is a *possible winner* if c is a winner for some completion of P.
- c is a *necessary winner* if c is a winner for each completion of P.

Thus, in the above example, c is a necessary winner for plurality; for Borda, the possible winners are a and c, and there is no necessary winner.

The computation of possible and necessary winners has received a significant attention, starting by Xia and Conitzer (2008). There also exists a probabilistic version, where one counts the extensions where a candidate wins (Bachrach et al, 2010).

Several classes of situations deserve a special attention:

1. *possible and necessary winners for the addition of voters*: some voters have expressed their votes entirely, whereas the others have not expressed anything: the partial profile is $P = \langle P_1, \ldots, P_{n-k} \rangle$, where P_i is a linear order on $\mathscr{X}$. This class of situations corresponds, with a different interpretation, to a coalitional manipulation problem: more precisely, let us consider the coalition A composed of the last k voters. Then x is a possible winner if the coalition A can make x win (or equivalently, A has a constructive manipulation for x), whereas x is a possible winner if A cannot prevent x from winning (or equivalently, A has no destructive manipulation against x).
2. *possible and necessary winners for the addition of candidates*: the voters have expressed their preferences on a fixed subset of candidates, and nothing on the other candidates: the partial profile is $P = \langle P_1, \ldots, P_n \rangle$, where P_i is a linear order on $\{x_1, \ldots, x_{m-k}\} \subseteq \mathscr{X}$. This class of situations occurs when new candidates declare in the curse of the process: one can for instance think of a Doodle poll for finding the date of a meeting, where voters have expressed their preferences on a first set of time slots, and when new time slots that were previously not considered can become possible after this first vote; or else, consider a hiring committee where a preliminary vote occurs between the candidates already interviewed and a new candidate is declared admissible (Chevaleyre et al, 2010). As an example, consider 12 voters, an initial set of candidates $X = \{a, b, c\}$ and a new candidate y. If the voting rule is plurality with tie-breaking priority order $a > b > c > y$, the partial profile is such that the plurality scores before y is taken into account are $a : 5$, $b : 4$, $c : 3$. On can check that a and b are possible winners, but not c. For instance, for b, it is enough that 2 of the voters who ranked a first now rank y first: the new plurality scores are $a : 3$, $b : 4$, $c : 3$, $y : 2$, and the winner is b.
3. *truncated ballots*: every voter has expressed a partial ranking consisting of her top k candidates.

4. *incomplete lists*: every voter i has expressed a ranking of an arbitrary subset $S_i \subseteq \mathscr{X}$ of alternatives (the candidates in $X \setminus X_i$ being incomparable with those in S_i, and incomparable with each other). This class of situations occurs when voters have an informed opinion on a subset of alternatives only: for instance, on a web application for evaluating restaurants, a voter can evaluate only those she has tried.

A problem closely related to the search of possible winners for the addition of candidates is that of manipulation by candidate cloning (Elkind et al, 2010); the difference is that for candidate cloning, although we don't know how the clones of a candidate will be ranked by a voter in the profile after cloning, still, we know that they will be ranked contiguously in each vote.

Another topic is the design of *communication protocols* for voting rules. The definition of a voting rule does not say anything on the way the winner will be determined by the central authority. On the other hand, a *communication protocol* for a voting rule specifies the pieces of information that each voter will communicate in each round, in such a way that at the end of the protocol, the result will be known. (More generally, a protocol can be seen as an algorithm where atomic instructions are replaced by *communication actions* between agents, in such a way that an agent, in a given round, communicates information based on *her knowledge*.) The cost of a protocol is the total number of bits exchanged in the worst case. The *(deterministic) communication complexity* of a voting rule F is the cost of the cheapest protocol for F: it measures the minimal amount of information to be communicated so that the result of the vote can be determined. The communication complexity of voting rules is studied in detail by Conitzer and Sandholm (2005).

A trivial protocol for an arbitrary voting rule F is the following: each voter i sends her vote V_i to the central autority (this requires $n \log m!$ bits), and then the central authority sends back the name of the winner to all the voters (this requires $n \log m$ bits). The communication complexity of a voting rule is thus in $O(n \log m!)$. However, for some voting rules there exists cheaper protocols. This is obvious for instance for plurality, since it suffices for each voter to send the name of their preferred candidate to the central authority: the communication complexity of plurality is therefore at most in the order of $n \log m$ (it is in fact *exactly* of this order; the proof of the lower bound is nontrivial (Conitzer and Sandholm, 2005)); but it is also the case for many other voting rules, such as plurality with runoff (in the order of $n \log m$), STV (in the order of $n(\log m)^2$), etc.

Another related problem is the *compilation of the votes of a subelectorate*. In a context where the votes do not come in a single round (consider for instance a political election where the votes of the citizens living abroad come with a few days delay, or to a Doodle pool where some persons are late in responding). In this case, it makes sense to compile the votes expressed so far, using as little space as possible, so as to "prepare the ground" for the time where the remaining votes are known. The *compilation complexity* of a voting rule is the minimal size for compiling a profile. It is identified, for some rules, by Chevaleyre et al (2009) and Xia and Conitzer (2010).

For more details about voting with incomplete preferences, communication protocols, vote compilation, as well as the learning of (some classes of) voting rules and the robustness of voting rules, see the chapter by Boutilier and Rosenschein (2016).

2.6 Some Other Issues

For the sake of brevity, there are a number of other research topics at the meeting point of voting and AI which we have not discussed in this section.

One which is especially relevant to this book is *group planning*. It is addressed for the fist time by Ephrati and Rosenschein (1993): each agent has her own goal; at each round, agents vote on the next action to perform without revealing their preferences explicitly. More generally, *collective combinatorial optimization* deals with the design of methods for the collective version of specific combinatorial optimization problems, such as shortest path finding (Klamler and Pferschy, 2007), minimum spanning tree (Darmann et al, 2009); egalitarian versions of some other combinatorial optimization problems are studied by Galand and Perny (2006); Escoffier et al (2013).

A few other topics, such as: *randomized voting*, *iterative voting*, *computer-assisted theorem proving in social choice*, *approximate notions of single-peakedness*, the *computational aspects of apportionment and districting*, *group classification*, *group recommendation*, *social choice and crowdsourcing*, and *dynamic social choice*, are briefly reviewed in the chapter by Brandt et al (2016c, Section 4), and the first four are reviewed in more detail by (Endriss, 2017). Lastly, new approaches to the *rationalization of voting rules*, partly originating in AI research, are reviewed in the chapter by Elkind and Slinko (2016).

3 Fair Allocation

Quand on partage le gâteau, l'important est : qui tient le couteau ?
(When the cake is divided, the main question is: who holds the knife?)
Bernard Maris, French economist, killed on January 7, 2015 in Paris

3.1 Fair Allocation Problems

Every CDM process is guided, explicitly or not, by the desirable properties that the collective decision must satisfy. We have seen in the introduction the most prominent of these properties: *efficiency*, which is often implemented by Pareto-efficiency. Another property which is very often required is *fairness*. Indeed, the essence of CDM is very often to look for admissible compromises between the agents antagonistic interests and preferences, which is a possible definition of fairness. We will later introduce several formal models of fairness.

The need for fairness is particularly strong in a kind of CDM problems called *fair allocation* or *fair division* problems, which are the subject of this section. Here, the alternatives are just *allocations* of goods (or resources) to agents. Even if the traditional CDM problem assumes that the agents have preferences over all the alternatives, it is commonly assumed that each agent only cares about what she receives (her own *share*). In other words, an agent will be indifferent between two allocations where she receives the same share.

Fair allocation problems can be divided into classes. The first one concerns *divisible* goods or resources, like money, time, water or land. For a long time this class has been explored by economists, using continuous mathematics in microeconomics. The second one is about *indivisible* objects or resources, like works of art, pieces of furniture, teaching time slots, cars or houses. Mixed problems exist, a classical example being fair division problems with indivisible goods, but monetary compensation (money is a special divisible good, others are indivisible).

The contribution of artificial intelligence to the field of fair allocation mainly concerns fair allocation of indivisible goods without monetary compensation, which are the most difficult from an algorithmic point of view, because of their strong combinatorial nature. Actually, consider the allocation of m objects to n agents (each object must be allocated to one and only one agent), the number of possible allocations is n^m: the size of the solution space increases exponentially with the size of the problem instance. For similar reasons, artificial intelligence is involved also in fair allocation of divisible and heterogeneous resources (*cake-cutting*), see Section 3.5.

3.2 Some Real World Fair Allocation Problems

Before going further and in order to emphasize their importance, we enumerate a set of real world fair allocation problems.

1. frequency allocation to radio stations, land division, mining and natural resource sharing (Antarctic, ocean floor, Moon), common exploitation of a scientific facility, such as Earth-observing satellites (Lemaître et al, 1999);
2. fair representation (apportionment) (Balinski and Young, 2001), setting up electoral boundaries;
3. fair allocation of kidneys or other organs to transplant;
4. allocation of positions in public entities;
5. sharing operating costs of international organisations, assessment of taxes;
6. allocation of permits to discharge pollutants;
7. sharing water treatment facilities between localities;
8. dividing estates in inheritance or divorce;
9. sharing time slots in schools, hospitals, airports,...;
10. allocating tasks or offices to employees, rooms to students, articles to reviewers, quotas of refugees to countries.

Notice that a lot of these problems concerns fair division of indivisible and non-shareable goods without monetary compensations.

3.3 How to Define Fairness?

Even if fairness is sometimes defined using the prominent Aristotelian's principle "equal treatment of equals, unequal treatment of unequals", there is no definitive and universal definition of fairness, but a number of properties corresponding to different formal definitions. None of these properties is universally considered to be the right notion of fairness, but each one conveys a different aspect of fairness. Some of these properties are defined on the collective preference (such as anonymity, separability, inequality reduction) while others apply to allocations (proportionality, envy-freeness).

Even if there is no universal notion of fairness, two properties are commonly required: *unanimity* and *anonymity*. Unanimity corresponds, in the context of fair division, to the Pareto-efficiency property, already discussed in Section 1.

3.3.1 Anonymity

If equal agents should be treated equally, then the *anonymity* property should be the first prerequisite. This properties conveys the fact that the collective preference should not depend on the agents' identities. Formally, for all permutation of agents σ, then the collective preorder aggregation function G, must satisfy

$$G(\succeq_1, \succeq_2, \dots \succeq_n) = G(\succeq_{\sigma(1)}, \succeq_{\sigma(2)}, \dots \succeq_{\sigma(n)}).$$

3.3.2 The Tension Between Unanimity and Strict Equality

In this paragraph we adopt the utilitarian model, with a common scale of utilities: for example, it is meaningful to say that a given allocation satisfies agent 1 more than agent 2.

In general, unanimity and strict equality cannot be satisfied at the same time. That is, there is in general no Pareto-efficient allocation giving to each agent the same amount of individual utility, as the following abstract situation involving two agents confronted to four possible allocations illustrates:

allocations	u_1	u_2
a	4	4
b	3	6
c	7	5
d	2	11

Allocation a is perfectly equitable, but is dominated by allocation c (c is better than a for each agent). Hence a is ruled out, in spite of its perfect equity. How to choose then between b, c and d? No one dominates another. Egalitarianism (maximizing the min) selects c whereas classical utilitarianism chooses the sum-efficient but obviously unfair d. From an equity point of view, we are inclined to select c as the "optimal" allocation. But there are less obvious situations. Imagine having to choose between two allocations associated to the following utility profiles: $\langle 1,49,50\rangle$ and $\langle 2,2,96\rangle$. Or what about the case $\langle 14,43,43\rangle$ and $\langle 15,15,70\rangle$?

3.3.3 The Priority Principle

Another principle is sometimes considered in situations where anonymity is not completely relevant: the *priority* principle. Following this principle, an allocation decision should be based on agents' characteristics. For example, in the kidney allocation problem, patients having waited longer than others could have priority, or those having a longer life expectancy after transplant. Birthright is also a typical example.

3.3.4 Independence of Unconcerned Agents (IUA)

This property, also called *separability*, applies to the collective preference. According to IUA, the collective preference should be such that an agent who is indifferent between two allocations can be ignored when choosing between these two allocations (she is not concerned by the choice). In other words, if this property is not satisfied, the collective preference between two allocations for which an agent is indifferent, depends on the precise individual utility of this agent for these allocations, which is hardly acceptable.

Consider the following example (Moulin, 1988). There are three agents and the collective utility aggregation function g is the median. Let a and b be two allocations with respective utility profiles $\langle 0,2,3\rangle$ and $\langle 0,1,4\rangle$. We have $g(\langle 0,2,3\rangle) > g(\langle 0,1,4\rangle)$ hence $a \succ_{col} b$. Now, consider two other allocations a' and b', with respective profiles $\langle 5,2,3\rangle$ and $\langle 5,1,4\rangle$ (utilities of agents 2 and 3 are not modified, but utility of agent 1 is raised from 0 to 5. We have now $g(\langle 5,2,3\rangle) < g(\langle 5,1,4\rangle)$, that is $b' \succ_{col} a'$. The preference is reversed. Agent 1 is not concerned by the choice between a and b, but her individual utility influences the choice between allocations which does not change utilities of others! The collective preorder represented by the median does not satisfies the IUA property.

This property is connected to the following important result. A collective preorder is continuous and satisfies the IUA property if and only if this preorder is represented by an additive collective utility aggregation function $g(\overrightarrow{u}) = \sum_i f(u_i)$ where f is continuous and increasing.

3.3.5 Inequality Reduction

This property supposes the utilitarian model with a common individual utility scale.

We have to define first what is called an *inequality reducing transfer*, or *Pigou-Dalton, transfer*. Let $\overrightarrow{u} = \langle u_1, u_2, \dots, u_n \rangle$ be a utility profile, with $u_1 < u_2$. Consider a utility transfer from agent 2 to agent 1 (from the richest to the poorest) such that after this transfer, $\overrightarrow{u}$ becomes $\overrightarrow{v}$ with $u_1 + u_2 = v_1 + v_2$ and $|v_2 - v_1| < |u_2 - u_1|$. Such transfer is said to reduce inequalities.

The inequality reduction property requires that any inequality reduction transfer does not strictly decrease the collective utility. Formally, the preorder $\succeq_{col}$ represented by the aggregation function g reduces inequalities when, for any pair of utility profiles $\overrightarrow{u}$ and $\overrightarrow{v}$ equal except on their first and second components, such that $u_1 < u_2$, $|v_2 - v_1| < |u_2 - u_1|$ and $u_1 + u_2 = v_1 + v_2$, we have $g(\overrightarrow{u}) \leq g(\overrightarrow{v})$.

Here is an example with three agents and $g(\langle u_1, u_2, u_3 \rangle) = u_1^2 + u_2^2 + u_3^2$. Let a and b be two allocations respectively associated to the utility profiles $\langle 0, 3, 4 \rangle$ and $\langle 1, 2, 4 \rangle$. From a to b, inequalities are reduced, however $g(\langle 0, 3, 4 \rangle) = 25 > g(\langle 1, 2, 4 \rangle) = 21$ which means that a is preferred to b. We conclude that the collective preference does not obey the inequality reduction property in this case.

An interesting fact is connected with this property and the separability (IUA) property: the preorder $\succeq_{col}$ represented by the additive aggregation function $g(\overrightarrow{u}) = \sum_i f(u_i)$ reduces inequalities if and only if f is a concave function. In the previous example, $f(x) = x^2$ is convex.

We now turn to properties which apply to allocations.

3.3.6 Proportionality (or Proportional Fair Share) and Max-min Fair Share

An allocation satisfies proportionality when each agent gets at least $1/n$ of the total utility she would have received if she receives alone all objects (n is the number of agents). This property was coined by Steinhaus in 1948 in the context of continuous fair division (cake-cutting) problems.

Proportionality has been adapted to indivisible goods without monetary compensation by Budish (2011), which defines the *max-min fair share* property. The original definition is purely ordinal. In utilitarian terms, the max-min fair share of an agent is the maximal utility that she can hope to get from an allocation if all the other agents have the same preferences as her, when she always receive the worst share (it is the best of the worst shares). An allocation satisfies the max-min fair share property if each agent receives a utility no less than the utility of her max-min fair share. Proportionality implies max-min fair share.

3.3.7 Envy-freeness

This very general and intuitive property does not require interpersonal comparisons of utility (just like proportionality). It both applies to the ordinal and cardinal (utilitarian) models.

An allocation is envy-free if no agent strictly prefers the share of another agent to her own share. It is a kind of stability property. Formally, let $a_{i/j}$ be an allocation identical to a except that agent i gets the share that was the share of agent j in a. We say that a is envy-free if $a \succeq_i a_{i/j}$ for all agents i, j.

Envy-freeness and Pareto-efficiency are generally not compatible. Furthermore, the problem of determining whether an allocation satisfying envy-freeness and Pareto-efficiency exists is quite complex; even in most reasonable settings (see for instance Bouveret and Lang, 2008).

Under additive preferences, envy-freeness implies proportionality. Other fairness properties and their relations to each other are presented in the paper by Bouveret and Lemaître (2016).

3.4 Main Aggregation Functions

In the utilitarian model, a family of aggregation functions is particularly interesting in the context of fair division, namely the *root-power quasi-arithmetic means* family, defined as follows (assuming strictly positive utilities):

$$g_p(\overrightarrow{u}) = \left(\frac{1}{n}\sum_i u_i^p\right)^{1/p},\ p \neq 0 \qquad g_0(\overrightarrow{u}) = \left(\prod_i u_i\right)^{1/n},\ \text{for } n = |\overrightarrow{u}|$$

The family is parameterized by the real number p. Functions of this family are additive[11] and hence induced preorders obey the separability (IUA) property (see page 23). When $p = 1$, g_1 is the standard arithmetic mean, corresponding to classical utilitarianism. The case $p = 0$ corresponds to the Nash function, which is independent of individual scale of agents utilities, a particularly interesting property. The collective preorder induced by g reduces inequalities if and only if $p < 1$. Finally, when p tends to $-\infty$, g tends toward the min function, and the induced preorder tends toward the leximin preorder. [12]

[11] Strictly speaking, these function are not additive. However the preorders they induce can be represented by additive functions, derived from original ones by increasing transformations. Even g_0 is additive in the broad sense of the term, because the additive function $\sum_i \log(u_i)$ represents the same preorder.

[12] The leximin preorder is a refinement of the preorder induced by the min function which satisfies unanimity (Pareto-efficiency). The leximin preorder is precisely the one which at the same time reduces inequalities and is independent of the common scale of utilities.

Notice that this family creates a continuum between the extremes classical utilitarianism (sum) and egalitarianism (min).

Another family of interest is the family of *ordered weighted averages* (OWA) (Yager, 1988), a variant of the weighted averages in which weights do not hold on the components but on the ranks. A n-OWA is a family of aggregation functions from $\mathbb{R}^n$ into $\mathbb{R}$, taking $\overrightarrow{w} = \langle w_1, \ldots, w_n \rangle \in [0,1]^n$ as parameter, with $\sum_i w_i = 1$, defined by $O_{\overrightarrow{w}}(\overrightarrow{a}) = \sum_{i=1}^{n} w_i \cdot a_i^\star$, where $\langle a_1^\star, a_2^\star, \ldots, a_n^\star \rangle$ is $\langle a_1, a_2, \ldots, a_n \rangle$ once sorted weakly increasing. OWAs can express: the mean ($w_i = 1/n$ for all i); the min ($w_1 = 1$, and $w_i = 0$ for all $i > 1$); the median ($w_{(n+1)/2} = 1$, and $w_i = 0$ for $i \neq (n+1)/2$); parameterizable compromises between min and mean, for example $w_i = \alpha^i$, $0 < \alpha < 1$ (with a suitable normalization); a function which tends towards a representation of the leximin preorder (the previous one when α tends towards 0).

3.5 *Procedural Allocation of a Divisible and Heterogeneous Resource (Cake-Cutting)*

The previous allocation model — choosing an allocation that maximizes an appropriate collective utility function — is based on two implicit assumptions. First, each agent should completely and honestly report her preferences (under the form of a utility function). Second, the agents rely on a central entity that is in charge of computing the allocation. In some cases, the agents do not wish to publicly report their preferences, and even if they accept to do so, nothing prevents them from acting strategically and misreporting them. Finally, the agents can simply refuse to trust a central authority. Hence the model based on the central optimization of a collective utility function is not adapted to every situation.

There exist a very different kind of allocation *procedures*, that have been studied for years. These procedures are by essence distributed and output a fair and efficient allocation from the preferences or a small fraction of them reported (honestly if possible) by the agents. These procedures — often called *mechanisms* — are particularly used in the context of the allocation of an infinitely divisible and heterogeneous good. The traditional metaphor is the *cake-cutting* situation: a rectangular cake (formally modeled has the $[0,1]$ real interval) is the common divisible and heterogeneous resource, and has to be shared among the n starving invitees, which all have a particular utility function on this interval.[13]

The allocation procedures that are studied in this context are similar to games in which the agents interact. The most prominent procedure in this context is the well-known "I cut, you choose" game, that can be used to cut a cake between two participants which can be roughly described as follows. One of the two participants is the divider. The other one is the chooser. Provided that the cake can be divided in all possible ways, that it is heterogeneous and that the participants can have different

[13] Or at least an ordinal function on intervals: between two intervals, each agent must be able to determine which of the two is better.

tastes for different parts of the cake, cutting the cake in two equal parts is in general not Pareto-optimal. The safest action for the divider is to cut the cake into two pieces that are equal *from her point of view*. Then the chooser will take the best of the two pieces. Under mild natural assumptions,[14] it is easy to see that the resulting allocation is Pareto-efficient, envy-free and proportional (satisfying the fair share property). However, as we shall see later, the generalization of this protocol to three agents or more is not straightforward.

The allocation problem of divisible and heterogeneous goods has a lot of real-world applications, among which we can mention the time-sharing of a common facility, or the land division problem.

A large number of works are dedicated to this problem, essentially in the field of economics. The seminal books by Robertson and Webb (1998), Young (1994, Chapters 8 and 9) and Brams and Taylor (1996) are good surveys of the topic. The interested reader can also have a look at the more recent paper by Brams et al (2006). Plenty of procedures are now well-known and well-studied, adapted to different contexts, and characterized by their fairness and efficiency properties. Some impossibility theorems have also been described.

More recently, researchers in artificial intelligence have contributed to the field of cake-cutting. They have mainly focused on the algorithmic complexity of the proposed procedures. The analysis of the complexity bounds requires the introduction of precise models of interaction between the agents. From this point of view, the researchers in AI also contribute: for instance, the classical model of interaction, that has been proposed by Robertson and Webb (1998), has been recently extended by a group of computer scientists (Brânzei et al, 2016). For interested readers, the paper by Procaccia (2009) proposes an interesting survey of the complexity bounds of cake-cutting procedures and the chapter by Procaccia (2016) gives an overview on algorithmic aspects of cake-cutting.

To give an idea of the difficulty of the mathematical problems we have to deal in this area and how computer scientists have contributed to the field, let us go back to the aforementioned problem of finding a protocol to find an envy-free cake cutting for three agents or more. In the early 60's, Selfridge and Conway independently came up with a protocol that returns an envy-free cake cutting in a bounded number of steps for three agents (Brams and Taylor, 1996). A few decades later, Brams and Taylor (1995) came up with a general envy-free protocol that works for any number of agents. This protocol is guaranteed to terminate in finite time; however, the number of queries needed can be unbounded, even for four agents. The problem of finding a protocol that returns an envy-free allocation in a bounded number of queries for any number of agents had been opened for decades until it was finally solved by two computer scientists, Aziz and Mackenzie (2016). They came up with an algorithm that finds an envy-free allocation in less than $n^{n^{n^{n^{n^{n}}}}}$ queries, hence closing what has been described by several researchers as one of the main challenges in the field.

[14] The agents are rational (they decide so as to maximize their satisfaction) and their utility is additive in the ordinal sense: if $A \succ B$, $C \succ D$ and $A \cap C = \emptyset$, then $A \cup C \succ B \cup D$.

To finish this overview, we can mention some extension of the cake-cutting problem (called online cake-cutting) in the case where some participants arrive and depart when the allocation process is ongoing (Walsh, 2010).

3.6 Fair Division and Computer Science

As we have seen earlier, resource allocation problems have long been mainly studied by economists, either from the normative or axiomatic point of view (as in the works by Young (1994) and Moulin (2003) for instance) or from the procedural point of view like in the works by Brams and Taylor (1996, 2000) about cake-cutting. However, like in voting, computer scientists and AI and OR researchers have started for a couple of years to investigate the computational aspects of resource allocation problems: compact preference representation, algorithmic or complexity issues... The vitality of the field is well illustrated by the chapters dedicated to fair division in the survey books on computational social choice (Brandt et al, 2016b, Chapters 11, 12 and 13) (Rothe and Rothe, 2015, Chapters 7 and 8).

3.6.1 Compact Preference Representation

Even if a lot of work has been done in recent years on the topic of compact preference representation (see Chapter 7 of this volume dedicated to this topic), only a small fraction of this work directly concerns resource allocation problems, with the notable exception of combinatorial auctions (as we will see in Section 4). On the one hand the domain of compact preference representation is quite young and resource allocation problems only represent a small fraction of the individual or collective decision making problems involving compact representation issues. On the other hand, a lot of works dedicated to algorithmic or complexity issues of fair division problem simply rule out these compact representation problems by assuming that the preferences are additive: see for instance the paper by Lipton et al (2004) mainly dedicated to additive preferences, or the works by Bezáková and Dani (2005), Bansal and Sviridenko (2006), Asadpour and Saberi (2007) on the Santa-Claus problem.

The first papers explicitly dedicated to compact preference representation in fair division problems date back to the works by Chevaleyre et al (2004) on k-additive functions and those by Bouveret et al (2005) and Bouveret and Lang (2008) mainly concerning logic-based compact representation. We can also mention an adaptation of the language of CP-nets for the compact preference representation in the context of fair division problems, that was proposed by Bouveret et al (2009). It seems however that not much work has been done since on compact representation in fair division, and that AI researchers have focused on simpler settings like additive preferences. One of the reasons might be that it is not so clear whether the benefit of using an expressive compact representation language is worth the additional complexity cost and elicitation burden or not, and that the use of simple additive preferences might well be enough for most applications.

3.6.2 Complexity and Algorithmic Issues

At the beginning of computational social choice, most works especially dedicated to algorithmic aspects of resource allocation are from the fields of combinatorial auctions or operations research. In the latter domain, fair division problems have been mostly considered from the point of view of "fair" multicriteria optimisation problems, that is, optimisation problems where the criteria should be maximized (or minimized) while being made as equal as possible. For instance, leximin, or (inequality-reducing) OWA optimization problems belong to this kind of problems. Among these works, to cite only two, Ogryczak (1997) applies fair optimization to fair facility location problems, and Luss (1999) to fair division.

Several works have followed the earlier paper about algorithmic aspects of fair division in different fields of research, both in artificial intelligence and operations research. For instance, a very active stream of works has concerned algorithmic and complexity aspects of a particular fair division problem, the *Santa-Claus Problem*. This problem can be formulated as follows: Santa-Claus has to allocate a set of m (indivisible, non shareable) toys to a set of n children. Each child has an additive utility function on the set of toys. The allocation must be made so as to maximize the utility of the least satisfied child. Of course, in spite of this special formulation, this problem is nothing else than a fair division problem with indivisible goods, additive preferences, and under the egalitarian criterion. Several papers (Bezáková and Dani, 2005; Bansal and Sviridenko, 2006; Asadpour and Saberi, 2007, just to name a few) have investigated the complexity and approximation algorithms of this problem. They have also drawn a interesting parallel with scheduling problems that have led to fruitful approximation approaches.

Still in the context of additive preferences, several works have also investigated other fairness criteria. The seminal paper by Lipton et al (2004) is one of the first works concerning the complexity and approximation of computing an allocation minimizing the envy between agents.[15] de Keijzer et al (2009) have extended the latter paper by proving that the problem of determining whether an envy-free and Pareto-efficient allocation was NP-complete. Later on, the case of ordinal separable (*i.e.* additive) preferences has been investigated by Bouveret et al (2010), and further by Aziz et al (2015) that have introduced interesting notions of envy-freeness and Pareto-efficiency based on stochastic dominance. Concerning utility maximization problems, Bouveret and Lemaître (2009) have focused on the computation of leximin-optimal solutions using constraint programming approaches. Golden and Perny (2010) and Lesca and Perny (2010) have also studied preference aggregation in particular in the context of fair division problems, focusing on fairness criteria like Lorenz optimality, maximization of an OWA or of a Choquet integral (exten-

[15] This work introduces an interesting extension of the aforementioned envy-freeness criterion, by proposing different *measures* of envy.

sion of the OWA that can take into account positive or negative interactions between agents).

Finally, several papers go beyond additive preferences and investigate theoretical complexity issues related to the use of compact preference representation languages in the context of fair division. For utility maximization problems, the paper by Dunne (2005) was one of the first works investigating the complexity of fair division problems with preference representation based on Straight-Line Programs. At the same time, Bouveret et al (2005) have focused on preference representation based on propositional logic. All these results have been extended to other social welfare functions, other languages like *k*-additive languages, and approximation issues by Nguyen et al (2014). Finally, the complexity of finding envy-free allocations with logic-based compact preference languages has been investigated by Bouveret and Lang (2008).

3.6.3 Distributed Allocation and Communication Complexity

Even if distributed allocation and negotiation issues in fair division will mainly be discussed in Chapter 20 of this volume, an overview of computer science aspects of resource allocation would be incomplete without evoking this domain. In the absence of any central authority, the natural way of computing an "optimal" allocation is to start from an initial allocation and then let the agents changing it using multilateral negotiation. In this framework, the main desirable properties are related to the convergence of the negotiation process, and the complexity is not defined in terms of computation, but in terms of communication costs (number of steps, size of the messages exchanged...).

The first theoretical results in this domain date back to Sandholm (1998). The notion of communication complexity has been imported in fair division in particular by Endriss and Maudet (2005) and Dunne et al (2005), that mainly focus on the number of swaps needed to reach an optimal allocation. We can also mention, among other works on the subject, the paper by Chevaleyre et al (2007) that focuses on a relaxation of the envy-freeness criterion, for which the agents only have a limited knowledge of the other agents. Finally, a paper by (Chevaleyre et al, 2017) analyzes the fairness properties (like envy-freeness or proportionality) of the allocations obtained after the convergence of the negotiation process with several kinds of deals, and also in the case where the set of possible swaps is constrained by a graph.

3.6.4 Recent Trends

We can observe an interesting recent trend in the community of computational fair division. As already noticed in the section dedicated to compact preference representation, a significant trend is to abandon complex theoretical frameworks and look for simple models that are inspired by practical applications and can be used in real situations. This has led to a stream of works trying to implement fair division in practice (see for instance the wonderful web application Spliddit[16]), or take inspiration from practice to find better ways of defining fairness or motivating fairness criteria. With simplicity in mind, most of these works are based on additive preferences.

Among the works proposing some alternative approaches to fairness, Bouveret and Lemaître (2016) introduce a scale of five fairness criteria of increasing strength. This scale can be used as a measure of the level of conflict of the agents' preferences: the higher criterion it is possible to satisfy, the less conflicting the preferences are, and the more likely it will be possible to find a satisfactory allocation. Among the five criteria, the maximin share and the Competitive Equilibrium from Equal Income (CEEI) criteria were already known in economics, but had been ignored so far by computer scientists. This is no longer the case, and it had led to fruitful works on theoretical properties, approximation and complexity, either about CEEI (Brânzei et al, 2015; Aziz, 2015), or about the maximin share (Procaccia and Wang, 2014; Amanatidis et al, 2015; Kurokawa et al, 2015). Caragiannis et al (2016b) have also revisited a long-standing criterion, the seminal Nash social welfare function, and shed a new light on its appealing fairness properties.

Finally, we can also mention an interesting work that perfectly characterizes the fruitful collaboration between economics and computer science. Dickerson et al (2014) analyze, both in practice and analytically, the probability of existence of an envy-free allocation, depending on the ratio between the number of objects and the number of agents: when this ratio is low, an envy-free allocation is very likely not to exist, and the opposite when the ratio is high enough. Furthermore, the simulations show a very interesting phase-transition phenomenon.

4 Combinatorial Auctions

4.1 From Classical to Combinatorial Auctions

Auctions is probably one of the most widely studied collective decision making problem in the economical literature in the last fifty years. In its most general definition, an auction is simply a structured mechanism in which some agents, the bidders, compete for some objects to buy. The mechanism (which is in practice implemented by a central entity, the auctioneer) is in charge of determining which objects each agent will get at which price. A wide variety of mechanisms are studied by economists and used in practice. Just to name a few, an auction can be *open* if the participants publicly announce their bids, or *sealed* if the bidders only reveal this information to the auctioneer and hide it from the other participants. An auction is *ascending* if the bidders iteratively increase their bids until no agent is willing to pay a higher price, and *descending* if the price proposed for an object decreases until at least one bidder declares to be interested. In a *first price* auction, the winner should

[16] `http://www.spliddit.org/`

pay the price corresponding to the highest bid, whereas in a *second price* auction, she should pay the second highest price.

The most common auctions types are the *English auction* (open first-price ascending) which is commonly used in artwork sales. The *Dutch auction* (open first-price descending, in which the auctioneer progressively decreases the price proposed for an object until a bidder accepts and pays this price) is traditionally used for perishable products like tulips in the Netherlands. The *Vickrey auction* (sealed second-price) is also called philatelic auction because it is used in the United States for collection stamps sales. Finally, the sealed first-price auction is classically used for the attribution of government contracts.

Auction theory has been studied for about 50 years mostly in economics — the first theoretical work on auctions is generally attributed to Vickrey (1961) — but computer scientists have recently paid an increasing attention to this field, mainly with the study of *combinatorial auctions*.

The study of combinatorial auctions in computer science dates back to the work of Rassenti et al (1982). The starting point of this work is that classical auctions mechanisms, sequential by nature (that is, selling objects one by one) can be quite inefficient and barely adapted to the situations where the bidders have non-modular preferences on the objects; in other words, when they have are preferential dependencies.

Let us consider a simple example where three objects are for sale: a vinyl disc player (p) and two (indivisible) sets of vinyl discs: the first set contains records from the Beatles (b) and the second one records from the Rolling Stones (s). Agent 1 is very interested in having one of the two sets of records (no matter which of the two), but does not own any disc player. Moreover, she is not interested in buying the disc player alone, because she does not have any vinyl disc to listen to. She is for instance ready to pay €100 for $\{p,b\}$ or $\{p,s\}$, €110 for $\{p,b,s\}$ but nothing for individual objects. In other words, p and b are complementary, p and s are as well, but s and b are substitutes. Agent 2 is also interested in the sets of records, but she already owns a disc player. Let us say that she is thus ready to pay €30 for $\{b\}$ or for $\{s\}$, €10 for $\{p\}$, €40 for $\{p,b\}$ or $\{p,s\}$ and €70 for $\{p,b,s\}$ (her preferences are additive).

If the objects are sold sequentially, the first agent will probably have difficulties to bid for them. Not knowing Agent 2's preferences, she will probably not take the risk of bidding for one of the two sets of records if she does not know for sure whether she will get the disc player (and the other way around). Agent 2 will not have the same difficulties: her preferences being additive, she can safely bid on each of the three objects separately. Not only sequential allocation can have a negative effect on bidding, but it can also harm the overall auction efficiency. For instance, in the latter auction, if Agent 1 is risk-adverse and chooses not to bid at all, the three objects will go to Agent 2 for a total price of €70 (provided that it is a first-price auction). If all three objects had been allocated to Agent 1, the auctioneer would have earned €110.

An obvious solution to this problem is to sell *bundles* of objects instead of selling them individually.[17] However, it it not obvious how to do so, because the way the bundles are formed should be related to the preferential dependencies of the agents. In some cases it is reasonable to assume that these dependencies are the same (in a shoe sale for instance, we can reasonably assume that the agents will only be interested in pairs of shoes, not individual shoes), but it is not always the case. For instance, in the latter auction, would it be more relevant to sell p and s together, or b and s together? The only solution to this problem is to sell all the objects simultaneously, and provide a way for the bidders to choose themselves the bundles they want to bid for. This is the basic idea behind combinatorial auctions. It has motivated the first works in this domain, concerning the allocation of take-off and landing slots in airports (Rassenti et al, 1982). In this application, the notion of preferential dependency is naturally present (what would an airline do with a take-off slot without the corresponding landing slot?).

It is not a surprise that this extension of classical auctions has been mainly developed and studied in the field of computer science and artificial intelligence. A lot of problems that arise in combinatorial auctions are well-known in computer science. As we shall see, the combinatorial blow-up induced by the representation of the allocation space calls for compact representation bidding languages. Furthermore, the problem of determining the optimal allocation is a lot more complex than in traditional auctions and induces intricate algorithmic issues. Finally, even if we will elude these aspects in the chapter, issues related to the design of truth-telling auctions mechanisms and their resistance to manipulation is a crucial topic. They are not only related to combinatorial auctions, but have a special formulation in this context. All these topics are covered in details in the reference book by Cramton et al (2006).

In what follows, we will denote by $\mathscr{O}$ the finite set of objects to be allocated to the agents (the set of objects the agents bid for). Given a set of n agents $\mathscr{N}$ and a set of objects $\mathscr{O}$, an *allocation* $\overrightarrow{\pi}$ is a vector $\langle \pi_1, \dots, \pi_n \rangle$, where for all i, $\pi_i \subset \mathscr{O}$ denotes the *share* received by agent i. In this section, we will only focus on allocations satisfying the preemption constraint, that is, such that $\forall i \neq j : \pi_i \cap \pi_j = \emptyset$ (an object cannot be allocated to two different agents).

4.2 Bidding Languages

As we have seen, the main difference between combinatorial and classical auctions is the bidding set, which is namely the set of objects $\mathscr{O}$ for classical auctions and the set of bundles $2^{\mathscr{O}}$ for combinatorial auctions. From the theoretical point of view, changing the bidding set does not make a huge difference in the formal definition of

[17] Other methods exist; see for instance simultaneous ascending auctions (Cramton, 2006).

the problem. However, in practice, the combinatorial dimension of the bidding set induces crucial representation[18] and computation issues.

The most prominent bidding languages used in combinatorial auctions are the ones from the family of XOR / OR / OR* languages (Nisan, 2006; Fujishima et al, 1999; Sandholm, 2002).

Definition 1 (XOR / OR / OR* languages). Let $\mathscr{O}$ be a finite set of objects. An *atomic bid* on $\mathscr{O}$ is a pair $\langle \mathscr{S}, w\rangle \in 2^{\mathscr{O}} \times \mathbb{R}^+$. A set $\{\langle \mathscr{S}_1, w_1\rangle, \ldots, \langle \mathscr{S}_p, w_p\rangle\}$ of atomic bids is said to be *admissible* if $\mathscr{S}_i \cap \mathscr{S}_j = \emptyset$ for all $i \neq j$ in $\{1, \ldots, p\}$.

A bid expressed in the *XOR language* is a finite set of atomic bids

$$\langle \mathscr{S}_1, w_1\rangle \ \mathsf{XOR} \ \ldots \ \mathsf{XOR} \ \langle \mathscr{S}_p, w_p\rangle.$$

The utility function associated to a bid $\mathscr{M}$ in the XOR language, mapping each set of objects π (in other words each possible share) to the price the agent is ready to pay for it, is defined as follows:

$$\begin{aligned} u : 2^{\mathscr{O}} &\to \mathbb{R}^+ \\ \pi &\mapsto \max_{\substack{\langle \mathscr{S}_i, w_i\rangle \in \mathscr{M} \\ \mathscr{S}_i \subseteq \pi}} w_i \end{aligned}$$

A bid expressed in the *OR language* is a finite set of atomic bids

$$\langle \mathscr{S}_1, w_1\rangle \ \mathsf{OR} \ \ldots \ \mathsf{OR} \ \langle \mathscr{S}_p, w_p\rangle.$$

The utility function associated to a bid $\mathscr{M}$ in the OR language is:

$$\begin{aligned} u : 2^{\mathscr{O}} &\to \mathbb{R}^+ \\ \pi &\mapsto \max_{\substack{\mathscr{M}' \subseteq \mathscr{M} \\ \mathscr{M}' \text{ admissible}}} \sum_{\substack{\langle \mathscr{S}_i, w_i\rangle \in \mathscr{M}' \\ \mathscr{S}_i \subseteq \pi}} w_i \end{aligned}$$

A bid expressed in the *OR*language* is a bid expressed in the OR language in which one or several dummy objects $d \notin \mathscr{O}$ can be present.

In the XOR language, an agent can cast a custom set of atomic bids. Each atomic bid gives the price the agent is ready to pay for the corresponding bundle. Given a set of objects the price an agent is ready to pay for it is the price of the best bundle it contains.

The OR language works the same way, except that in this language, the prices are interpreted additively.

By adding dummy objects to the bids in the OR* language the agents can make several bids incompatible, as in the XOR language, even if they do not overlap otherwise.

---

[18] A simplistic representation of an agent's utility function requires $2^m - 1$ values, corresponding to the number of non-empty subsets of $\mathscr{O}$.

Some authors (see for instance the work by Sandholm, 1999) have proposed to combine the OR and XOR language to benefit from the expressivity of the XOR language and the compactness of the OR language. Several languages have been developed and used, among which we can mention the OR-of-XOR, XOR-of-OR and OR / XOR languages.

Some other kinds of bidding language have also been developed, among others logical ones. For instance, Boutilier and Hoos (2001) have proposed a language mixing logic and numerical weights (representing utilities) associated to subformulas. The main interest of such a language is to combine the approach based on objects and the approach based on bundles, by proposing a way to logically combine the weighted formulas (that can be seen as the atomic "bids" of this language).

4.3 The Winner Determination Problem

4.3.1 Formulation and Theoretical Complexity

The Winner Determination Problem (WDP for short) is the central problem in combinatorial auctions. The objective is to decide, among the set of bids, which ones will be selected, coming down to determine which objects will be allocated to which agents. The main allocation criterion used in combinatorial auctions is the utilitarian criterion, that is, we look for the allocation that maximizes the revenue of the auctioneer.

Definition 2 (Winner Determination Problem).

- **Input :** A set of agents $\mathscr{N}$, a set of objects $\mathscr{O}$, and a set of utility functions $(u_1, \ldots, u_n)$ expressed as bids in a combinatorial auction language.
- **Output :** An allocation $\overrightarrow{\pi}$ of the objects that maximizes $\sum_{i=1}^{n} u_i(\pi_i)$.

Notice that this formulation of WDP implicitly assumes that the auctioneer can freely dispose objects (in other words, the allocation can be incomplete), which is a common assumption in combinatorial auctions.

The Winner Determination Problem has mainly been studied in the context of OR or XOR bids, for which there is a natural formulation in 0–1 linear programming. The idea is to create a variable $\mathbf{x_i^j} \in \{0,1\}$ for each atomic bid $\langle \mathscr{S}_j, w_j \rangle \in \mathscr{M}_i$. $\mathbf{x_i^j} = 1$ if and only if this atomic bid is selected in the allocation.

$$
\begin{aligned}
\mathbf{max}\ & \textstyle\sum_{i \in \mathscr{N}} \sum_{\mathscr{S}_j \in \mathscr{M}_i} w_j \times \mathbf{x_i^j} \\
\mathbf{s.t.}\ & \mathbf{x_i^j} \in \{0,1\} \\
& \textstyle\sum_{i \in \mathscr{N}} \sum_{\substack{\mathscr{S}_j \in \mathscr{M}_i \\ o \in \mathscr{S}_j}} \mathbf{x_i^j} \leq 1 \text{ for all } o \in \mathscr{O} \quad \text{(OR constraint)} \\
\text{or}\ & \textstyle\sum_{\mathscr{S}_j \in \mathscr{M}_i} \mathbf{x_i^j} \leq 1 \text{ for all } i \in \mathscr{N} \qquad \text{(XOR constraint)}
\end{aligned}
$$

We can notice that this formulation of the OR and XOR Winner Determination Problem makes the problem strictly equivalent to the well-known knapsack problem. It implies that the general decision version of the problem is NP-complete (Rothkopf et al, 1998), but it also remains NP-complete even with very restrictive assumption about the values and the kind of bids allowed, and also on the number of agents (Lehmann et al, 2006).

4.3.2 Optimal Solving

In spite of the complexity of the WDP for OR, XOR languages and their variants, quite large instances can be nevertheless efficiently solved by state-of-the-art linear solvers running on the previous formulation of the WDP. However, the use of *ad hoc* branching approaches (see for instance Chapter 1 of Volume 2) that are tailored to this particular problem give even better results.

There are two natural ways of solving the WDP with a branching algorithm. The first possibility is to branch on objects, that is, to choose an object at each node of the search tree and to decide to which bid this object will be allocated. To take into account the free disposal assumption, a classical approach is to create a dummy bid containing all the objects and to which will be allocated all the disposed items. This branching approach can be used in combination with several methods that drastically reduce the size of the search space. For instance, some parts of the search tree can be pruned by only allocating objects to bids that have not already been considered in the previous branches. The second way of solving the WDP is to branch on bids, that is, to choose at each node of the search tree an atomic bid and decide whether it will be satisfied or not. Maintaining a conflict graph between bids, that updates when the bids are selected or discarded, dramatically improves the algorithm efficiency.

5 Conclusion

In this chapter, we have presented the foundations of (mainly) centralized collective decision making.[19]

This domain, that has originally been mostly studied by political scientists and economists, has recently met computer science and more specifically artificial intelligence. This very active scientific domain born from this convergence has been called *computational social choice*. To illustrate the scientific activity in this domain, we have presented in this chapter three prominent centralized collective making problems: voting, fair division and combinatorial auctions. For each of these domains, we have presented the main works related to artificial intelligence.

[19] The cake-cutting setting, that has also been introduced in this chapter, is an exception: the resolution of this problem is based on interactions between agents, which is by definition not a centralized process.

Centralized collective decision making proceeds by directly aggregating the agents preferences into a collective preferred decision that is not supposed to be changed afterwards. A different approach to collective decision making is to let the agents interact and negotiate: this is the case in *distributed* collective decision making, presented in Chapter 20 of this volume, that will complement the overview given in this chapter.

References

Ailon N, Charikar M, Newman A (2005) Aggregating inconsistent information: ranking and clustering. In: Proc. of ACM Symposium on Theory of computing (STOC'05)

Alon N (2006) Ranking tournaments. SIAM J Discrete Math 20(1):137–142

Amanatidis G, Markakis E, Nikzad A, Saberi A (2015) Approximation algorithms for computing maximin share allocations. In: ICALP (1), Springer, Lecture Notes in Computer Science, vol 9134, pp 39–51

Arrow K (1951) Social Choice and Individual Values. John Wiley and Sons, second edition 1963

Asadpour A, Saberi A (2007) An approximation algorithm for max-min fair allocation of indivisible goods. In: Proc. of ACM Symposium on Theory of computing (STOC'07), pp 114–121

Aziz H (2015) Competitive equilibrium with equal incomes for allocation of indivisible objects. Operations Research Letters 43(6):622–624

Aziz H, Mackenzie S (2016) A discrete and bounded envy-free cake cutting protocol for any number of agents. In: Foundations of Computer Science (FOCS), 2016 IEEE 57th Annual Symposium on, IEEE, pp 416–427

Aziz H, Gaspers S, Mackenzie S, Walsh T (2015) Fair assignment of indivisible objects under ordinal preferences. Artificial Intelligence 227:71–92

Bachrach Y, Betzler N, Faliszewski P (2010) Probabilistic possible-winner determination. In: Proc. AAAI Conf. on Artificial Intelligence (AAAI'10)

Balinski ML, Young HP (2001) Fair Representation : Meeting the Ideal of One Man One Vote, seconde edn. Brookings Institution Press

Bansal N, Sviridenko M (2006) The Santa Claus problem. In: Proc. of ACM Symposium on Theory of computing (STOC'06), pp 31–40

Bartholdi J, Orlin J (1991) Single transferable vote resists strategic voting. Social Choice and Welfare 8(4):341–354

Bartholdi J, Tovey C, Trick M (1989a) The computational difficulty of manipulating an election. Social Choice and Welfare 6(3):227–241

Bartholdi J, Tovey C, Trick M (1989b) Voting schemes for which it can be difficult to tell who won the election. Social Choice and Welfare 6(3):157–165

Bartholdi J, Tovey C, Trick M (1992) How hard is it to control an election? Mathematical and Computer Modeling 16(8/9):27–40

Betzler N, Niedermeier R, Woeginger GJ (2011) Unweighted coalitional manipulation under the Borda rule is NP-hard. In: Proc. Int. Joint Conf. on Artificial Intelligence (IJCAI'11)

Bezáková I, Dani V (2005) Allocating indivisible goods. SIGecom Exch 5(3):11–18

Boutilier C, Hoos HH (2001) Bidding languages for combinatorial auctions. In: Proc. Int. Joint Conf. on Artificial Intelligence (IJCAI'01), pp 1211–1217

Boutilier C, Rosenschein J (2016) Incomplete information and communication in voting. In: Brandt F, Conitzer V, Endriss U, Lang J, Procaccia AD (eds) Handbook of Computational Social Choice, Cambridge University Press, chap 10

Bouveret S, Lang J (2008) Efficiency and envy-freeness in fair division of indivisible goods: Logical representation and complexity. J of Artificial Intelligence Research 32:525–564

Bouveret S, Lemaître M (2009) Computing leximin-optimal solutions in constraint networks. Artificial Intelligence 173(2):343 – 364

Bouveret S, Lemaître M (2016) Characterizing conflicts in fair division of indivisible goods using a scale of criteria. Autonomous Agents and Multi-Agent Systems 30(2):259–290

Bouveret S, Fargier H, Lang J, Lemaître M (2005) Allocation of indivisible goods: a general model and some complexity results. In: Proc. Int. Joint Conf. on Autonomous Agents and Multiagent Systems (AAMAS'05)

Bouveret S, Endriss U, Lang J (2009) Conditional importance networks: A graphical language for representing ordinal, monotonic preferences over sets of goods. In: Proc. Int. Joint Conf. on Artificial Intelligence (IJCAI'09), pp 67–72

Bouveret S, Endriss U, Lang J (2010) Fair division under ordinal preferences: Computing envy-free allocations of indivisible goods. In: Proc. European Conf. on Artificial Intelligence (ECAI'10)

Brams S, Fishburn P (2004) Voting procedures. In: Arrow K, Sen A, Suzumura K (eds) Handbook of Social Choice and Welfare, Elsevier, chap 4

Brams S, Kilgour DM, Zwicker W (1998) The paradox of multiple elections. Social Choice and Welfare 15:211–236

Brams S, Jones MA, Klamler C (2006) Better ways to cut a cake. Notices of the American Mathematical Society 53(11):13141321

Brams S, Kilgour M, Sanver R (2007) A minimax procedure for electing committees. Public Choice 3-4(132):401–420

Brams SJ, Taylor AD (1995) An envy-free cake division protocol. The American Mathematical Monthly 102(1):9–18

Brams SJ, Taylor AD (1996) Fair Division — From Cake-cutting to Dispute Resolution. Cambridge University Press

Brams SJ, Taylor AD (2000) The win-win solution. Guaranteeing fair shares to everybody. W. W. Norton & Company

Brandt F, Brill M, Harrenstein P (2016a) Tournament solutions. In: Brandt F, Conitzer V, Endriss U, Lang J, Procaccia AD (eds) Handbook of Computational Social Choice, Cambridge University Press, chap 3

Brandt F, Conitzer V, Endriss U, Lang J, Procaccia AD (2016b) Handbook of Computational Social Choice. Cambridge University Press

Brandt F, Conitzer V, Endriss U, Lang J, Procaccia AD (2016c) Introduction to computational social choice. In: Brandt F, Conitzer V, Endriss U, Lang J, Procaccia AD (eds) Handbook of Computational Social Choice, Cambridge University Press, chap 1

Brânzei S, Hosseini H, Miltersen PB (2015) Characterization and computation of equilibria for indivisible goods. In: International Symposium on Algorithmic Game Theory, Springer, pp 244–255

Brânzei S, Caragiannis I, Kurokawa D, Procaccia AD (2016) An algorithmic framework for strategic fair division. In: AAAI, pp 418–424

Budish E (2011) The combinatorial assignment problem : approximate competitive equilibrium from equal incomes. J of Political Economy 119(6)

Caragiannis I, Covey JA, Feldman M, Homan CM, Kaklamanis C, Karanikolas N, Procaccia AD, Rosenschein JS (2009) On the approximability of Dodgson and Young elections. In: ACM-SIAM Symposium on Discrete Algorithms (SODA'09), pp 1058–1067

Caragiannis I, Hemaspaandra E, Hemaspaandra LA (2016a) Dodgson's Rule and Young's Rule. In: Brandt F, Conitzer V, Endriss U, Lang J, Procaccia AD (eds) Handbook of Computational Social Choice, Cambridge University Press, chap 5

Caragiannis I, Kurokawa D, Moulin H, Procaccia AD, Shah N, Wang J (2016b) The unreasonable fairness of maximum Nash welfare. In: Proceedings of the 2016 ACM Conference on Economics and Computation, ACM, pp 305–322

Chevaleyre Y, Endriss U, Estivie S, Maudet N (2004) Multiagent resource allocation with k-additive utility functions. In: Proc. DIMACS-LAMSADE Workshop on Computer Science and Decision Theory, vol 3, pp 83–100

Chevaleyre Y, Endriss U, Maudet N (2007) Allocating goods on a graph to eliminate envy. In: Proc. AAAI Conf. on Artificial Intelligence (AAAI'07)

Chevaleyre Y, Lang J, Maudet N, Ravilly-Abadie G (2009) Compiling the votes of a subelectorate. In: Proc. Int. Joint Conf. on Artificial Intelligence (IJCAI'09), pp 97–102

Chevaleyre Y, Lang J, Maudet N, Monnot J (2010) Possible winners when new candidates are added: the case of scoring rules. In: Proc. AAAI Conf. on Artificial Intelligence (AAAI'10)

Chevaleyre Y, Endriss U, Maudet N (2017) Distributed Fair Allocation of Indivisible Goods. Artificial Intelligence 242:1–22

de Condorcet N (1735) Essai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix. Imprimerie Royale, Paris

Conitzer V (2006) Computing Slater rankings using similarities among candidates. In: Proc. AAAI Conf. on Artificial Intelligence (AAAI'06)

Conitzer V, Sandholm T (2002a) Complexity of manipulating elections with few candidates. In: Proc. AAAI Conf. on Artificial Intelligence (AAAI'98)

Conitzer V, Sandholm T (2002b) Vote elicitation: Complexity and strategy-proofness. In: Proc. AAAI Conf. on Artificial Intelligence (AAAI'98), pp 392–397

Conitzer V, Sandholm T (2005) Communication complexity of common votiong rules. In: Proc. ACM Conf. on Electronic Commerce (EC'05)

Conitzer V, Walsh T (2016) Barriers to manipulation. In: Brandt F, Conitzer V, Endriss U, Lang J, Procaccia AD (eds) Handbook of Computational Social Choice, Cambridge University Press, chap 6

Conitzer V, Rognlie M, Xia L (2009) Preference functions that score rankings and maximum likelihood estimation. In: Proceedings of IJCAI-09, pp 109–115

Cramton P (2006) Simultaneous ascending auctions. In: Cramton P, Shoham Y, Steinberg R (eds) Combinatorial auctions, MIT Press, chap 4

Cramton P, Shoham Y, Steinberg R (eds) (2006) Combinatorial Auctions. MIT Press

Darmann A, Klamler C, Pferschy U (2009) Maximizing the minimum voter satisfaction on spanning trees. Mathematical Social Sciences 58(2):238–250

Davies J, Katsirelos G, Narodystka N, Walsh T (2011) Complexity of and algorithms for Borda manipulation. In: Proc. AAAI Conf. on Artificial Intelligence (AAAI'11)

Dickerson JP, Goldman JR, Karp J, Procaccia AD, Sandholm T (2014) The computational rise and fall of fairness. In: AAAI, vol 14, pp 1405–1411

Dunne PE (2005) Multiagent resource allocation in the presence of externalities. In: Proc. Int. Central and Eastern European Conf. on Multi-Agent Systems (CEEMAS'2005), pp 408–417

Dunne PE, Wooldridge M, Laurence M (2005) The complexity of contract negotiation. Artificial Intelligence 164(1-2):23–46

Dwork C, Kumar R, Naor M, Sivakumar D (2001) Rank aggregation methods for the web. In: Int. World Wide Web Conf. (WWW10), pp 613–622

Elkind E, Slinko A (2016) Rationalizations of voting rules. In: Brandt F, Conitzer V, Endriss U, Lang J, Procaccia AD (eds) Handbook of Computational Social Choice, Cambridge University Press, chap 8

Elkind E, Faliszewski P, Slinko AM (2010) Cloning in elections. In: Proc. AAAI Conf. on Artificial Intelligence (AAAI'10)

Endriss U (ed) (2017) Trends in Computational Social Choice. AI Access, to appear

Endriss U, Maudet N (2005) On the communication complexity of multilateral trading: Extend report. J of Autonomous Agents and Multi-Agent Systems 11(1):91–107

Ephrati E, Rosenschein JS (1993) Multi-agent planning as a dynamic search for social consensus. In: Proc. Int. Joint Conf. on Artificial Intelligence (IJCAI'93), pp 423–431

Escoffier B, Gourvès L, Monnot J (2013) Fair solutions for some multiagent optimization problems. Autonomous Agents and Multi-Agent Systems 26(2):184–201

Faliszewski P, Rothe J (2016) Control and bribery in voting. In: Brandt F, Conitzer V, Endriss U, Lang J, Procaccia AD (eds) Handbook of Computational Social Choice, Cambridge University Press, chap 7

Fischer F, Hudry O, Niedermeier R (2016) Weighted tournament solutions. In: Brandt F, Conitzer V, Endriss U, Lang J, Procaccia AD (eds) Handbook of Computational Social Choice, Cambridge University Press, chap 4

Fujishima Y, Leyton-Brown K, Shoam Y (1999) Taming the computational complexity of combinatorial auctions: Optimal and approximate approaches. In: Proc. Int. Joint Conf. on Artificial Intelligence (IJCAI'99)

Galand L, Perny P (2006) Search for compromise solutions in multiobjective state space graphs. In: Proc. European Conf. on Artificial Intelligence (ECAI'06), pp 93–97

Gibbard A (1973) Manipulation of voting schemes: A general result. Econometrica 41:587–601

Golden B, Perny P (2010) Infinite order Lorenz dominance for fair multiagent optimization. In: Proc. Int. Joint Conf. on Autonomous Agents and Multiagent Systems (AAMAS'10), pp 383–390

Harsanyi JC (1955) Cardinal welfare, individualistic ethics, and interpersonal comparisons of utility. J of Political Economy 63:309–321

Hemaspaandra E, Hemaspaandra LA, Rothe J (1997) Exact analysis of Dodgson elections: Lewis Carroll's 1876 system is complete for parallel access to NP. J of the ACM 44(6):806–825

Hudry O (2004a) Computation of median orders: Complexity results. In: Proc. DIMACS-LAMSADE Workshop on Computer Science and Decision Theory, vol 3, pp 179–214

Hudry O (2004b) A note on "Banks winners in tournaments are difficult to recognize" by G. J. Woeginger. Social Choice and Welfare 23(1):113–114

de Keijzer B, Bouveret S, Klos T, Zhang Y (2009) On the complexity of efficiency and envy-freeness in fair division of indivisible goods with additive preferences. In: Proc. Int. Conf. on Algorithmic Decision Theory (ADT'09)

Klamler C, Pferschy U (2007) The travelling group problem. Social Choice and Welfare 3(29):429–452

Konczak K, Lang J (2005) Voting procedures with incomplete preferences. In: Proc. IJCAI'05 Multidisciplinary Workshop on Advances in Preference Handling

Kurokawa D, Procaccia AD, Wang J (2015) When can the maximin share guarantee be guaranteed? Tech. rep., Carnegie Mellon University

Lacy D, Niou E (2000) A problem with referenda. J of Theoretical Politics 12(1):5–31

Lang J, Xia L (2009) Sequential composition of voting rules in multi-issue domains. Mathematical Social Sciences 57(3):304–324

Lang J, Xia L (2016) Voting in combinatorial domains. In: Brandt F, Conitzer V, Endriss U, Lang J, Procaccia AD (eds) Handbook of Computational Social Choice, Cambridge University Press, chap 9

Laslier JF, Sanver MR (eds) (2010) Handbook on Approval Voting. Studies in Choice and Welfare, Springer

Lehmann D, Mller R, Sandholm TW (2006) The winner determination problem. In: Cramton P, Shoham Y, Steinberg R (eds) Combinatorial auctions, MIT Press, chap 12

Lemaître M, Verfaillie G, Bataille N (1999) Exploiting a Common Property Resource under a Fairness Constraint: a Case Study. In: Proc. Int. Joint Conf. on Artificial Intelligence (IJCAI'99), pp 206–211

Lesca J, Perny P (2010) LP solvable models for multiagent fair allocation problems. In: Proc. European Conf. on Artificial Intelligence (ECAI'10), pp 387–392
Lipton R, Markakis E, Mossel E, Saberi A (2004) On approximately fair allocations of divisible goods. In: Proc. ACM Conf. on Electronic Commerce (EC'04)
Luss H (1999) On equitable resource allocation problems: a lexicographic minimax approach. Operations Research 47(3):361–378
May K (1952) A set of independent necessary and sufficient conditions for simple majority decisions. Econometrica 20:680–684
McCabe-Dansted J, Pritchard G, Slinko A (2008) Approximability of Dodgson's rule. Social Choice and Welfare 31(2):311–330
Moulin H (1988) Axioms of Cooperative Decision Making. Cambridge University Press
Moulin H (2003) Fair division and collective welfare. MIT Press
Nguyen NT, Nguyen TT, Roos M, Rothe J (2014) Computational complexity and approximability of social welfare optimization in multiagent resource allocation. Autonomous agents and multi-agent systems 28(2):256–289
Nisan N (2006) Bidding languages for combinatorial auctions. In: Cramton P, Shoham Y, Steinberg R (eds) Combinatorial auctions, MIT Press, chap 9
Ogryczak W (1997) On the lexicographic minimax approach to location problems. European J of Operational Research 100:566–585
Procaccia AD (2009) Thou shalt covet thy neighbors cake. In: Proc. Int. Joint Conf. on Artificial Intelligence (IJCAI'09), pp 239–244
Procaccia AD (2016) Cake cutting algorithms. In: Brandt F, Conitzer V, Endriss U, Lang J, Procaccia AD (eds) Handbook of Computational Social Choice, Cambridge University Press, chap 13
Procaccia AD, Rosenschein JS (2007) Junta distributions and the average-case complexity of manipulating elections. J of Artificial Intelligence Research 28:157–181
Procaccia AD, Wang J (2014) Fair enough: Guaranteeing approximate maximin shares. In: Proc. 14th ACM Conference on Economics and Computation (EC'14), forthcoming
Rassenti S, Smith VL, Bulfin RL (1982) A combinatorial auction mechanisms for airport time slot allocation. Bell J of Economics pp 402–417
Rawls J (1971) A Theory of Justice. Harvard University Press, Cambridge, Mass., traduction franaise disponible aux éditions du Seuil
Robertson J, Webb W (1998) Cake-Cutting Algorithms: Be Fair If You Can. AK Peters Ltd
Rothe J, Rothe I (2015) Economics and Computation: An Introduction to Algorithmic Game Theory, Computational Social Choice, and Fair Division. Springer
Rothe J, Spakowski H, Vogel J (2003) Exact complexity of the winner for Young elections. Theory of Computing Systems 36(4):375–386
Rothkopf MH, Pekeč A, Harstad RM (1998) Computationally manageable combinatorial auctions. Management Science 44(8):1131–1147
Sandholm TW (1998) Contract types for satisficing task allocation: I. theoretical results. In: Proc. AAAI Spring Symposium: Satisficing Models, pp 68–75

Sandholm TW (1999) An algorithm for optimal winner determination in combinatorial auctions. In: Proc. Int. Joint Conf. on Artificial Intelligence (IJCAI'99), pp 452–547

Sandholm TW (2002) Algorithm for optimal winner determination in combinatorial auctions. Artificial Intelligence 134:1–54

Satterthwaite MA (1975) Strategy-proofness and arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. J of Economic Theory 10(2):187–217

Sen AK (1970) Collective choice and social welfare. North-Holland

Vickrey W (1961) Counterspeculation, auctions, and competitive sealed tenders. J of Finance 16:8–37

Walsh T (2008) Complexity of terminating preference elicitation. In: Proc. Int. Joint Conf. on Autonomous Agents and Multiagent Systems (AAMAS'08), pp 967–974

Walsh T (2010) Online cake cutting. In: Third Int. Workshop on Computational Social Choice

Woeginger GJ (2003) Banks winners in tournaments are difficult to recognize. Social Choice and Welfare 20(3):523–528

Xia L, Conitzer V (2008) Determining possible and necessary winners under common voting rules given partial orders. In: Proc. AAAI Conf. on Artificial Intelligence (AAAI'08), pp 196–201

Xia L, Conitzer V (2010) Compilation complexity of common voting rules. In: Proc. AAAI Conf. on Artificial Intelligence (AAAI'10)

Yager RR (1988) On ordered weighted averaging aggregation operators in multicriteria decision making. IEEE Transactions on Systems, Man, and Cybernetics 18:183–190

Young HP (1994) Equity in Theory and Practice. Princeton University Press

Zwicker WS (2016) Introduction to the theory of voting. In: Brandt F, Conitzer V, Endriss U, Lang J, Procaccia AD (eds) Handbook of Computational Social Choice, Cambridge University Press, chap 2