CHAPTER 1

Introduction to Computational Social Choice

Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia

1.1 Computational Social Choice at a Glance

Social choice theory is the field of scientific inquiry that studies the aggregation of individual preferences toward a collective choice. For example, social choice theorists who hail from a range of different disciplines, including mathematics, economics, and political science—are interested in the design and theoretical evaluation of voting rules. Questions of social choice have stimulated intellectual thought for centuries. Over time, the topic has fascinated many a great mind, from the Marquis de Condorcet and Pierre-Simon de Laplace, through Charles Dodgson (better known as Lewis Carroll, the author of *Alice in Wonderland*), to Nobel laureates such as Kenneth Arrow, Amartya Sen, and Lloyd Shapley.

Computational social choice (COMSOC), by comparison, is a very young field that formed only in the early 2000s. There were, however, a few precursors. For instance, David Gale and Lloyd Shapley's algorithm for finding stable matchings between two groups of people with preferences over each other, dating back to 1962, truly had a computational flavor. And in the late 1980s, a series of papers by John Bartholdi, Craig Tovey, and Michael Trick showed that, on the one hand, computational complexity, as studied in theoretical computer science, can serve as a barrier against strategic manipulation in elections, but on the other hand, it can also prevent the efficient use of some voting rules altogether. Around the same time, a research group around Bernard Monjardet and Olivier Hudry also started to study the computational complexity of preference aggregation procedures.

Assessing the computational difficulty of determining the output of a voting rule, or of manipulating it, is a wonderful example of the importation of a concept from one field, theoretical computer science, to what at that time was still considered an entirely different one, social choice theory. It is this interdisciplinary view on collective decision making that defines computational social choice as a field. But, importantly, the contributions of computer science to social choice theory are not restricted to the design and analysis of algorithms for preexisting social choice problems. Rather, the arrival of computer science on the scene led researchers to revisit the old problem of 2

Cambridge University Press & Assessment 978-1-107-06043-2 — Handbook of Computational Social Choice Edited by Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, Ariel D. Procaccia, Foreword by Hervé Moulin Excerpt <u>More Information</u>

1 INTRODUCTION TO COMPUTATIONAL SOCIAL CHOICE

social choice from scratch. It offered new perspectives, and it led to many new types of questions, thereby arguably contributing significantly to a revival of social choice theory as a whole.

Today, research in computational social choice has two main thrusts. First, researchers seek to apply computational paradigms and techniques to provide a better analysis of social choice mechanisms, and to construct new ones. Leveraging the theory of computer science, we see applications of computational complexity theory and approximation algorithms to social choice. Subfields of artificial intelligence, such as machine learning, reasoning with uncertainty, knowledge representation, search, and constraint reasoning, have also been applied to the same end.

Second, researchers are studying the application of social choice theory to computational environments. For example, it has been suggested that social choice theory can provide tools for making joint decisions in multiagent system populated by heterogeneous, possibly selfish, software agents. Moreover, it is finding applications in group recommendation systems, information retrieval, and crowdsourcing. Although it is difficult to change a political voting system, such low-stake environments allow the designer to freely switch between choice mechanisms, and therefore they provide an ideal test bed for ideas coming from social choice theory.

This book aims to provide an authoritative overview of the field of computational social choice. It has been written for students and scholars from both computer science and economics, as well as for others from the mathematical and social sciences more broadly. To position the field in its wider context, in Section 1.2, we provide a brief review of the history of social choice theory. The structure of the book reflects the internal structure of the field. We provide an overview of this structure by briefly introducing each of the remaining 18 chapters of the book in Section 1.3. As computational social choice is still rapidly developing and expanding in scope every year, naturally, the coverage of the book cannot be exhaustive. Section 1.4 therefore briefly introduces a number of important active areas of research that, at the time of conceiving this book, were not yet sufficiently mature to warrant their own chapters. Section 1.5, finally, introduces some basic concepts from theoretical computer science, notably the fundamentals of computational complexity theory, with which some readers may not be familiar.

1.2 History of Social Choice Theory

Modern research in computational social choice builds on a long tradition of work on collective decision making. We can distinguish three periods in the study of collective decision making: early ideas regarding specific rules going back to antiquity; the classical period, witnessing the development of a general mathematical theory of social choice in the second half of the twentieth century; and the "computational turn" of the very recent past. We briefly review each of these three periods by providing a small selection of illustrative examples.

1.2.1 Early Ideas: Rules and Paradoxes

Collective decision-making problems come in many forms. They include the question of how to fairly divide a set of resources, how to best match people on the basis of their

1.2 HISTORY OF SOCIAL CHOICE THEORY

3

preferences, and how to aggregate the beliefs of several individuals. The paradigmatic example, however, is voting: how should we aggregate the individual preferences of several voters over a given set of alternatives so as to be able to choose the "best" alternative for the group? This important question has been pondered by a number of thinkers for a long time. Also the largest part of this book, Part I, is devoted to voting. We therefore start our historic review of social choice theory with a discussion of early ideas pertaining to voting.¹

Our first example for the discussion of a problem in voting goes back to Roman times. Pliny the Younger, a Roman senator, described in A.D. 105 the following problem in a letter to an acquaintance. The Senate had to decide on the fate of a number of prisoners: acquittal (A), banishment (B), or condemnation to death (C). Although option A, favored by Pliny, had the largest number of supporters, it did not have an absolute majority. One of the proponents of harsh punishment then strategically moved to withdraw proposal C, leaving its former supporters to rally behind option B, which easily won the majority contest between A and B. Had the senators voted on all three options, using the *plurality rule* (under which the alternative ranked at the top by the highest number of voters wins), option A would have won. This example illustrates several interesting features of voting rules. First, it may be interpreted as demonstrating a lack of fairness of the plurality rule: even though a majority of voters believes A to be inferior to one of the other options (namely, B), A still wins. This and other fairness properties of voting rules are reviewed in Chapter 2. Second, Pliny's anecdote is an instance of what nowadays is called *election control by deleting candidates*. By deleting C, Pliny's adversary in the senate was able to ensure that B rather than A won the election. Such control problems, particularly their algorithmic aspects, are discussed in Chapter 7. Third, the example also illustrates the issue of *strategic manipulation*. Even if option C had not been removed, the supporters of C could have manipulated the election by pretending that they supported B rather than C, thereby ensuring a preferred outcome, namely, B rather than A. Manipulation is discussed in depth in Chapters 2 and 6.

In the Middle Ages, the Catalan philosopher, poet, and missionary Ramon Llull (1232–1316) discussed voting rules in several of his writings. He supported the idea that election outcomes should be based on direct majority contests between pairs of candidates. Such voting rules are discussed in detail in Chapter 3. What exact rule he had in mind cannot be unambiguously reconstructed anymore, but it may have been the rule that today is known as the *Copeland rule*, under which the candidate who wins the largest number of pairwise majority contests is elected. Whereas Pliny specifically discussed the subjective interests of the participants, Llull saw voting as a means of revealing the divine truth about who is the objectively best candidate, for example, to fill the position of abbess in a convent. The mathematical underpinnings of this *epistemic perspective* on voting are discussed in Chapter 8.

Our third example is taken from the period of the Enlightenment. The works of the French engineer Jean-Charles de Borda (1733–1799) and the French philosopher and mathematician Marie Jean Antoine Nicolas de Caritat (1743–1794), better known

¹ There are also instances of very early writings on other aspects of social choice. A good example is the discussion of fair division problems in the Talmud, as noted and analyzed in modern terms by game theorists Aumann and Maschler (1985).

4

Cambridge University Press & Assessment 978-1-107-06043-2 — Handbook of Computational Social Choice Edited by Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, Ariel D. Procaccia, Foreword by Hervé Moulin Excerpt <u>More Information</u>

1 INTRODUCTION TO COMPUTATIONAL SOCIAL CHOICE

as the Marquis de Condorcet—and particularly the lively dispute between them—are widely regarded as the most significant contributions to social choice theory in the early period of the field. In 1770, Borda proposed a method of voting, today known as the *Borda rule*, under which each voter ranks all candidates, and each candidate receives as may points from a given voter as that voter ranks other candidates below her. He argued for the superiority of his rule over the plurality rule by discussing an example similar to that of Pliny, where the plurality winner would lose in a direct majority contest to another candidate, while the Borda winner does not have that deficiency. But Condorcet argued against Borda's rule on very similar grounds. Consider the following scenario with 3 candidates and 11 voters, which is a simplified version of an example Condorcet described in 1788:

4	3	2	2
Peter	Paul	Paul	James
Paul	James	Peter	Peter
James	Peter	James	Paul

In this example, four voters prefer candidate Peter over candidate Paul, whom they prefer over candidate James, and so forth. Paul wins this election both under the plurality rule (with 3 + 2 = 5 points) and the Borda rule (with $4 \cdot 1 + 3 \cdot 2 + 2 \cdot 2 + 2 \cdot 0 = 14$ points). However, a majority of voters (namely, 6 out of 11) prefer Peter to Paul. In fact, Peter also wins against James in a direct majority contest, so there arguably is a very strong case for rejecting voting rules that would not elect Peter in this situation. In today's terminology, we call Peter the *Condorcet winner*.

Now suppose two additional voters join the election, who both prefer James, to Peter, to Paul. Then a majority prefers Peter to Paul, and a majority prefers Paul to James, but now also a majority prefers James to Peter. This, the fact that the majority preference relation may turn out to be cyclic, is known as the *Condorcet paradox*. It shows that Condorcet's proposal, to be guided by the outcomes of pairwise majority contests, does not always lead to a clear election outcome.

In the nineteenth century, the British mathematician and story teller Charles Dodgson (1832–1898), although believed to have been unaware of Condorcet's work, suggested a voting rule designed to circumvent this difficulty. In cases where there is a single candidate who beats every other candidate in pairwise majority contests, he proposed to elect that candidate (the Condorcet winner). In all other cases, he proposed to count how many elementary changes to the preferences of the voters would be required before a given candidate would become the Condorcet winner, and to elect the candidate for which the required number of changes is minimal. In this context, he considered the swap of two candidates occurring adjacently in the preference list of a voter as such an elementary change. The *Dodgson rule* is analyzed in detail in Chapter 5.

This short review, it is hoped, gives the reader some insight into the kinds of questions discussed by the early authors. The first period in the history of social choice theory is reviewed in depth in the fascinating collection edited by McLean and Urken (1995).

1.2.2 Classical Social Choice Theory

While early work on collective decision making was limited to the design of specific rules and on finding fault with them in the context of specific examples, around the

1.2 HISTORY OF SOCIAL CHOICE THEORY

5

middle of the twentieth century, the focus suddenly changed. This change was due to the seminal work of Kenneth Arrow, who, in 1951, demonstrated that the problem with the majority rule highlighted by the Condorcet paradox is in fact much more general. Arrow proved that there exists no reasonable preference aggregation rule that does not violate at least one of a short list of intuitively appealing requirements (Arrow, 1951). That is, rather than proposing a new rule or pointing out a specific problem with an existing rule, Arrow developed a mathematical framework for speaking about and analyzing all possible such rules.

Around the same time, in related areas of economic theory, Nash (1950) published his seminal paper on the bargaining problem, which is relevant to the theory of fair allocation treated in Part II of this book, and Shapley (1953) published his groundbreaking paper on the solution concept for cooperative games now carrying his name, which plays an important role in coalition formation, to which Part III of this book is devoted. What all of these classical papers have in common is that they specified philosophically or economically motivated requirements in mathematically precise terms, as so-called *axioms*, and then rigorously explored the logical consequences of these axioms. As an example of this kind of axiomatic work of this classical period, let us review Arrow's result in some detail.

Let $N = \{1, ..., n\}$ be a finite set of *individuals* (or *voters*, or *agents*), and let A be a finite set of *alternatives* (or *candidates*). The set of all *weak orders* \succeq on A, that is, the set of all binary relations on A that are complete and transitive, is denoted as $\mathcal{R}(A)$, and the set of all *linear orders* \succeq on A, which in addition are antisymmetric, is denoted as $\mathcal{L}(A)$. In both cases, we use \succ to denote the strict part of \succeq . We use weak orders to model preferences over alternatives that permit ties and linear orders to model strict preferences. A *social welfare function* (SWF) is a function of the form $f : \mathcal{L}(A)^n \to \mathcal{R}(A)$. That is, f is accepting as input a so-called *profile* $P = (\succeq_1, \ldots, \succeq_n)$ of preferences, one for each individual, and maps it to a single preference order, which we can think of as representing a suitable compromise. We allow ties in the output, but not in the individual preferences. When f is clear from the context, we write \succeq for $f(\succeq_1, \ldots, \succeq_n)$, the outcome of the aggregation, and refer to it as the *social preference order*.

Arrow argued that any reasonable SWF should be *weakly Paretian* and *independent* of irrelevant alternatives (IIA). An SWF f is weakly Paretian if, for any two alternatives $a, b \in A$, it is the case that, if $a \succ_i b$ for all individuals $i \in N$, then also $a \succ b$. That is, if everyone strictly prefers a to b, then also the social preference order should rank a strictly above b. An SWF f is IIA if, for any two alternatives $a, b \in A$, the relative ranking of a and b by the social preference order \succeq only depends on the relative rankings of a and b provided by the individuals—but not, for instance, on how the individuals rank some third alternative c. To understand that it is not straightforward to satisfy these two axioms, observe that, for instance, the SWF that ranks alternatives in the order of frequency with which they appear in the top position of an individual preferable is not Paretian. The majority rule, while easily seen to be both Paretian and IIA, is not an SWF, because it does not always return a weak order, as the Condorcet paradox has shown.

An example of an SWF that most people would consider rather unreasonable is a *dictatorship*. We say that the SWF f is a dictatorship if there exists an individual

6

Cambridge University Press & Assessment 978-1-107-06043-2 — Handbook of Computational Social Choice Edited by Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, Ariel D. Procaccia, Foreword by Hervé Moulin Excerpt <u>More Information</u>

1 INTRODUCTION TO COMPUTATIONAL SOCIAL CHOICE

 $i^* \in N$ (the dictator) such that, for all alternatives $a, b \in A$, it is the case that $a \succ_{i^*} b$ implies $a \succ b$. Thus, f simply copies the (strict) preferences of the dictator, whatever the preferences of the other individuals. Now, it is not difficult to see that every dictatorship is both Paretian and IIA. The surprising—if not outright disturbing—result due to Arrow is that the converse is true as well:

Theorem 1.1 (Arrow, 1951). When there are three or more alternatives, then every SWF that is weakly Paretian and IIA must be a dictatorship.

Proof. Suppose $|A| \ge 3$, and let f be any SWF that is weakly Paretian and IIA. For any profile P and alternatives $a, b \in A$, let $N_{a \succ b}^P \subseteq N$ denote the set of individuals who rank a strictly above b in P. We call a coalition $C \subseteq N$ of individuals a *decisive coalition* for alternative a versus alternative b if $N_{a \succ b}^P \supseteq C$ implies $a \succ b$, that is, if everyone in C ranking a strictly above b is a sufficient condition for the social preference order to do the same. Thus, to say that f is weakly Paretian is the same as to say that the grand coalition N is decisive, and to say that f is dictatorial is the same as to say that there exists a singleton that is decisive. We call C weakly decisive for avs. b if we have at least that $N_{a \succ b}^P = C$ implies $a \succ b$.

We first show that *C* being weakly decisive for *a* versus *b* implies *C* being (not just weakly) decisive for *all* pairs of alternatives. This is sometimes called the *Contagion Lemma* or the *Field Expansion Lemma*. So let *C* be weakly decisive for *a* versus *b*. We show that *C* is also decisive for *a'* versus *b'*. We do so under the assumption that a, b, a', b' are mutually distinct (the other cases are similar). Consider any profile *P* such that $a' \succ_i a \succ_i b \succ_i b'$ for all $i \in C$, and $a' \succ_j a, b \succ_j b'$, and $b \succ_j a$ for all $j \notin C$. Then, from weak decisiveness of *C* for *a* versus *b* we get $a \succ b$; from *f* being weakly Paretian, we get $a' \succ a$ and $b \succ b'$, and thus from transitivity, we get $a' \succ b'$. Hence, in the specific profile *P* considered, the members of *C* ranking *a'* above *b'* was sufficient for *a'* getting ranked above *b'* also in the social preference order. But note that, first, we did not have to specify how individuals outside of *C* rank *a'* versus *b'*, and that, second, due to *f* being IIA, the relative ranking of *a'* versus *b'* can only depend on the individual rankings of *a'* versus *b'*. Hence, the only part of our construction that actually mattered was that everyone in *C* ranked *a'* above *b'*. So *C* really is decisive for *a'* versus *b'* as claimed.

Consider any coalition $C \subseteq N$ with $|C| \ge 2$ that is decisive (for some pair of alternatives, and thus for all pairs). Next, we will show that we can always split *C* into two nonempty subsets C_1, C_2 with $C_1 \cup C_2 = C$ and $C_1 \cap C_2 = \emptyset$ such that one of C_1 and C_2 is decisive for all pairs as well. This is sometimes called the *Splitting Lemma* or the *Group Contraction Lemma*. Recall that $|A| \ge 3$. Consider a profile *P* in which everyone ranks alternatives a, b, c in the top three positions and, furthermore, $a \succ_i b \succ_i c$ for all $i \in C_1, b \succ_j c \succ_j a$ for all $j \in C_2$, and $c \succ_k a \succ_k b$ for all $k \notin C_1 \cup C_2$. As $C = C_1 \cup C_2$ is decisive, we certainly get $b \succ c$. By completeness, we must have either $a \succ c$ or $c \succeq a$. In the first case, we have a situation where exactly the individuals in C_1 rank *a* above *c*. That is, C_1 is weakly decisive for *a* versus *c*. Hence, by the Contagion Lemma, C_1 is in fact decisive for all pairs. In the second case

1.2 HISTORY OF SOCIAL CHOICE THEORY

7

 $(c \succeq a)$, transitivity and $b \succ c$ imply that $b \succ a$. Hence, by an analogous argument as before, C_2 must be decisive for all pairs.

Recall that, due to f being weakly Paretian, N is a decisive coalition. We can now apply the Splitting Lemma again and again, to obtain smaller and smaller decisive coalitions, until we obtain a decisive coalition with just a single member. This inductive argument is admissible, because N is finite. But the existence of a decisive coalition with just one element means that f is dictatorial.

Arrow's Theorem is often interpreted as an impossibility result: it is impossible to devise an SWF for three or more alternatives that is weakly Paretian, IIA, and nondictatorial. The technique we have used to prove it is also used in Chapter 2 on voting theory and in Chapter 17 on judgment aggregation. These chapters also discuss possible approaches for dealing with such impossibilities by weakening our requirements somewhat.

The authoritative reference on classical social choice theory is the two-volume *Handbook of Social Choice and Welfare* edited by Arrow et al. (2002, 2010). There also are several excellent textbooks available, each covering a good portion of the field. These include the books by Moulin (1988a), Austen-Smith and Banks (2000, 2005), Taylor (2005), Gaertner (2006), and Nitzan (2010).

1.2.3 The Computational Turn

As indicated, Arrow's Theorem (from 1951) is generally considered the birth of modern social choice theory. The work that followed mainly consisted in axiomatic, or normative, results. Some of these are negative (Arrow's Theorem being an example). Others have a more positive flavor, such as the characterization of certain voting rules, or certain families of voting rules, by a set of properties. However, a common point is that these contributions (mostly published in economics or mathematics journals) neglected the *computational* effort required to determine the outcome of the rules they sought to characterize, and failed to notice that this computational effort could sometimes be prohibitive. Now, the practical acceptability of a voting rule or a fair allocation mechanism depends not only on its normative properties (who would accept a voting rule that is considered unfair by society?), but also on its implementability in a reasonable time frame (who would accept a voting rule that needs years for the outcome to be computed?). This is where computer science comes into play, starting in the late 1980s. For the first time, social choice became a field investigated by computer scientists from various fields (especially artificial intelligence, operations research, and theoretical computer science) who aimed at using computational concepts and algorithmic techniques for solving complex collective decision making problems.

A paradigmatic example is *Kemeny's rule*, studied in detail in Chapter 4. Kemeny's rule was not explicitly defined during the early phase of social choice, but it appears implicitly in Condorcet's works, as discussed, for instance, in Chapter 8. It played a key role in the second phase of social choice: it was defined formally by John G. Kemeny in 1959, characterized axiomatically by H. Peyton Young and Arthur B. Levenglick in 1978, and rationalized as a maximum likelihood estimator for recovering the ground truth by means of voting in a committee by Young in 1988. Finally, it was recognized

8

1 INTRODUCTION TO COMPUTATIONAL SOCIAL CHOICE

as a computationally difficult rule, independently and around the same time (the "early phase of computational social choice") by John Bartholdi, Craig Tovey, and Michael Trick, as well as by Olivier Hudry and others. None of these papers, however, succeeded in determining the exact complexity of Kemeny's rule, which was done only in 2005, at the time when computational social choice was starting to expand rapidly. Next came practical algorithms for computing Kemeny's rule, polynomial-time algorithms for approximating it, parameterized complexity studies, and applications to various fields, such as databases or "web science." We took Kemeny's rule as an example, but there are similar stories to be told about other preference aggregation rules, as well as for various fair allocation and matching mechanisms.

Deciding when computational social choice first appeared is not easy. Arguably, the Gale-Shapley algorithm (1962), discussed in Chapter 14, deals both with social choice and with computation (and even with communication, since it can also be seen as an interaction protocol for determining a stable matching). Around the same time, the Dubins-Spanier Algorithm (Dubins-Spanier, 1961), discussed in Chapter 13, was one of the first important contributions in the formal study of cake cutting, that is, of fairly partitioning a divisible resource (again, this "algorithm" can also be seen as an interaction protocol). Just as for preference aggregation, the first computational studies appeared in the late 1980s. Finally, although formal computational studies of the fair allocation of indivisible goods appeared only in the early 2000s, they are heavily linked to computational issues in combinatorial auctions, the study of which dates back to the 1980s.

By the early 2000s this trend toward studying collective decision making in the tradition of classical social choice theory, yet with a specific focus on computational concerns, had reached substantial momentum. Researchers coming from different fields and working on different specific problems started to see the parallels to the work of others. The time was ripe for a new research community to form around these ideas. In 2006 the first edition of the COMSOC Workshop, the biannual International Workshop on Computational Social Choice, took place in Amsterdam. The announcement of this event was also the first time that the term "*computational social choice*" was used explicitly to define a specific research area.

Today, computational social choice is a booming field, carried by a large and growing community of active researchers, making use of a varied array of methodologies to tackle a broad range of questions. There is increasing interaction with representatives of classical social choice theory in economics, mathematics, and political science. There is also increasing awareness of the great potential of computational social choice for important applications of decision-making technologies, in areas as diverse as policy making (e.g., matching junior doctors to hospitals), distributed computing (e.g., allocating bandwidth to processes), and education (e.g., aggregating student evaluations gathered by means of peer assessment methods). Work on computational social choice is regularly published in major journals in artificial intelligence, theoretical computer science, operations research, and economic theory—and occasionally also in other disciplines, such as logic, philosophy, and mathematics. As is common practice in computer science, a lot of work in the field is also published in the archival proceedings of peer-reviewed conferences, particularly the major international conferences on artificial intelligence, multiagent systems, and economics and computation.

1.3 BOOK OUTLINE

9

1.3 Book Outline

This book is divided into four parts, reflecting the structure of the field of computational social choice. Part I, taking up roughly half of the book, focuses on the design and analysis of voting rules (which aggregate individual preferences into a collective decision). The room given to this topic here mirrors the breadth and depth with which the problem of voting has been studied to date.

The remaining three parts consist of three chapters each. Part II covers the problem of allocating goods to individuals with heterogeneous preferences in a way that satisfies rigorous notions of fairness. We make the distinction between divisible and indivisible goods. Part III addresses questions that arise when agents can form coalitions and each have preferences over these coalitions. This includes two-sided matching problems (e.g., between junior doctors seeking an internship and hospitals), hedonic games (where agents' preferences depend purely on the members of the coalition they are part of), and weighted voting games (where coalitions emerge to achieve some goal, such as passing a bill in parliament).

Much of classical (noncomputational) social choice theory deals with voting (Part I). In contrast, fair allocation (Part II) and coalition formation (Part III) are not always seen as subfields of (classical) social choice theory, but, interestingly, their intersection with computer science has become part of the core of computational social choice, due to sociological reasons having to do with how the research community addressing these topics has evolved over the years.

Part IV, finally, covers topics that did not neatly fit into the first three thematic parts. It includes chapters on logic-based judgment aggregation, on applications of the axiomatic method to reputation and recommendation systems found on the Internet, and on knockout tournaments (as used, for instance, in sports competitions). Next, we provide a brief overview of each of the book's chapters.

1.3.1 Part I: Voting

Chapter 2: Introduction to the Theory of Voting (Zwicker). This chapter provides an introduction to the main classical themes in voting theory. This includes the definition of the most important voting rules, such as Borda's, Copeland's, and Kemeny's rule. It also includes an extensive introduction to the axiomatic method and proves several characterization and impossibility theorems, thereby complementing our brief exposition in Section 1.2.2. Special attention is paid to the topic of strategic manipulation in elections.

Chapter 2 also introduces Fishburn's classification of voting rules. Fishburn used this classification to structure the set of Condorcet extensions, the family of rules that respect the principle attributed to the Marquis de Condorcet, by which any alternative that beats all other alternatives in direct pairwise contests should be considered the winner of the election. Fishburn's classification groups these Condorcet extensions into three classes—imaginatively called C1, C2, and C3—and the following three chapters each present methods and results pertaining to one of these classes.

Chapter 3: Tournament Solutions (Brandt, Brill, and Harrenstein). This chapter deals with voting rules that only depend on pairwise majority comparisons, so-called

10 1 INTRODUCTION TO COMPUTATIONAL SOCIAL CHOICE

C1 functions. Pairwise comparisons can be conveniently represented using directed graphs. When there is an odd number of voters with linear preferences, these graphs are tournaments, that is, oriented complete graphs. Topics covered in this chapter include McGarvey's Theorem, various tournament solutions (such as Copeland's rule, the top cycle, or the bipartisan set), strategyproofness, implementation via binary agendas, and extensions of tournament solutions to weak tournaments. Particular attention is paid to the issue of whether and how tournament solutions can be computed efficiently.

Chapter 4: Weighted Tournament Solutions (Fischer, Hudry, and Niedermeier). This chapter deals with voting rules that only depend on weighted pairwise majority comparisons, so-called C2 functions. Pairwise comparisons can be conveniently represented using weighted directed graphs, where the weight of an edge from alternative x to alternative y is the number of voters who prefer x to y. Prominent voting rules of type C2 are Kemeny's rule, the maximin rule, the ranked pairs method, Schulze's method, and—anecdotally—Borda's rule. The chapter focusses on the computation, approximation, and fixed-parameter tractability of these rules, while paying particular attention to Kemeny's rule.

Chapter 5: Dodgson's Rule and Young's Rule (Caragiannis, Hemaspaandra, and Hemaspaandra). This chapter focuses on two historically significant voting rules belonging to C3, the class of voting rules requiring strictly more information than a weighted directed graph, with computationally hard winner determination problems. The complexity of this problem is analyzed in depth. Methods for circumventing this intractability—approximation algorithms, fixed-parameter tractable algorithms, and heuristic algorithms—are also discussed.

The remaining five chapters in Part I all focus on specific methodologies for the analysis of voting rules.

Chapter 6: Barriers to Manipulation in Voting (Conitzer and Walsh). This chapter concerns the manipulation problem, where a voter misreports her preferences in order to obtain a better result for herself, and how to address it. It covers the Gibbard-Satterthwaite impossibility result, which roughly states that manipulation cannot be completely avoided in sufficiently general settings, and its implications. It then covers some ways of addressing this problem, focusing primarily on erecting computational barriers to manipulation—one of the earliest lines of research in computational social choice, as alluded to before.

Chapter 7: Control and Bribery in Voting (Faliszewski and Rothe). Control and bribery are variants of manipulation, typically seen as carried out by the election organizer. Paradigmatic examples of control include adding or removing voters or alternatives. Bribery changes the structure of voters' preferences, without changing the structure of the entire election. This chapter presents results regarding the computational complexity of bribery and control problems under a variety of voting rules. Much like Chapter 6, the hope here is to obtain computational hardness in order to prevent strategic behavior.

Chapter 8: Rationalizations of Voting Rules (Elkind and Slinko). While the bestknown approach in social choice to justify a particular voting rule is the axiomatic one,