Voting and Judgment Aggregation: An Axiomatic and Algorithmic Analysis

Inaugural-Dissertation

zur Erlangung des Doktorgrades
der Mathematisch-Naturwissenschaftlichen Fakultät
der Heinrich-Heine-Universität Düsseldorf

vorgelegt von

Ann-Kathrin Selker
aus Georgsmarienhütte

Düsseldorf, Dezember 2021

aus dem Institut für Informatik
der Heinrich-Heine-Universität Düsseldorf

Gedruckt mit der Genehmigung der
Mathematisch-Naturwissenschaftlichen Fakultät der
Heinrich-Heine-Universität Düsseldorf

Berichterstatter:

1. Jun.-Prof. Dr. Dorothea Baumeister

2. Prof. Dr. Jörg Rothe

Tag der mündlichen Prüfung:
16.12.2021

Abstract

This thesis deals with preference and judgment aggregation in the context of computational social choice, a subfield of multiagent systems and artificial intelligence. In preference aggregation, agents (typically called voters) have preferences over a given set of candidates and the goal is to aggregate ballots derived from these preferences to determine an output (e.g., a single candidate, a set of candidates, or a linear order over the candidates). In judgment aggregation, agents (typically called judges) have judgments over a given set of issues and the goal is to aggregate these judgments into a collective set of judgments.

In the context of multiwinner elections that elect a committee of candidates, this thesis studies a new type of ballot called ℓ-ballots. In contrast to existing ballot types, ℓ-ballots are a compromise between ordinal and cardinal ballots. This thesis focuses on the axiomatic properties of two newly defined types of multiwinner voting rules using these ℓ-ballots as input, and proposes a generalization of ℓ-ballots to fully cardinal ballots.

Furthermore, this thesis explores the computational complexity of strategic attacks on voting rules. For several prominent iterative voting rules, i.e., voting rules that proceed in rounds, it is shown that shift bribery is NP-complete in all considered cases. In the context of iterative voting, voters are encouraged to repeatedly update their ballots as a reaction to the current state of the election. This thesis uses a model where voters are connected via an underlying social network and compute their information about the current state of the election both by observing their neighbors' ballots and an opinion poll announced by a polling agency. This thesis explores the manipulation power of the polling agency for the voting rules plurality and veto and shows that manipulation is para-NP-hard even for very restricted social networks and permitted ballot deviations. In particular, the thesis focuses on distance restrictions for the polling agency in regard to the manipulated opinion poll and for the voters in regard to their deviations.

Finally, this thesis deals with control in judgment aggregation under various preference types of the chair where a chair tries to achieve a better outcome by changing the structure of the aggregation process, e.g., by adding or deleting judges. This thesis shows NP-completeness for most considered problems for the judgment aggregation procedures uniform (constant) premise-based quota rules.

Contents

Chapter 1

Introduction

Computational social choice is an interdisciplinary field that encompasses among others political science, economics, mathematics, and computer science. The aim is to aggregate, e.g, preferences, opinions, or judgments over a set of, e.g., candidates, items, or issues to reach a collective outcome. This process of collective decision making is present in everyday life, be it political elections or choosing a selection of menu items. See the book edited by Brandt et al. (2016) for an overview on the field of computational social choice, and the book edited by Endriss (2017) for an outlook on current research directions.

In *preference aggregation*, the goal is to aggregate voters' preferences over candidates to either elect a winning candidate, elect a winning committee, or create a ranking over the given candidates. Examples include electing a president, electing a parliament, or creating a ranking over politicians of a party stating in which order the party will assign them seats in parliament. Well-known contributors to the field of preference aggregation include, for example, the mathematician Jean-Charles de Borda (1733–1799), the Marquis de Condorcet (1743–1794), a philosopher and mathematician, and economist and Nobel laureate Kenneth Arrow (1921–2017). But preference aggregation is not confined to political elections. Preference aggregation provides models that can be used to describe and design the decision making process for autonomous (computer) systems, and is of particular interest to the field of artificial intelligence and multiagent systems. Computer scientists are not only concerned with combinatorial problems that result from large-scale applications, but also provide tools to analyze, among others, the suitability of an aggregation procedure for the given application. Important contributions include the design of algorithms and the study of computational complexity. The latter deals with quantifying the amount of time or space a problem takes to solve and provides, for example, barriers against strategic attacks on aggregation procedures by proving that attacks need an unfeasible span of time to execute in the worst case. The theory of parameterized complexity is able to give a more fine-grained analysis than the classical complexity setting that focuses on worst-case complexity.

In contrast to preference aggregation, the field of *judgment aggregation* only emerged quite recently. Kornhauser and Sager (1986) note that in a court of three judges that have to decide whether the defendant is guilty, the intuitive approach of aggregating consistent judgments can lead to an inconsistent collective outcome. The reformulation of this paradox by Pettit (2001) marked the starting point of a series of impossibility theorems regarding the incompatibility of desirable axioms for judgment aggregation procedures. Judgment aggregation generalizes preference aggregation and is able to express various scenarios in one framework. However, this expressiveness often comes with the toll of a high computational complexity.

This thesis is organized as follows. Chapter 2 formally introduces complexity theory, preference aggregation, and judgment aggregation, and provides a short overview of related work. The following chapters then illustrate my contribution to the field of voting and judgment aggregation. In multiwinner voting, the goal is to elect a committee of candidates. In this thesis, one focus for multiwinner voting is on the representation of voters' preferences. Chapter 3 deals with a new type of ballot that generalizes existing models for representing voters' preferences and studies the axiomatic properties of corresponding multiwinner rules. The next chapters deal with strategic attacks on aggregation procedures where an attacker tries to reach a more favorable outcome by interfering with the aggregation process. Chapter 4 extends the study of shift bribery—a type of attack where voters are influenced to change their ballots—to iterative voting rules. Chapter 5 proposes distance-based manipulation problems for iterative elections with polls, where an attacker can influence the information voters have about the election, and gives parameterized complexity results. Chapter 6 introduces the concept of control to judgment aggregation where an attacker is able to change the structure of the judgment aggregation process, for example by adding or deleting judges, and presents complexity results for various notions of what constitutes a more favorable outcome. Finally, Chapter 7 concludes the thesis and provides perspectives for future work in my field.

Chapter 2

Background and Related Work

This chapter provides the necessary background for the following chapters. Section 2.1 shortly introduces the concept of computational complexity that is used in all following chapters. Then Section 2.2 gives an overview on preference aggregation. In particular, Section 2.2.1 deals with singlewinner elections, Section 2.2.2 introduces the concept of strategic behavior in collective decision making with a focus on preference aggregation, Section 2.2.3 extends singlewinner to multiwinner preference aggregation, and Section 2.2.4 gives the background for iterative elections. Finally, Section 2.3 introduces judgment aggregation which generalizes preference aggregation.

2.1 Complexity Theory

Computer science is, among others, concerned with solving problems where problems are defined by the input they receive and the question (or task) they answer (or execute). An important factor for a given problem is the question how fast the best algorithm is able to solve it (i.e., the *computational complexity* of the problem (Hartmanis and Stearns, 1965)), and therefore how "hard" the problem is. However, if no fast algorithm for a problem is known, does that mean that a fast algorithm has just not been discovered yet, or does it not exist? While the specific running time is dependent on the algorithm, there are upper and lower bounds for a best algorithm that are due to the nature of the problem. For example, there are problems that can provably be solved in at most linear time (because there is an algorithm with that running time), or that are impossible to solve without using an exponential length of time (and therefore no algorithm can exist that solves the problem faster), where "linear" and "exponential" are seen in relation to the size of the input. This thesis mostly uses worst-case (computational) complexity, where the given upper bound is for the input that takes the most time to solve by a best (i.e., fastest) algorithm for the given problem. Problems can be grouped in so-called *complexity classes* depending on the upper bound for the worst-case input.

This thesis focuses on *decision problems*, i.e., problems that ask a yes/no question. Formally, a decision problem is a set S of words over the alphabet $\{0,1\}$,[1] where S is considered to be the set of inputs for which the answer is yes. An algorithm solves S if—given an input $x \in \{0,1\}^*$—it always correctly decides whether $x \in S$. An algorithm is *deterministic* if for each input, each step the algorithm takes is predetermined by the previous step, which among others implies that executing the algorithm several times for the same input always results in the same outcome. In contrast to a deterministic algorithm, a *nondeterministic* algorithm can take different steps in executions of the same input. In particular, a nondeterministic algorithm solves a decision problem S if for a given input x, at least one series of steps leads to the output 'yes' if $x \in S$, and the algorithm does not output 'yes' for any series of steps if $x \notin S$. The complexity class P contains all decision problems that can be solved in deterministic polynomial time (i.e., can be solved by a deterministic algorithm in polynomial time), whereas NP contains all decision problems that can be solved in nondeterministic polynomial time. Alternatively, one can define NP as the class of decision problems where a solution can be verified in deterministic polynomial time. Obviously, P is a subset of NP. One of the most prominent open problems in the field of complexity theory is the question whether P is equal to NP.

Since complexity classes only give upper bounds, but not lower bounds for the problems contained in them, the actual complexity between the problems in one class can vary. One idea to compare two problems in terms of complexity is to show that one problem is at least as hard as the other problem. This thesis uses many-one reducibility for the comparison of problems.

Definition 2.1 (polynomial-time many-one reducibility)**.** Let A and B be sets of words over $\{0,1\}$. A is *polynomial-time many-one reducible* to a problem B—denoted by $A \leq_{\mathrm{m}}^{\mathrm{P}} B$—if there exists a polynomial-time computable function $f : \{0,1\}^* \to \{0,1\}^*$ so that $x \in A \Leftrightarrow f(x) \in B$ for all $x \in \{0,1\}^*$.

If $A \leq_{\mathrm{m}}^{\mathrm{P}} B$, then A can be solved given the input $x \in \{0,1\}^*$ by solving B with the input $f(x)$. That means that problem B is at least as hard as problem A. Note that $\leq_{\mathrm{m}}^{\mathrm{P}}$ is transitive, so that $A \leq_{\mathrm{m}}^{\mathrm{P}} B$ and $B \leq_{\mathrm{m}}^{\mathrm{P}} C$ implies $A \leq_{\mathrm{m}}^{\mathrm{P}} C$ for all $A, B, C \subseteq \{0,1\}^*$.

One goal of complexity theory is to find problems that can be seen as representatives of

[1] Note that it is always possible to encode words over other, more complex alphabets into words in $\{0,1\}^*$.

their complexity class because they are a member of this complexity class and are at least as hard to solve as every other member of this class.

Definition 2.2 (hardness and completeness)**.** Let B be a set of words over $\{0,1\}^*$ and let $\mathcal{C}$ be a complexity class. B is called $\mathcal{C}$-*hard* if for all $A \in \mathcal{C}$, it holds that $A \leq_{\mathrm{m}}^{\mathrm{P}} B$, and $\mathcal{C}$-*complete* if B is $\mathcal{C}$-hard and $B \in \mathcal{C}$.

Problems in P are called *tractable*, whereas NP-hard problems are called *intractable* under the assumption that P and NP are not equal. Since NP is closed under many-one reducibility, i.e., $A \leq_{\mathrm{m}}^{\mathrm{P}} B$ and $B \in \mathrm{NP}$ implies $A \in \mathrm{NP}$, the existence of a deterministic polynomial-time algorithm for an NP-complete problem would imply deterministic polynomial-time algorithms for all problems in NP. Building on the work by Cook (1971) who proved for the first time that a decision problem—namely, deciding whether a given Boolean formula is satisfiable (SAT for short)—is NP-complete, Karp (1972) proved the NP-completeness for various other natural decision problems. While intractability is a huge disadvantage for many problems arising in natural applications, in this thesis it is mostly seen as a positive aspect in the context of problems that deal with strategic behavior (see Section 2.2.2 for a short introduction). Therefore, the following chapters explore the complexity of the considered problems in detail. See the book by Garey and Johnson (1979) for an introduction to NP-completeness and the book by Arora and Barak (2009) for an extensive introduction to complexity theory in general.

A disadvantage of classical complexity theory is the sole focus on the worst case. A brute-force search approach does not have to be the best way to find a solution for an intractable problem. In practice, the worst case might appear rarely so that algorithms might still solve the problem efficiently in the average case, or there might be efficient heuristics. Furthermore, in many applications the considered instances are bounded in size, for example because there are only a few data sets or the considered graphs are sparse. If such a bound is guaranteed to be small or even fixed, an algorithm might exploit it to find solutions efficiently. The study of problems that include a parameter corresponding to a bound on the input is called *parameterized complexity theory*.

A part of this thesis focuses on parameterized decision problems. In contrast to the classical decision problems, these problems also include a (numerical) parameter that is assumed to be small, e.g., the number of clauses in a Boolean formula for a parameterized version of SAT. Formally, a parameterized decision problem is a set $S \subseteq \{0,1\}^*$ of inputs for which the answer is yes and a parameterization κ that assigns each instance $x \in \{0,1\}^*$ a numerical value k as a parameter. The parameterized complexity class FPT (fixed-parameter tractable) is the parameterized equivalent of P (Downey and Fellows, 1995a), whereas para-NP is the parameterized equivalent of NP (Flum and Grohe, 2003). Note that the membership in a parameterized complexity class is obviously dependent on the parameter, i.e., the same classical decision problem can be a member of different parameterized complexity classes for different choices of parameterization.

Definition 2.3 (FPT, fpt-algorithm, and para-NP)**.** Let $S \subseteq \{0,1\}^*$, let κ be a parameterization $\kappa : \{0,1\}^* \to \mathbb{N}$, let $f : \mathbb{N} \to \mathbb{N}$ be a computable function, and let $g : \{0,1\}^* \to \mathbb{N}$ be a polynomial-time computable function.

1. (S,κ) is in FPT in regard to parameterization κ if there exists a deterministic algorithm that solves S in time $f(\kappa(x)) \cdot g(x)$ for each $x \in \{0,1\}^*$. Such an algorithm is called an *fpt-algorithm*.
2. (S,κ) is in para-NP in regard to parameterization κ if there is a nondeterministic algorithm that solves S in time $f(\kappa(x)) \cdot g(x)$ for each $x \in \{0,1\}^*$.

Analogously to classical complexity theory, a parameterized (polynomial-time many-one) reduction is a tool to compare the hardness of two parameterized problems.

Definition 2.4 (parameterized reduction)**.** Let A, B be sets of words over $\{0,1\}$ and let $\kappa, \kappa' : \{0,1\}^* \to \mathbb{N}$ be parameterizations. (A,κ) is *parameterized polynomial-time many-one reducible* to (B,κ') if there exists an fpt-algorithm $f : \{0,1\}^* \to \{0,1\}^*$ so that $x \in A \Leftrightarrow f(x) \in B$ for all $x \in \{0,1\}^*$, and there exists a computable function g so that $\kappa'(f(x)) \leq g(\kappa(x))$ for all $x \in \{0,1\}^*$.

Note that in contrast to the classical setting, the algorithm that transforms an instance of a problem A into an instance of the problem B is not a polynomial-time, but an fpt-algorithm. That means that the existence of a parameterized reduction does not imply the existence of a polynomial-time reduction and cannot be used to show NP-hardness. The definitions for hardness and completeness as in Definition 2.2 carry over to the parameterized context.

Downey and Fellows (1995a,b) introduce the W-hierarchy including the complexity classes W[1] and W[2]. A parameterized decision problem is in the complexity class W[t] if there exists a parameterized reduction to a certain weighted circuit satisfiability problem.

Note that FPT $= W[0] \subseteq \mathrm{W}[1] \subseteq \mathrm{W}[2] \subseteq \cdots \subseteq$ para-NP and that it is an open question whether FPT is equal to W[1]. Parameterized problems in FPT are called tractable, whereas problems that are W[t]-hard, $t > 0$, or para-NP-hard are called intractable. See the books by Flum and Grohe (2006) and Cygan et al. (2015) for an introduction to parameterized complexity theory.

In contrast to the yes/no questions by decision problems, optimization problems ask for a solution with certain "optimal" properties, e.g., a satisfying truth assignment to a Boolean formula with the maxinum number of 1's possible. Finding an optimal solution for a problem can often be intractable. However, for most applications it is sufficient if the solution is always provably close to the optimal solution. Johnson (1974) starts the comprehensive study of *approximation algorithms* and, among others, gives examples of algorithms that guarantee that the size of a solution is within a factor c the size of the optimal solution while still computing the solution in deterministic polynomial time. See the book by Vazirani (2013) for an overview on approximation.

2.2 Preference Aggregation

In the field of preference aggregation, the goal is to aggregate preferences of agents called *voters* over a set of alternatives (or *candidates*) to elect a winner. One focus in this thesis is the computational complexity of aggregation and of strategic behavior in preference aggregation.

Formally, let $C = \{c_1, \dots, c_m\}$ be a set of candidates and let $N = \{1, \dots, n\}$ be a set of voters, where each voter j has a (private, not further specified) preference over the candidates and submits a ballot v_j. In many applications, these ballots are seen as "sincere", meaning that they reflect the voters' preferences over the candidates in some way. In these contexts, the notions ballot and preference are often used interchangeably in the literature. See Section 2.2.2 and Section 2.2.4 for examples of "insincere" or "strategic" ballots. A list of ballots $P = (v_1, \dots, v_n)$ is called a *profile*.

In this thesis, the following types of ballots will be used. Note that these types can also be used to model preferences (see, e.g., Section 2.2.4), but in general, preferences are assumed to be far more complex. Using *ordinal* ballots, each voter ranks the candidates from best to worst, e.g., each ballot v_j is a linear order $>_j$ over the candidates where $c_i >_j c_k$ can be interpreted as v_j strictly preferring candidate c_i to candidate c_k. Note that a linear order $>$ is

- complete ($a > b$ or $b > a$ holds for all $a, b \in C$, $a \neq b$),
- transitive ($a > b$ and $b > c$ implies $a > c$ for all $a, b, c \in C$), and
- asymmetric (if $a > b$, then $b > a$ does not hold).

There are also models where the ballot is a weak order $\succsim_j$, where $c_i \succsim_j c_k$ denotes that voter j either prefers c_i over c_k, or does not distinguish between them. Weak orders do not have to be asymmetric. Let $pos_j(c)$ denote the rank of candidate c in ballot v_j. For example, for the ballot $c_3 >_2 c_1 >_2 c_2$, it holds that $pos_2(c_3) = 1$.

For *approval ballots*, voters only distinguish between approved and disapproved candidates, and the ballot $v_j \subseteq C$ is a set denoting the approved candidates of voter j. Note that this dichotomous model can also be generalized to allow for more groups of candidates (see Chapter 3 for my contribution to this generalization). Furthermore, for some applications, the number of candidates in an approval ballot might be fixed. The model studied in Chapter 5 uses so-called plurality ballots, where each voter can only approve of exactly one candidate, and veto ballots, where each voter can only disapprove of exactly one candidate and v_j denotes the *disapproved* candidate of voter j.

Finally, voters can assign a numerical value to each candidate. For these *cardinal ballots*, the ballot is a set $v_j = \{(c_i, k_i) \mid c_i \in C, k_i \in \mathbb{Q}, 1 \leq i \leq m\}$. Depending on the application, there might be upper or lower bounds for the values of k_i, or the ballot might be normalized so that the values for the most preferred candidate and for the most disliked candidate are fixed for all voters. See Section 3.4 for voting rules designed for cardinal ballots. In computational social choice, cardinal ballots (or *cardinal utilities*) are mostly used in the field fair division of indivisible goods, where agents (the equivalent to voters) assign values to items (the equivalent to candidates) and the goal is to find an assignment of items to agents that maximizes the utility in some way.

Example 2.5 (ballot types)**.** Let $C = \{a, b, c, d\}$ be a candidate set. Assume that a voter is a big fan of a, likes both b and c equally, and slightly dislikes d. A possible (sincere) ordinal ballot using linear orders is then $v = a > b > c > d$, a weak order $v = a \succ b \sim c \succ d$,[2] a

[2] Here, $a \succ b$ denotes that $a \succsim b$ and not $b \succsim a$, and $b \sim c$ denotes that $b \succsim c$ and $c \succsim b$.

possible approval ballot $v = \{a,b,c\}$, a plurality ballot $v = \{a\}$, a veto ballot $v = \{d\}$, and a possible cardinal ballot $v = \{(a,10),(b,3),(c,3),(d,1)\}$.

2.2.1 Singlewinner Elections

This section deals with singlewinner elections where the goal is to elect a winner among the set of candidates. To determine the outcome of the election, an *aggregation function* is needed to aggregate the voters' ballots (and therefore their underlying preferences). In this chapter, the following two types of functions are used.

Definition 2.6 (voting rule, social welfare function)**.** Let C be a set of candidates, let $N = \{1,\dots,n\}$ be the set of voters, let $\mathcal{B}(C)$ be the set of all possible ballots over C,[3] and let $\mathcal{L}(C)$ be the set of all possible linear orders over C.

- A (singlewinner) *voting rule* $\mathcal{R}$ maps a profile to a subset of candidates called the winners of the election. Formally, $\mathcal{R} : \mathcal{B}(C)^n \to 2^C$.
- A *social welfare function* $\mathcal{R}$ maps a profile to a set of linear orders over the candidates. Formally, $\mathcal{R} : \mathcal{B}(C)^n \to 2^{\mathcal{L}(C)}$.

Note that the above definition allows for an empty set of winners. However, most aggregation functions are designed to always output a winner. If the set of winners of a voting rule is a singleton, the respective candidate is called a unique winner of the election, and else a nonunique winner. In the case where the output of the aggregation function is always a singleton, the function is said to be resolute, and irresolute otherwise. A *tiebreaking scheme* is a function that maps a set to a single member from this set and can be used to make an irresolute aggregation function resolute. See also the book chapter by Baumeister and Rothe (2015) for a detailed introduction to singlewinner preference aggregation.

In this thesis, one important family of voting rules with ordinal ballots are the (positional) scoring rules. Each scoring rule is associated with a family of scoring vectors of the form $\alpha = (\alpha_1,\dots,\alpha_m)$ for each number m of candidates, where $\alpha_i \in \mathbb{N}$, $\alpha_1 \geq \cdots \geq \alpha_m$, and $\alpha_1 > \alpha_m$. For example, the scoring rule *plurality* uses scoring vectors of the form $(1,0,\dots,0)$, *veto* (also called anti-plurality) has scoring vectors of the form $(1,\dots,1,0)$, and *Borda Count* uses scoring vectors of the form $(m-1,m-2,\dots,1,0)$ for m candidates

[3]Note that most aggregation functions are only defined for a single type of ballot.

(see Example 2.9). For each voter, a candidate receives the points as indicated in the scoring vector according to the rank in the respective voter's ordinal ballot. For example, the highest-ranked candidate in a ballot receives α_1 points. These points are added to compute the resulting score for each candidate, and the winners of the election are then the candidates with the highest score (also called plurality winners, veto winners, ..., respectively). Let r_α denote the scoring rule with scoring vector α, let C be a set of candidates, and let P be a profile over C, then

$$r_\alpha(P) = \operatorname*{argmax}_{c \in C} \sum_{v_i \in P} \alpha_{pos_i(c)}.$$

In the context of approval ballots, the rule *approval voting* elects the candidates that appear in the most ballots, i.e., have the highest approval score. See the book by Laslier and Sanver (2010) for a detailed analysis of approval voting. Note that approval voting in the context of plurality ballots (respectively, veto ballots) is also called plurality (respectively, veto). Here, a plurality winner (respectively, veto winner) is the candidate with the most approvals (respectively, least disapprovals). Scoring rules can also be modified to allow for several consecutive rounds. These *iterative (positional) scoring rules* proceed in a (variable or fixed) number of rounds, where the candidates with the lowest scores are eliminated each round. After elimination, the voters' ballots are reduced to only include the candidates that are still participating. The election ends when all remaining candidates have the same score (and are therefore considered the winners of the election) or—in the case of a fixed number of rounds—after the last round (where the candidates with the highest score in the reduced profile after the final elimination are considered the winners). My contribution in Chapter 4 deals with the iterative scoring rules iterated veto and veto with runoff. For *iterated veto* (see Example 2.12), all but the candidates with the highest veto score are eliminated. *Veto with runoff* (see Example 2.10) proceeds in two fixed rounds. In the first round, all candidates that do not have the highest veto score are eliminated, unless this leaves only one candidate, then this candidate and the candidate(s) with the second-highest veto score proceed to the next round. Further iterative scoring rules in this thesis include the Hare rule, iterated plurality, and plurality with runoff (all based on plurality), the Baldwin and the Nanson rule (both based on Borda), and another veto-based rule called Coombs rule, see Chapter 4.

There are various aggregation functions using ordinal ballots that rely on the concept of pairwise comparisons. Given a profile of ordinal ballots, a candidate c wins a pairwise comparison against a candidate d if c is ranked higher than d in more ballots than the other way round. A candidate that wins the pairwise comparisons against all other candidates is called a *Condorcet winner*. Note that a Condorcet winner does not always have to exist. The voting rule *Condorcet* (Condorcet, 1785) elects the Condorcet winner, if such a candidate exists. The voting rule *Copeland* (see Example 2.8) based on a social welfare function proposed by Copeland (1951) elects the candidates with the most wins in pairwise comparisons against other candidates where a tie with another candidate counts as half a win. However, there are several variants depending on how tied candidates are treated. Other examples include counting a tied comparison as a win or as a loss. Define the Kendall tau distance d_τ (Kendall, 1938) between two linear orders $>_i$ and $>_j$ as the number of candidates where both orders disagree, i.e.,

$$d_\tau(>_i, >_j) = |\{(a,b) \mid a,b \in C \wedge a >_i b \wedge b >_j a\}|.$$

The social welfare function *Kemeny rule* (Kemeny, 1959) returns the linear orders that minimize the Kendall tau distance to the profile. Note that the rule Kemeny originally proposed uses weak orders. However, the version with linear orders is more prevalent.

For large-scale applications, an extremely important aspect of a voting rule is the computational complexity of the winner determination. In these cases, it is essential to be able to decide in polynomial time whether a candidate is a winner of the election. Formally, the winner determination decision problem is defined as follows.

$\mathcal{R}$-WINNER-DETERMINATION	
Given:	An election (C,P), where $C = \{c_1, \dots, c_m\}$ is a set of candidates and $P = (v_1, \dots, v_n)$ is a profile, and a designated candidate $w \in C$.
Question:	Is w a winner of the election using voting rule $\mathcal{R}$, i.e., is $w \in \mathcal{R}(P)$?

Note that it is possible in polynomial time to compute the set of winning candidates for all scoring rules including iterative scoring rules, so the respective winner determination problem is in P. However, the winner determination problem for the Kemeny rule is intractable (Bartholdi III et al., 1989b; Hemaspaandra et al., 2005).

Another important aspect of aggregation functions is the properties they satisfy. For example, for the voting rule that always returns the highest-ranked candidate of the first voter as the winner, the winner determination is tractable, but it definitely does a bad job at aggregating the ballots and therefore the underlying preferences of the voters. The aforementioned voting rule is dictatorial, so a *non-dictatorial* aggregation function is certainly desirable. A voting rule $\mathcal{R}$ is *non-imposed* . if for each candidate $c \in C$, there exists at least one profile P in the domain of $\mathcal{R}$ so that $\{c\} = \mathcal{R}(P)$. Based on the aforementioned notion of a Condorcet winner, a voting rule is *Condorcet consistent* if it always returns the Condorcet winner as the unique winner if such a candidate exists. The Condorcet rule is obviously Condorcet consistent, whereas the positional scoring rules are not (Fishburn, 1974). For ordinal ballots, a voting rule $\mathcal{R}$ is *monotone* (Arrow, 1950) if a winner $w \in \mathcal{R}$ remains a winner whenever a voter improves the rank of w in her ballot (and does not change any other relative rankings), and *positive responsive* (May, 1952) if this improvement leads to w being a unique winner. See Chapter 3 for an adaption of these axioms to multiwinner elections as introduced in Section 2.2.3.

Some (intuitively desirable) properties for social welfare functions for ordinal ballots include the following axioms: The axiom *universal domain* demands that the domain of the social welfare function has to consist of all possible profiles over the given candidate set. A social welfare function $\mathcal{R}$ is *Pareto optimal* (see, e.g., the book by Arrow (1963)) if for all candidates $a, b \in C$, $a >_i b$ for all $v_i \in P$ implies that $a > b$ for all $> \in \mathcal{R}(P)$. This axiom ensures that $\mathcal{R}$ actually considers the voters' ballots when returning a winning linear order, by following the voters' lead when they all agree on a relative ranking between two candidates. Furthermore, $\mathcal{R}$ is *independent of irrelevant alternatives* (see, e.g., the book by Arrow (1963)) if for each pair of candidates $a, b \in C$ and each pair of profiles $P = (p_1, \dots, p_n)$ and $Q = (q_1, \dots, q_n)$ that agree on the relative ranking of a and b for all voters $i \in N$,[4] all linear orders in $\mathcal{R}(P)$ and $\mathcal{R}(Q)$ also agree on the relative ranking of a and b. That means that $\mathcal{R}$ should decide on the relative ranking of a and b in the winning linear order solely by considering the relative rankings of them in the voters' ballots and not rankings of the other candidates.

However, not all aforementioned properties are compatible with each other in the sense that they can be satisfied by a single aggregation function. In the context of ordinal ballots, the famous impossibility result by Arrow (1963) states that for at least three candidates, there exists no social welfare function using ordinal ballots that has a universal domain, is Pareto optimal, independent of irrelevant alternatives, and is non-dictatorial, i.e., the result does

[4]This means that either $a >_i b$ in both p_i and q_i, or $b >_i a$ in both p_i and q_i, for all voters $i \in N$.

not correspond to a fixed voter's ballot for each given profile. Note that Arrow's theorem can also be restated using voting rules and adapted notions of the relevant axioms (Taylor, 2005). One way to circumvent Arrow's theorem is to restrict the domain of the considered aggregation function. Black (1958) introduces the concept of single-peaked domains for ordinal ballots. Given an ordering of the candidates π, the corresponding single-peaked domain contains all linear orders $>_i$ where $top(i)\ \pi\ c_j\ \pi\ c_k$ or $c_k\ \ \pi\ c_j\ \pi\ top(i)$ implies that $c_j >_i c_k$ for all $c_j, c_k \in C$, where $top(i)$ denotes the highest ranked candidate in $>_i$. This ordering corresponds to arranging the candidates on an axis, e.g., left to right in political elections. To illustrate this property, each single-peaked order only has one peak when the positions of the candidates in the ballots are plotted on axis π (see Example 2.7). Bartholdi and Trick (1986) show that it is possible in polynomial time to decide whether a given profile is single-peaked, i.e., whether each ballot in the given profile is single-peaked with respect to an ordering.

Another well-studied domain restriction are the single-crossing domains as introduced by Mirrlees (1971) in the context of income taxation, which contain the profiles where, given an ordering π of the voters, there are no voters v_i, v_j, v_k and candidates c, d so that $v_i\ \pi\ v_j\ \pi\ v_k$, but $c >_i d$, $d >_j c$, and $c >_k d$. This corresponds to the idea that the voters are ordered in a way so that for each pair of candidates $c, d \in C$, the voters ranking c higher than d form an interval in the profile. To illustrate this property, draw a line between each occurence of a candidate in the profile. Then for each pair of candidates, the respective lines between them are allowed to only cross once (see Example 2.7). Elkind et al. (2012) and Bredereck et al. (2013) provide polynomial algorithms to decide whether a given profile is single-crossing, and Bredereck et al. also characterize the single-crossing domain by identifying two forbidden substructures of the profiles. Faliszewski et al. (2011) adapt the notion of single-peaked and Elkind and Lackner (2015) adapt the notion of single-crossing to approval ballots, and Elkind et al. (2020b) give a characterization for domains that are single-peaked and single-crossing at the same time.

Example 2.7 (single-peaked and single-crossing profile)**.** Let $C = \{a, b, c, d\}$ be a candidate set and let P be the profile in Figure 2.1a of ordinal ballots over C. Figure 2.1b plots the candidates against their rank in the respective ballots for the candidate ordering π where $a\ \pi\ b\ \pi\ c\ \pi\ d$. Since each plot has exactly one peak, p is single-peaked. Furthermore, Figure 2.1c plots the ballots for the voter ordering π with $v_1\ \pi\ v_2\ \pi\ v_3$ and shows that the lines connecting the candidates cross each other line at most once. Therefore, P is also

$$
\begin{array}{lrl}
P\colon & v_1 = & a > b > c > d \\
 & v_2 = & b > c > d > a \\
 & v_3 = & c > d > b > a
\end{array}
$$

(a) Profile P

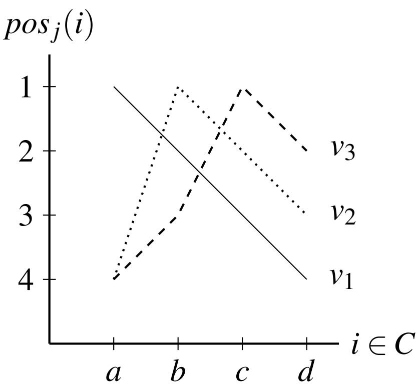

(b) Profile P is single-peaked.

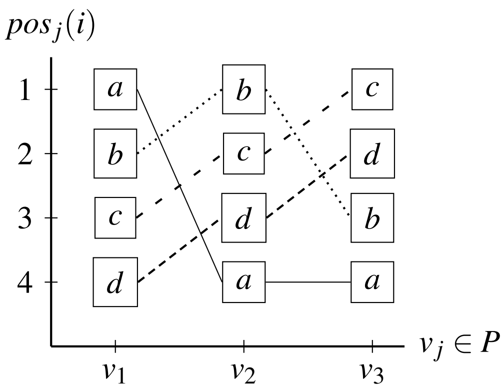

(c) Profile P is single-crossing.

Figure 2.1: The profile P in Example 2.7 is single-peaked and single-crossing.

single-crossing.

Restricted domains not only allow for non-dictatorial aggregation functions that satisfy the axioms in Arrow's theorem (and allow for non-manipulable aggregation functions that satisfy the axioms in the Gibbard-Satterthwaite Theorem, see Section 2.2.2), but can also have an impact on the computational complexity of the winner determination of other aggregation functions. For example, Brandt et al. (2015) show that the highly intractable winner determination of the Dodgson rule[5] and the Kemeny rule become tractable for single-peaked profiles. See the book chapter by Elkind et al. (2017b) for an overview on restricted domains.

[5]Given a profile of linear orders, the *Dodgson rule* returns the candidate(s) for which the number of swaps between adjacent candidates in the ballots needed to make this candidate the Condorcet winner is minimal (Dodgson, 1876). The complexity of the winner determination for the Dodgson rule was studied by Bartholdi III et al. (1989b) and Hemaspaandra et al. (1997).

2.2.2 Strategic Behavior

Unfortunately, not all actors in an election are honest. This section explores strategic behavior in elections where an attacker changes some aspects of the preference aggregation in order to reach a more desirable outcome.

There are several ways for strategic attacks on preference aggregation procedures.

- In *manipulation* (see, e.g., the book chapter by Zwicker (2016)), an agent that participates in the aggregation procedure reports an insincere ballot, i.e., a ballot that does not correspond to her underlying preferences over the candidates, in order to reach a more desired outcome of the election (see Example 2.8). In some problem variants, there can be a group of manipulators that work together to reach a joint objective, or there are (distance) restrictions on which insincere ballots may be reported. This thesis studies a slightly different notion of manipulation in the context of iterative elections as introduced in Section 2.2.4 where an attacker does not report an insincere ballot, but an insincere poll. See Chapter 5 for my contribution to the topic manipulation in iterative elections.

- In *bribery*, an external agent is able to bribe a given number of voters to submit insincere ballots in order to change the election result in his or her favor (see Example 2.9). Variations of the problem include prices for the voters, i.e., the briber has a budget and each voter (or even each change in a ballot) has a (possibly different) price attached. This thesis studies a model called shift bribery as introduced on page 23. See Chapter 4 for my contribution to the topic bribery for iterative scoring rules.

- In *control*, an external agent called the chair does not have any influence on the agents' ballots, but on the structure of the aggregation procedure (see Example 2.10). This can include adding or deleting candidates or voters, or partitioning the voters or candidates in groups where the aggregation takes place for each group separately. This thesis studies control for the related setting of judgment aggregation as introduced in Section 2.3. See the book chapter by Faliszewski and Rothe (2016) for an overview of bribery and control in voting, and see Chapter 6 for my contribution to the study of judge control problems in judgment aggregation.

Note that it is necessary to define beforehand what constitutes a more desirable outcome for an attacker. One way is via a preference (see, e.g., the book by Taylor (2005)), where for a resolute voting rule the new outcome is preferred to the old one if and only if the attacker prefers the new winner to the old winner. In the case of irresolute voting rules, the notion of a "better outcome" is more complicated and requires assumptions about the attacker's preferences over sets of candidates. For an overview over these so-called set extensions see, e.g., the book chapter by Barberà et al. (2004). In the models used in this section, there is a target candidate given. In a *constructive* attack, the attacker wants the target candidate to win the election, whereas in a *destructive* attack, the attacker tries to prevent the target candidate from winning.

Unfortunately, in most cases, voting rules are susceptible to strategic attacks. Gibbard (1973) and Satterthwaite (1975) independently proved the so-called Gibbard-Satterthwaite Theorem stating that there are no resolute, non-dictatorial and non-imposing voting rules for at least three candidates that cannot be successfully manipulated, assuming there are no restrictions on the voters' ballots.[6] This result started research on how manipulation can nevertheless be prevented. One important aspect in this research is the computational complexity of deciding whether a successful strategic attack is possible in the given election. If it turns out that the underlying problem is computationally hard for a voting system, then this might discourage an attacker.

Bartholdi III et al. (1989a) pioneered the approach of studying the computational complexity of manipulation by studying the following decision problem for voting rules using ordinal ballots.

$\mathcal{R}$-MANIPULATION	
Given:	A set of candidates C, a profile $P = (v_1, \ldots, v_n)$ over C, and a designated candidate $c \in C$.
Question:	Is there a ballot v_{n+1} so that $c \in \mathcal{R}((v_1, \ldots, v_n, v_{n+1}))$?

They show that for all voting rules that satisfy certain properties and for a single manipulator, there exists a polynomial-time greedy algorithm that computes a ballot for the manipulator to successfully make a designated candidate the winner of the election (or

[6] Black (1958) (respectively, Saporiti and Tohmé (2006)) show the existence of a non-manipulable rule when the input profile is single-peaked (respectively, single-crossing).

P:	$v_1 =$	$a > b > c > d$
	$v_2 =$	$a > b > c > d$
	$v_3 =$	$b > c > d > a$

(a) Original profile P

P':	$v_1 =$	$a > b > c > d$
	$v_2 =$	$a > b > c > d$
	$v_3 =$	$b > c > d > a$
	$v_4 =$	$b > a > c > d$

(b) Manipulated profile P'

Table 2.1: Original and manipulated profile in Example 2.8.

determines that such a ballot does not exist). Examples of such voting rules include the family of positional scoring rules. In contrast to this result, they show that the computational complexity of manipulating a voting rule called second-order Copeland (a variant of the Copeland rule) is NP-hard, although the winner can be computed in polynomial time. Bartholdi III and Orlin (1991) show that is computationally hard to manipulate the widely used iterative voting rule called STV (Single Transferable Vote, in the singlewinner version also called the Hare rule, see Chapter 4). Conitzer and Sandholm (2003) add a preround to several well-known voting rules to increase the complexity of manipulation. Conitzer et al. (2007) introduce the concept of coalitional manipulation where a group of manipulators cast ballots to achieve a common goal. A slightly different formulation of the manipulation decision problem includes the manipulator's sincere ballot in the given profile and asks whether the manipulator can successfully change her ballot to achieve a better outcome. Instead of allowing a complete change in the manipulator's ballot, one can restrict the new ballot to have at most a certain distance to the original ballot (Obraztsova and Elkind, 2012). See the survey by Faliszewski and Procaccia (2010) and the book chapter by Conitzer and Walsh (2016) for an overview of results for the computational complexity of manipulating voting rules.

Example 2.8 (manipulation)**.** Let $C = \{a,b,c,d\}$ be a candidate set and let $P = (v_1, v_2, v_3)$ be a profile over C defined in Table 2.1a. Consider the voting rule *Copeland* (r_{Cope}) where candidates acquire one point for each win in a pairwise comparison and half a point for each tie, and the candidates with the highest score win the election. Assume that a manipulator wants candidate b to win the election in profile P. Then she can submit the ballot $v_4 = b > a > c > d$. The election with profile $P' = (v_1, v_2, v_3, v_4)$ (see Table 2.1b) proceeds as follows: Candidate a wins the pairwise comparison against c and d and ties with b, so that a has a score of 2.5. Candidate b has a score of 2.5, candidate c a score of 1, and candidate d a score of 0. Therefore, $\{a,b\} = r_{Cope}(P')$ and the manipulation was successful.

Faliszewski et al. (2009a) were the first to study the computational complexity of bribery in elections, including the following decision problem.

$\mathcal{R}$-BRIBERY	
Given:	A set of candidates C, a profile P over C, a designated candidate $c \in C$, and an integer $k \in \mathbb{N}$.
Question:	Is there a profile P' where at most k voters change their ballots compared to P, so that $c \in \mathcal{R}(P')$?

In the priced variant, voters do not have an equal price like in the above scenario, so the problem includes a price function for voters and asks whether in the profile P' only voters changed their ballots whose total price does not exceed the budget k. Hemaspaandra et al. (2007) focus on destructive bribery where the briber wants to prevent the victory of a designated candidate. Dey et al. (2017) introduce a mix of manipulation and bribery where a briber can only bribe voters that profit from this action, i.e., voters who prefer the outcome after the bribery to the original one. As for manipulation, it is possible to restrict the new ballots of the bribed voters, for example by a distance restriction to the sincere ballot. This models scenarios where voters are not willing to deviate too much from their sincere ballot, be it for ethical reasons or the fear of negative consequences for casting an insincere ballot if there is a risk of having their ballot exposed. Dey (2021) studies this problem for bribery. Baumeister et al. (2019) introduce a generalized version of distance bribery where the resulting bribed profile (and not the individual votes) is subject to a distance restriction.

Example 2.9 (bribery)**.** Let $C = \{a,b,c,d\}$ be a candidate set and let P be a profile over C defined in Table 2.2a. Consider the scoring rule Borda Count (r_{Borda}). Recall that under Borda, each candidate receives a score based on the positions in the voters' ballots, and the candidates with the highest score win. Assume that a briber wants candidate c to win in profile P and is able to bribe at most one voter. In P, candidate a receives a score of $|C|-1+|C|-1+|C|-4 = 3+3+0 = 6$, whereas candidate b has a score of 7, candidate c has a score of 4, and candidate d has a score of 1, therefore $\{b\} = r_{Borda}(P)$. If the briber

P:	$v_1 =$	$a > b > c > d$
	$v_2 =$	$a > b > c > d$
	$v_3 =$	$b > c > d > a$

(a) Original profile P

P' :	$v'_1 =$	$c > d > a > b$
	$v_2 =$	$a > b > c > d$
	$v_3 =$	$b > c > d > a$

(b) Bribed profile P'

Table 2.2: The original and the bribed profile in Example 2.9.

convinces voter 1 to change v_1 into $v'_1 = c > d > a > b$, then c wins in the resulting profile P' (see Table 2.2b) with a score of 6 in contrast to the scores of 4, 5, and 3 of candidates a, b, and d, respectively. Therefore, the bribery was successful.

The computational complexity of control by adding, deleting, and partitioning candidates or voters was first studied by Bartholdi III et al. (1992) for plurality and the Condorcet rule. For example, the following decision problem shows the notion of control by adding candidates. Let P_S denote the restricted profile of P where each ballot in P is restricted to the candidates in S.

$\mathcal{R}$-CONTROL-BY-ADDING-CANDIDATES	
Given:	A set of candidates C, a set of additional candidates D, a profile P over $C \cup D$, a designated candidate $c \in C$, and an integer $k \in \mathbb{N}$.
Question:	Is there a set $D' \subseteq D$ of size at most k so that $c \in \mathcal{R}(P_{C \cup D'})$?

Faliszewski et al. (2009b) study constructive control for these control types for the Copeland rule for all variants of ties in the pairwise comparisons. Furthermore, Elkind et al. (2011) provide, among others, results for the constructive control of Borda by adding candidates. Analogous to bribery, Miasko and Faliszewski (2016) introduce prices to control actions.

Example 2.10 (control by adding candidates)**.** Let $C = \{a, b, c\}$ and $D = \{d, e\}$ be candidate sets, let P be the profile over $C \cup D$ defined in Table 2.3a and let P_S be the restricted profile of P where each ballot is restricted to candidates in S.

Consider the iterative voting rule veto with runoff (r_{VRO}). Recall that veto with runoff proceeds in two fixed rounds. After the first round, all candidates that do not have the highest veto score are eliminated, unless that leaves just one candidate, then the candidates

P:	$v_1 =$	$a > e > b > c > d$
	$v_2 =$	$e > a > b > c > d$
	$v_3 =$	$e > b > c > d > a$

(a) The profile P over $C \cup D$.

P_C:	$v_1 =$	$a > b > c$
	$v_2 =$	$a > b > c$
	$v_3 =$	$b > c > a$

$P_{C\setminus\{c\}}$:	$v_1 =$	$a > b$
	$v_2 =$	$a > b$
	$v_3 =$	$b > a$

(b) Election without adding candidates.

$P_{C\cup\{d\}}$:	$v_1 =$	$a > b > c > d$
	$v_2 =$	$a > b > c > d$
	$v_3 =$	$b > c > d > a$

$P_{(C\cup\{d\})\setminus\{a,d\}}$:	$v_1 =$	$b > c$
	$v_2 =$	$b > c$
	$v_3 =$	$b > c$

(c) Election after adding candidate d.

Table 2.3: Profiles in Example 2.10.

that do not have the highest or second-highest veto score are eliminated. The election winners are the candidates with the highest veto score in the second round. Assume that the chair wants candidate b to win and is able to add at most one candidate from D to the candidate set. In the election using profile P_C (see Table 2.3b), b is the unique candidate with the highest veto score, namely a score of 3, whereas a has the second-highest score of 2. Therefore c is eliminated after the first round, resulting in the profile $P_{C\setminus\{c\}}$. Here, a has a veto score of 2 in contrast to the veto score of 1 of b, therefore $b \notin \{a\} = r_{VRO}(P_C)$.

Next, consider the profile $P_{C\cup\{d\}}$ where the chair added candidate d (see Table 2.3c). Both b and c have the highest veto score of 3, so that a and d are eliminated after the first round, resulting in the profile $P_{(C\cup\{d\})\setminus\{a,d\}}$. Here, b has a veto score of 3, whereas c has a veto score of 0, so $b \in \{b\} = r_{VRO}(P_{C\cup\{d\}})$. Therefore, the control action was successful.

Note that for all aforementioned results, an attacker is assumed to have complete information about the election, notably about the exact ballots of all voters. This is a highly unlikely assumption, but widely accepted since it cannot be easier to favorably change the election result in the presence of incomplete information than in an election with complete information. Therefore, the hardness results carry over to the case where the information of the attacker is limited. However, limiting the information an attacker has can make influencing an election difficult even for voting rules where favorably changing the election result is a tractable problem in the complete information model. Conitzer et al. (2011) show for many common voting rules that, given the manipulator only knows partial ordinal ballots of the other voters, it is NP-hard to decide whether a dominating manipulation strategy exists (i.e., a vote that will not be detrimental to the manipulator in comparison to his truthful vote regardless of how the full preferences of the other voters look like). Dey et al. (2018) extend this model and introduce the concept of weak (respectively, strong) manipulation where the manipulator has to be successful in at least one (respectively, in all) extensions of the partial ballots. Inspired by typical limited information a manipulator has in real life, Endriss et al. (2016b) introduce different models for incomplete information (e.g., knowledge about the candidate scores, but not about ballots) and study the effect of limited information on the computational complexity of manipulation for several voting rules including scoring rules.

The aforementioned results assume that strategic attacks are infeasible when the corresponding decision problems are intractable. However, such a computational barrier as protection against strategic attacks might still not be effective in practice. Recall that the notion of NP-hardness relates to the worst-case complexity of a decision problem. Typical real-world scenarios might involve certain restricted preferences in contrast to the unrestricted preferences in theory, or might have certain small parameters such as the number of candidates.

Faliszewski et al. (2011) show that many manipulation and control problems become tractable when the given profile is single-peaked, whereas Brandt et al. (2015) show the same for bribery. Magiera and Faliszewski (2017) show that several control problems become tractable for restricting the input to single-crossing profiles. Procaccia and Rosenschein (2007) show that for most profiles, it suffices to compute the fraction of untruthful and truthful votes to make a decision about the success of manipulation. Conitzer et al. (2007) introduce the problem of weighted coalitional manipulation (i.e., manipulation by a coalition of voters in the presence of weighted voters[7]) and determine for a variety of common voting rules the number of candidates for which weighted coalitional manipula-

[7] In an election with weighted voters, there is a given weight function $\omega : N \to \mathbb{N}$ that maps each voter to a positive integer called the *weight* w_i of the voter i. The aggregation function then uses the modified profile P' as input, where P' denotes the profile where the ballot of each voter i is duplicated $w_i - 1$ times.

tion becomes intractable. Along these lines, Chen et al. (2017) study the parameterized complexity of control by adding or deleting candidates for common voting rules when parameterized by the number of voters, and Bredereck et al. (2015) and Knop et al. (2020) study the complexity of bribery when the number of candidates is small, the latter for a form of bribery called multi-bribery (where voters can be bribed to abstain, swap adjacent candidates in their preference list, or change the number of approved candidates). Bredereck et al. (2014a) identify research challenges in voting regarding parameterized complexity. See the survey by Rothe and Schend (2013) for an overview of typical-case, parameterized, and approximation results for manipulation and control.

There are other, more positive interpretations of the aforementioned strategic attacks in voting. Manipulation can increase the satisfaction of voters with the outcome as described in Section 2.2.4. The concept of destructive bribery is closely related to the *margin of victory*, first defined by Xia (2012). Here, the goal is to compute how robust the election result is by computing the minimal number of voters that would have to report a different preference in order to change the election outcome. Reisch et al. (2014) show that the corresponding decision problem is intractable for several tournament voting rules, and Dey and Narahari (2015) provide sampling algorithms to estimate the margin of victory. Another interpretation of bribery is *campaign management*, where the bribing action can be modeled by real-life election campaigns, e.g., spending money to promote a designated candidate in the ballots. In the constructive case, campaign managers want to promote their candidate, for example by running advertisements and therefore positively influencing the voter's preference about the candidate, which leads to changes in their ballot. Furthermore, campaigns are bound by a budget and therefore need to carefully evaluate which actions yield the desired results. In this context, it makes sense to define a price per change in a voter's ballot rather than a price for each voter. Schlotter et al. (2017) focus on approval-based voting rules and give classical and parameterized complexity results for campaign management, and Elkind and Faliszewski (2010) give approximation algorithms for several voting rules including scoring rules.

One important topic for this thesis in the context of campaign management is the notion of *shift bribery* for ordinal ballots. Shift bribery is a special case of *swap bribery* (Elkind et al. (2009); introduced in another context by Faliszewski et al. (2009b) as microbribery), where a price function $\rho_j : C \times C \to \mathbb{N}$ denotes the price for swapping two adjacent candidates c_i and c_k in the ballot of voter j, and a swap bribery is successful if a given candidate wins the election after a series of swaps in the voters' ballots whose total price do not exceed the given budget. The price functions for shift bribery are defined in Definition 2.11.

Definition 2.11 (shift bribery price function)**.** Let C be a set of candidates, let $c \in C$ be a target candidate, and let v_j be a linear order over C. A *shift bribery price function* is a function $\rho_j : \mathbb{N} \rightarrow \mathbb{N}$ for a voter j where $\rho_j(0) = 0$, $\rho_j(x) \geq \rho_j(y)$ for $x > y$, and $\rho_j(x) = \rho_j(x-1)$ for all $x \geq pos_j(c)$.

Note that $\rho(i)$ indicates the price for shifting the target candidate c forward by i positions in the ballot v_j and that it is not possible to shift the candidate further than the first position in the voter's ballot. For example, $\rho_j(1) = 2$ indicates that changing voter j's ballot from $\ldots > c_k > c_i > \ldots$ to $\ldots > c_i > c_k > \ldots$ costs two units. The shift bribery problem in the constructive case is then defined as follows.

$\mathcal{R}$-SHIFT-BRIBERY	
Given:	A set of candidates C, a profile P of n linear orders over C, a target candidate $c \in C$, a budget $B \in \mathbb{N}$, and a list of price functions $\rho = (\rho_1, \ldots, \rho_n)$.
Question:	Is it possible to shift c in the ballots in P such that the total price does not exceed the budget b and $c \in \mathcal{R}(P')$ for the new profile P'?

In particular, it is only possible to shift the target candidate rather than all candidates as in swap bribery, and the shift can only be forwards in the constructive case. This model assumes that campaign managers do not run so-called smear campaigns where they negatively influence the voters' opinions about opposing candidates. To define the destructive case where the goal is to prevent a target candidate from winning, the shift bribery price functions are interpreted as the price for shifting the designated candidate *backwards*.[8]

Example 2.12 (shift bribery)**.** Let $C = \{a,b,c,d\}$ be the candidate set and let P be the profile over C defined in Table 2.4a. Note that P_S denotes the restricted profile of P where each ballot is restricted to candidates in S. Consider the iterative scoring rule iterated veto (r_{IV}). Recall that each round, the candidates that do not have the highest veto score are

[8] Note that in the destructive case, the condition $\rho_j(x) = \rho_j(x-1)$ for all $x \geq pos_j(c)$ changes to $\rho_j(x) = \rho_j(x-1)$ for all $x \geq m - pos_j(c)$ in order to prevent shifting the candidate too far backwards.

P:	$v_1 =$	$a > b > c > d$
	$v_2 =$	$a > b > c > d$
	$v_3 =$	$b > c > d > a$
$P_{C\setminus\{a,d\}}$:	$v_1 =$	$b > c$
	$v_2 =$	$b > c$
	$v_3 =$	$b > c$

(a) Election with the original profile P.

P':	$v_1 =$	$a > b > c > d$
	$v_2 =$	$a > b > c > d$
	$v'_3 =$	$b > c > a > d$
$P'_{C\setminus\{d\}}$:	$v_1 =$	$a > b > c$
	$v_2 =$	$a > b > c$
	$v'_3 =$	$b > c > a$

(b) Election with the bribed profile P'.

Table 2.4: The election with the original and the bribed profile in Example 2.12.

eliminated. Assume that a briber wants candidate a to win the election and that he has a budget of 2. Further assume that the price function ρ_3 for voter 3 is defined as $\rho_3(0) = 0$, $\rho_3(1) = 1$, and $\rho_3(2) = \rho_3(3) = 3$. In P, the election proceeds as follows. In the first round, candidates b and c have the highest veto score with a score of 3, so that candidates a and d are eliminated, resulting in the profile $P_{C\setminus\{a,d\}}$. Here, candidate b has the highest veto score of 3 whereas c only has a veto score of 0. Therefore, $a \notin \{b\} = r_{IV}(P)$, so the briber has to employ a bribing action.

Recall that the briber can only shift the designated candidate a forward. In the ballots of voters 1 and 2, a is already in the top position and cannot be shifted forward. Shifting a at least two positions forward in ballot v_3 costs 3 units and therefore exceeds the budget. Consider the profile P' in Table 2.4b where a was shifted forward in v_3 by one position. The election then proceeds as follows. In the first round, d has the lowest veto score and is eliminated, resulting in the profile $P'_{C\setminus\{d\}}$. In the second round, a, b, and c have a veto score of 2, 3, and 1, respectively, therefore a and c are eliminated. It holds that $a \notin \{b\} = r_{IV}(P')$, so there does not exist a successful shift bribing action.

Elkind et al. (2020a) consider swap and shift bribery in single-peaked and single-crossing domains, whereas Dorn and Schlotter (2012) analyze the parameterized complexity for swap bribery. Bredereck et al. (2016a) show that different classes of price functions lead to different complexity results for shift bribery when parameterized by the budget for a number of voting rules. They also study the parameterized complexity regarding other parameters. Faliszewski et al. (2021) investigate the approximability of shift bribery for positional scoring rules and Copeland. For the destructive variant of the problem, Kaczmarczyk and Faliszewski (2019) give efficient algorithms for solving the destructive shift bribery problems and identify cases where the complexity between the constructive and destructive problem variant differ. Similar to the margin of victory, Shiryaev et al. (2013) study the robustness of elections by viewing the destructive swap bribery problem as a mean to measure the maximal number of errors that voters can make in their ballot before the election result changes. Here, an error in a voter's ballot is defined as a pair of adjacent candidates that are swapped in the ballot in contrast to the true preference of the voter. In another variant of shift bribery called *combinatorial shift bribery*, it is possible to affect several votes at once with just a single bribery action (Bredereck et al., 2016c). See Chapter 4 for my contribution to the study of the computational complexity of shift bribery.

2.2.3 Multiwinner Elections

In contrast to singlewinner elections as presented in Section 2.2.1, in the context of multiwinner elections, the goal is to elect a group of candidates called a *committee* as the winner of the election. That means that a multiwinner voting rule maps a profile and a desired committee size k to a set that consists of the winning committees, i.e., subsets of candidates of size exactly k.

Definition 2.13 (multiwinner voting rule)**.** Let C be a set of candidates, let P be a profile over C, let $k < |C|$ be a positive integer, and let $W_k = \{S \mid S \subseteq C \wedge |S| = k\}$ be the set of k-sized committees. A *multiwinner voting rule* $\mathcal{F}$ maps a tuple (P, k) to sets in W_k.

In the above definition of a multiwinner voting rule, the committee size is part of the input. However, there are a few scenarios where it makes sense to not fix the size of a winning committee, called the *variable number of winners* setting. For example, this setting can be motivated in the context of a Hall of Fame where each year, only the best of the best candidates (e.g., athletes) are supposed to be elected into the Hall of Fame. However, this number can vary: Some years, there might be several outstanding candidates, whereas in other years, there might even be none that are deemed exceptional enough to be honored. Kilgour (2016) provides an in-depth study of multi-winner voting rules with a variable number of winners in the approval-based setting and explores the axiomatic properties of corresponding voting rules. Faliszewski et al. (2020) study the complexity of these approval-based rules. The rest of this thesis only considers the fixed number of winners model as defined in Definition 2.13, where the desired committee size is predetermined. For an overview on multiwinner elections, see the book chapter by Faliszewski et al. (2017b).

There are several possible directions for the design of multiwinner voting rules. One prominent design concept is individual excellence where the winning committee consists of the best candidates according to certain criteria, e.g., approval scores. One application for these rules is shortlisting, i.e., preselecting a short list of candidates for the next stage (e.g., of a hiring process). An example of an excellence-based multiwinner voting rule is the rule by Debord (1992) that extends Borda Count to multiwinner voting with a fixed committee size. Lackner and Maly (2021) study shortlisting procedures using approval ballots.

A second design concept is diversity, where the goal is to elect candidates in a way that preferably each voter has a candidate she approves of. A possible application is the facility location problem, where several facilities (e.g., hospitals) have to be placed (e.g., in a city) according to certain criteria. If candidates are possible locations and voters approve of locations that are (easily) accessible to them, multiwinner voting rules based on diversity help in finding optimal locations so that each voter has hospital access. See the article by Skowron et al. (2016) for other applications for diversity. A voting rule aiming at diversity was proposed by Chamberlin and Courant (1983). The Chamberlin-Courant rules are a family of multiwinner rules that elect the committee that—given a misrepresentation function—minimizes the misrepresentation of voters. Here, each voter is assigned their most preferred candidate from the committee as a representative.

Definition 2.14 (family of Chamberlin-Courant rules)**.** Let C be a set of candidates, let N be the set of voters, let $\mathcal{B}(C)$ be the set of all possible ballots over C, let $P \in \mathcal{B}(C)$ be a profile, let $k > 0$ be the desired committee size, and let W_k be the set of all k-sized committees over C.

1. A *misrepresentation function* $\mu : N \times C \to \mathbb{N}$ computes how much a candidate misrepresents a voter.
2. For a committee $S \subseteq C$, an *assignment function* $\phi_S : N \to S$ assigns each voter a representative from S. For the Chamberlin-Courant rules, $\phi_S(i) = \operatorname{argmin}_{c \in S} \mu(i, c)$, i.e., voters are represented by the best candidate in a given committee.

3. A *committee misrepresentation function* $d : \mathcal{B}(C) \times W_k \to \mathbb{N}$ computes for a given profile P and k-sized committee S how misrepresented the voters in P are when S is elected.

4. A *Chamberlin-Courant rule* is a multiwinner voting rule f_{CC} so that

$$f_{CC}(P,k) = \underset{S \in W_k}{\text{argmin}}\ d(P,S).$$

Note that the Chamberlin-Courant rules are used for both ordinal and approval ballots and that they can also be stated in terms of satisfaction instead of dissatisfaction (or misrepresentation) of the voters. A widely used misrepresentation function μ for ordinal ballots outputs the inverse Borda score, i.e., $\mu(i,c) = pos_i(c) - 1$ for a voter v_i and a candidate c, whereas the variant for approval ballots returns 0 if the voter approves of the candidate, and 1 otherwise. The choice for the committee misrepresentation function d depends on the application. The utilitarian approach that minimizes the total misrepresentation of each of the n voters with the winning committee (i.e., that uses the committee misrepresentation function $d_{sum}(P,S) = \sum_{1 \leq i \leq n} \mu(i, \phi_S(i)))$ is called the *minisum* principle, whereas the egalitarian approach that minimizes the maximum misrepresentation for a voter (i.e., that uses the committee misrepresentation function $d_{max}(P,S) = \max_{1 \leq i \leq n} \mu(i, \phi_S(i)))$ is called the *minimax* principle (see, e.g., the book chapter by Kilgour et al. (2006)).

Example 2.15 (Chamberlin-Courant rule)**.** Let $C = \{a,b,c\}$ be the candidate set, let $k = 2$ be the committee size, and let P be the profile in Table 2.5a of ordinal ballots over C.

For each committee $W \subset C$, the assignment function Φ_W assigns the respective best-ranked

P:	$v_1 =$	$a > b > c$
	$v_2 =$	$b > a > c$
	$v_3 =$	$b > a > c$
	$v_4 =$	$c > b > a$

(a) Profile P

	$\{a,b\}$	$\{a,c\}$	$\{b,c\}$
v_1 :	$\mu(1,a) = 0$	$\mu(1,a) = 0$	$\mu(1,b) = 1$
v_2 :	$\mu(2,b) = 0$	$\mu(2,a) = 1$	$\mu(2,b) = 0$
v_3 :	$\mu(3,b) = 0$	$\mu(3,a) = 1$	$\mu(3,b) = 0$
v_4 :	$\mu(4,b) = 1$	$\mu(4,c) = 0$	$\mu(4,c) = 0$
$\sum$:	1	2	1

(b) Misrepresentation of each voter for committees of size 2.

Table 2.5: Profile and misrepresentation values in Example 2.15.

candidate from W to each voter. For example, $\Phi_{\{a,b\}}$ assigns candidate a to voter 1, and candidate b to voters 2, 3, and 4. Table 2.5b states the misrepresentation for each voter with a given committee and computes $d(P,W)$ for each committee W of size 2 using the minisum principle. It follows that $f_{CC}(P,2) = \{\{a,b\},\{b,c\}\}$.

Finally, the third important design concept in multiwinner elections is the concept of proportional representation. In many settings, e.g., in the election of a parliament, it is important to ensure that voters' opinions are adequately represented. Skowron (2018) studies the proportionality of common multiwinner rules and ranks them in a hierarchy. Lackner and Skowron (2019) classify multiwinner voting rules by how close they are to the design concepts individual excellence and proportionality. The voting rules by Monroe (1995) aim exactly at fully representing voters. The Monroe rules use the same underlying idea as the Chamberlin-Courant rules, but additionally require that each candidate in a committee represents the same number of voters (barring rounding differences), therefore the assignment function does not always assign a voter's most preferred committee member as their representative.

However, even determining whether there exists a committee with a misrepresentation of at most a given integer is already intractable for both the Chamberlin-Courant and the Monroe rules (Procaccia et al., 2007). Deciding whether a given committee is a winning committee for these voting rules (which is one variant of the *winner determination problem*) is at least as hard as the aforementioned problem. In contrast to the winner determination of the Monroe rules, the winner determination for the Chamberlin-Courant rules become tractable in the presence of profiles from restricted domains (see the article by Betzler et al. (2013) for parameterized complexity results and single-peaked profiles, and see the article by Skowron et al. (2015b) for single-crossing profiles). Peters (2018b) introduces a new technique to compute winning committees in polynomial time for single-peaked profiles for several voting rules whose winner determination is NP-hard, including the Chamberlin-Courant rules. Skowron and Faliszewski (2017) give approximation algorithms for Chamberlin-Courant under approval ballots, whereas Skowron et al. (2015a) prove that the winner determination of both the Chamberlin-Courant and the Monroe rules is inapproximable under linear orders.

To be able to properly decide on a multiwinner rule in a given context, it is important to study the properties of multiwinner rules. Felsenthal and Maoz (1992) adapt four singlewinner rules to elect more than one winner and study how this change affected their properties. Elkind et al. (2017a) introduce a variety of axioms—partly inspired by axioms for singlewinner rules, see Section 2.2.1—and show which of several well-known multiwinner rules satisfy them. Lackner and Skowron (2021) study the properties of several approval-based multiwinner rules. There are also a variety of axioms that deal with the *stability* of multiwinner rules. In the context of approval ballots, one axiom called *justified representation* (Aziz et al., 2017a) and variants thereof have received a lot of attention. A multiwinner rule satisfies justified representation if it always elects a committee so that there does not exist a large enough group of unrepresented voters that have a common approved candidate.

Definition 2.16 (justified representation (Aziz et al., 2017a))**.** Let C be the set of candidates, let $N = \{1,\dots,n\}$ be the set of voters, let $P = (v_1,\dots,v_n)$ be a profile over C using approval ballots, and let $k > 0$ be a committee size.

1. A committee S of size k satisfies *justified representation under P and k* if there does not exist a subset of voters $N' \subseteq N$ with size at least n/k, so that $\bigcap_{i \in N'} v_i \neq \emptyset$ and $v_i \cap S = \emptyset$ for each $i \in N'$.
2. A multiwinner voting rule $\mathcal{F}$ using approval ballots satisfies *justified representation* if for each profile P and each $k > 0$, each winning committee $W \in \mathcal{F}(P,k)$ satisfies justified representation under P and k.

Aziz et al. (2017a) also define a more demanding axiom named extended justified representation (EJR), whereas Sánchez-Fernández et al. (2017) introduce an in-between axiom called proportional justified representation (PJR). Aziz et al. (2018) give complexity results for EJR and PJR. An option to measure the stability of an elected committee for multiwinner rules based on ordinal ballots relies in principle on the notion of a Condorcet winner (see Section 2.2.1). A committee is *Gehrlein-stable* (Gehrlein, 1985; Ratliff, 2003) if there are no candidates outside of the committee that a strict majority of voters prefers to any candidate in the committee. Based on the this concept, Barberà and Coelho (2008) study which multiwinner rules used for shortlisting are stable in the sense that the elected committee contains a (weak) Condorcet winner if one exists. Aziz et al. (2017b) study Gehrlein stability and local stability (the latter based on the Condorcet winning sets introduced by Elkind et al. (2015)) and determine, among other things, the computational complexity of mutiwinner rules that satisfy these types of stability. Gupta et al. (2019) complement these results by studying the parameterized complexity of finding a Gehrlein stable committee. See Chapter 3 for my contribution to the study of properties for multiwinner voting rules.

Bredereck et al. (2018) introduce a model where not only the voters' preferences, but also certain attributes of candidates such as gender or skill level play a part in finding the best committee. In the model by Kagita et al. (2021), voters do not approve candidates, but only such attributes. Izsak et al. (2018) study the case where there are certain relations between candidates that have to be taken into account when finding a good committee. For example, there might be candidates that work exceptionally well (respectively, exceptionally poorly) with each other. Further, Faliszewski et al. (2017a) introduce voting rules where the winning committee is a good compromise between several extremes, for example individual excellency and proportionality. Brill et al. (2019) give approximation algorithms for these type of balanced rules, and Kocot et al. (2019) adapt these rules to guarantee specific score results of committee scoring rules. Note that these *committee scoring rules* defined by Elkind et al. (2017a) are the multiwinner analogues of the (singlewinner) scoring rules introduced in Section 2.2.1 (see the articles by Skowron et al. (2019) and Faliszewski et al. (2019) for an axiomatic characterization of committee scoring rules).

There are a multitude of papers about strategic attacks on multiwinner rules. Lackner and Skowron (2018) define the axioms independence of irrelevant alternatives, monotonicity, and SD-strategyproofness to study the susceptibility of approval-based multiwinner rules. Continuing this work, Peters (2018a) proves an impossibility result that says that no multiwinner rule can satisfy a form of proportionality and a form of strategy-proofness at the same time. Meir et al. (2008) study the computational complexity of manipulation and control for SNTV (Single Non-Transferable Vote; a multiwinner equivalent of Plurality), Bloc (a multiwinner equivalent of k-approval, where k is the desired committee size), approval voting, and cumulative voting (a voting rule that uses cardinal ballots). Yang (2019) also focuses on manipulation and control and, among other things, studies the parameterized complexity of these attacks for approval-based rules. Obraztsova et al. (2013) study the complexity of manipulating scoring rules and give several polynomial-time algorithms. Further, they focus on the role of tie-breaking rules for the success of manipulation. In the context of bribery for approval-based multiwinner rules, Faliszewski et al. (2017c) study the computational complexity and approximability of bribery actions where only a single approval might be added, deleted, or moved within a vote, and Yang (2020) focuses on destructive bribery and gives classical and parameterized complexity results. Bredereck et al. (2016b) study the computational complexity of shift bribery in multiwinner elections for linear orders. Related to this line of research is the question of the robustness of a multiwinner voting rule. Bredereck et al. (2021) study what impact a swap of neighboring candidates has in the linear order of a voter for several rules. Further, they show that it is NP-hard to decide whether the election result can be changed by a given number of swaps.

Baumeister and Dennisen (2015) generalize the concept of the multiwinner rules minisum and minimax from dichotomous ballots to trichotomous ballots[9] and to complete and incomplete linear orders. Then Baumeister et al. (2015a) study the complexity of winner determination and manipulation for these rules, whereas Liu and Guo (2016) focus on the parameterized complexity of winner determination for the generalized minimax rules. Cygan et al. (2018) extend the results for winner determination for the minimax approval rule including approximation and parameterized complexity results. See Chapter 3 for my contribution on the research on minisum and minimax based rules.

2.2.4 Iterative Elections

As seen in Section 2.2.2, all "reasonable" voting rules can be manipulated. This section deals with iterative elections where voters may change their reported ballot in each round, i.e., repeatedly manipulate the election. This corresponds to the idea of *white manipulation*, where manipulation by voters is encouraged as it leads to a better outcome for the voters. Iterative elections can also model the case where voters submit their ballot sequentially, for example in Doodle polls.

Formally, let $N = \{1,\dots,n\}$ be the set of voters, let C be the set of candidates, let $P = (p_1,\dots,p_n)$ be a *preference profile*[10] (usually ordinal preferences) over C, let $P^0 = (a_1,\dots,a_n)$ be an (optional) profile of ballots over C, and let $\mathcal{R}$ be a voting rule. Iterative elections can be modeled as a game (see, e.g., the book chapter by Faliszewski et al. (2016) for an introduction to noncooperative game theory), where the voters are the players and the set of *actions* A_i (also called *strategies*) that each player i can take corresponds to the set of possible ballots over C of a type applicable to $\mathcal{R}$, and are therefore equal

[9] In a trichotomous ballot, the set of candidates is partitioned into three sets corresponding to approved, indifferent, and disapproved candidates.

[10] A *preference profile* is a list of preferences over a candidate set.

for all voters. The game proceeds in turns, where—depending on the model—voters either all act simultaneously or each turn t, a voter j is singled out to take an action so that $a_i^{t-1} = a_i^t$ for all voters $i \in N \setminus \{j\}$. A strategy profile for turn t is then a profile $P^t = (a_1^t, \dots, a_n^t) \in A_1 \times \cdots \times A_n = \mathcal{A}$, where in most applications, P^0 is a profile of initial ballots that are assumed to be truthful, i.e., where the ballots $a_1^0, \dots, a_n^0$ correspond to the preferences $p_1, \dots, p_n$. Each voter further has a utility function $u_j : \mathcal{A} \to \mathbb{R}$ where $u_j(P^t)$ denotes the utility voter j gains from the the strategy profile P^t in turn t. In the context of iterative elections, $u_j(P^t) > u_j(P^s)$ if voter j prefers $\mathcal{R}(P^t)$ to $\mathcal{R}(P^s)$ according to preference p_j. Note that in this context, $\mathcal{R}$ is generally assumed to be resolute, for example by applying a tie-breaking scheme to the outcome. Therefore, preferences over election outcomes can be determined straightforward. Assume that voters only use best-response dynamics to change their ballots i.e., they submit a ballot that will result in the best outcome for them.

Definition 2.17 (best response)**.** Let $N = \{1, \dots, n\}$ be the set of players, let $\mathcal{A} = A_1 \times \cdots \times A_n$ be the set of strategy profiles, and let $u_j : \mathcal{A} \to \mathbb{R}$ be the utility function of player $j \in N$. An action $a_j \in A_j$ is a *best response* to the strategy profile $(a_1, \dots, a_{j-1}, a_{j+1}, \dots, a_n) \in A_1 \times \cdots \times A_{j-1} \times A_{j+1} \times \cdots \times A_n$ if for all $a'_j \in A_j$,

$$u_j(a_1, \dots, a_{j-1}, a_j, a_{j+1}, \dots, a_n) \geq u_j(a_1, \dots, a_{j-1}, a'_j, a_{j+1}, \dots, a_n).$$

Note that—depending on the information model—a voter's computed best response does not have to coincide with the actual best response in that situation since voters may not have access to the current strategy profile. See Chapter 5 for a model where voters can only see the submitted ballots of their neighbors in a social network. Further, voters are assumed to have no memory, i.e., their best response does only depend on the current state of the election, not past states, and they are assumed to be myopic, i.e., they do not have access to the preference profile and do therefore not predict any future deviations by a voter.

In the basic model of Meir et al. (2010), a voter only deviates if he is *pivotal*, i.e., if his best response changes the election outcome. In contrast to this, Obraztsova et al. (2016) study the model where a voter not only changes her ballot when she knows that this change will impact the outcome positively, but also deviates optimistically, i.e., when she believes that other voters will also change their ballot in her favor. The winner is announced as

P:	$p_1 =$	$a > b > c$
	$p_2 =$	$b > c > a$
	$p_3 =$	$c > b > a$

(a) Preference profile P

P^0:	$a_1^0 =$	$\{a\}$
	$a_2^0 =$	$\{b\}$
	$a_3^0 =$	$\{c\}$

(b) Ballot profile P^0

Table 2.6: The preference profile and ballot profile in Example 2.18.

P_a^1:	$a_1^1 =$	$\{a\}$
	$a_2^1 =$	$\{b\}$
	$a_3^1 =$	$\{a\}$

P_b^1:	$a_1^1 =$	$\{a\}$
	$a_2^1 =$	$\{b\}$
	$a_3^1 =$	$\{b\}$

P_c^1:	$a_1^1 =$	$\{a\}$
	$a_2^1 =$	$\{b\}$
	$a_3^1 =$	$\{c\}$

Table 2.7: Possible profiles in the first turn in Example 2.18.

soon as the profile does not change anymore. Iterative elections where voters take turns are path-dependent, i.e., the election winner depends on the order in which voters are allowed to deviate. See the book chapter by Meir (2017) for an overview of iterative elections.

Example 2.18 (iterative elections)**.** Let r_{Plu} be the plurality rule, let $C = \{a, b, c\}$ be a candidate set, let P in Table 2.6a be the preference profile over C, and let P^0 in Table 2.6b be the profile of sincere plurality ballots over C. Assume that ties are broken lexicographically and that voters change their ballot sequentially, starting with voter 3. The plurality winner of P^0 is $r_{Plu}(P^0) = \{a\}$. It holds that $A_1 = A_2 = A_3 = \{\{a\}, \{b\}, \{c\}\}$, so consider the profiles P_a^1, P_b^1, and P_c^1 in Table 2.7, where voter 3 submits the ballot $\{a\}$, $\{b\}$, and $\{c\}$, respectively, and the other voters do not change their ballots. It holds that $r_{Plu}(P_a^1) = r_{Plu}(P_c^1) = \{a\}$ and $r_{Plu}(P_b^1) = \{b\}$, so that

$$u_3(P_b^1) > u_3(P_a^1) = u_3(P_c^1) ,$$

therefore only the action $\{b\}$ is a best response for strategy profile $(\{a\}, \{b\})$. In the profile P_b^1, the best-response of voter 2 and 3 is to not change ballots, whereas voter 1 does not change ballots because she is not pivotal. Therefore the winner of the iterative election is candidate b. Now assume that given profile P^0, it is the turn of voter 2. The only best response for voter 2 is to submit the ballot $\{c\}$, after which no other voter changes ballots.

P^0:	$a_1^0 =$	$\{a\}$
	$a_2^0 =$	$\{b\}$
	$a_3^0 =$	$\{c\}$

(a) Initial profile P^0.

P^1:	$a_1^1 =$	$\{a\}$
	$a_2^1 =$	$\{c\}$
	$a_3^1 =$	$\{b\}$

(b) Profile P^1 after the first turn.

Table 2.8: Profiles after turn 0 and 1 in Example 2.19.

Therefore, in this case the winner of the iterative election is candidate c, illustrating the path-dependency of iterative elections when voting sequentially.

An important topic for iterative elections is convergence. Since voters may change their ballots repeatedly, the resulting profiles can cycle, i.e., there may be a sequence of profiles that repeat over and over again.

Example 2.19 (convergence)**.** Consider the setting from Example 2.18, but this time assume that voters submit their ballots simultaneously. The best response for voter 2 and 3 in profile P^0 in Table 2.8a is to submit the ballots $\{c\}$ and $\{b\}$, respectively, whereas voter 1 does not change ballots. In the resulting profile P^1 in Table 2.8b, the best response for voter 2 and 3 is to submit the ballots $\{b\}$ and $\{c\}$, respectively, whereas voter 1 again does not change ballots, resulting again in profile P^0. Recall that voters do not have memories, therefore this cycle repeats over and over again and the election does not converge.

Meir et al. (2010) initiate the study of conditions and voting rules for which elections are guaranteed to converge, i.e., result in a profile where no voter has any incentive to change their ballot. Such a profile is also called an *equilibrum*. Note that—depending on the deviations allowed for the voters and their ballots—there are different kinds of equilibria, e.g., a Nash equilibrium defined as follows.

Definition 2.20 (Nash equlibrium)**.** Let $N = \{1, \ldots, n\}$ be the set of players and let $\mathcal{A} = A_1 \times \cdots \times A_n$ be the set of strategy profiles. A strategy profile $P = (a_1, \ldots, a_n) \in \mathcal{A}$ is in a *Nash equilibrium* if action a_j is a best response to the strategy profile $(a_1, \ldots, a_{j-1}, a_{j+1}, \ldots, a_n) \in A_1 \times \cdots \times A_{j-1} \times A_{j+1} \times \cdots \times A_n$ for each player $j \in N$.

For example, when voters deviate sequentially, Meir et al. (2010) prove that pluality always converges to a Nash equilibrium when the inital ballot profile is sincere. Reyhani and

Wilson (2012) and Lev and Rosenschein (2016) show that no scoring rule except plurality and veto converges in iterative elections. The latter also study tie-breaking schemes and show that the result holds for all tie-breaking schemes and that even restricting the tie-breaking to a certain scheme does not guarantee convergence. Since iterative elections do not converge for many common voting rules, Grandi et al. (2013) and Obraztsova et al. (2015b) restrict the allowed deviations even more than the best-response dynamics in order to achieve guaranteed convergence. Gourvès et al. (2016) consider a model where voters are embedded in a social network. They introduce *considerate equilibria* to iterative voting, where voters do not selfishly update their ballot, but consider their neighbors in the network. A profile is then a considerate equilibrium when a *coalition* of voters consisting of a clique in a graph cannot deviate without harming themselves or their neighbors.

Iterative voting can improve the quality of the election outcome, i.e., lead to a higher utility of the voter. Brânzei et al. (2013) analyze the price of anarchy in iterative elections, which is defined as the ratio between the quality of the outcome in a Nash equilibrium that can be reached when starting a sequence of best responses from the original profile, and the quality of the outcome in the original profile. Meir et al. (2014) also analyze the quality of the outcome in iterative elections. They conduct experiments in the situation where voters are uncertain about the current state of the election. However, in real-world elections with underlying social networks, this effect of a better outcome is not as pronounced. Tsang and Larson (2016) explain this phenomenon with the fact that many voters are only connected to voters with a similar view to their own which lessens the possibilities for pivotal deviations.

Reijngoud and Endriss (2012) consider opinion polls as a source of information for voters. Fairstein et al. (2019) compare different models for strategic voting in the presence of opinion polls and introduce a heuristic to predict how voters will behave in this scenario. Sina et al. (2015) consider a social network that the voters are embedded in. Voters obtain their information by an opinion poll and by their neighbors in the social network and can act accordingly to this information. They introduce the concept of network control where the chair can introduce new edges to the social network to obtain a desired outcome. Wilczynski (2019) introduces the concept of manipulation by the agency that publishes the opinion poll. Here, the agency publishes an incorrect poll to influence the voters to change their ballots so that a designated candidate wins the election. See Chapter 5 for my contribution to the study of the computational complexity of poll manipulation.

2.3 Judgment Aggregation

This section generalizes the notion of preference aggregation. In the field of judgment aggregation, agents called *judges* have to come to collective decisions in the form of yes/no answers concerning several possibly logically related issues. Formally, let $\Phi = \{\varphi_1, \neg\varphi_1, \dots, \varphi_m, \neg\varphi_m\}$ be the agenda consisting of the φ_i (called *issues*, $1 \leq i \leq m$) to be decided over. Note that Φ is closed under complementation, consists of propositional formulas φ_i over a set of propositional variables built by using the standard boolean connectives, does not contain any doubly-negated formulas, and is assumed to be finite by a large part of the existing literature. Let $N = \{1, \dots, n\}$ be the set of judges. Each judge $i \in N$ has an individual *judgment set* $J_i \subseteq \Phi$ that is required to be complete (i.e., for each $\varphi \in \Phi$, J_i contains φ or its complement),[11] and J_i is required to be consistent (i.e., there exists a truth assignment for the underlying set of propositional variables so that each issue in J_i evaluates to true). Then $\mathcal{J}(\Phi)$ denotes the set of all possible individual judgment sets and $\mathbf{J} = (J_1, \dots, J_n) \in \mathcal{J}(\Phi)^n$ is called a *profile*.

Definition 2.21 (judgment aggregation procedure)**.** Let Φ be an agenda and let $N = \{1, \dots, n\}$ be the set of judges. A *judgment aggregation procedure* $\mathcal{P}$ is a function

$$\mathcal{P} : \mathcal{J}(\Phi)^n \to 2^{2^{\Phi}}$$

that maps a profile $\mathbf{J} = (J_1, \dots, J_n) \in \mathcal{J}(\Phi)^n$ to a set of (possibly incomplete and inconsistent) judgment sets over Φ.

Note that most of the prevalent procedures are resolute, i.e., the collective outcome is always a singleton. The formula-based framework based on Boolean algebra dominates the literature, but Dietrich (2007) shows that the concepts of judgment aggregation can also be employed with different types of logic. Ågotnes et al. (2011) also introduce a new framework for judgment aggregation based on modal logic. Endriss et al. (2016a) compare the formula-based framework with a framework where issues are propositional variables whose logical relations are expressed in an external integrity constraint. This model is a special case of the field of *binary aggregation*, see, e.g., the article by Dokow and Holzman (2010). For a detailed introduction to judgment aggregation, see the book chapters by Baumeister et al. (2015c), List and Puppe (2009), and Endriss (2016).

[11] See Terzopoulou et al. (2018) for a model where the judges' judgment sets are not required to be complete.

	c	b	ℓ
J_1	1	1	1
J_2	0	1	0
J_3	1	0	0
Maj	1	1	0

Table 2.9: Illustration of the doctrinal paradox/discursive dilemma

The famous *doctrinal paradox*—first presented by Kornhauser and Sager (1986) and later generalized to the *discursive dilemma* by Pettit (2001) and List and Pettit (2002)—shows that when deciding each issue majority-wise, the resulting majority outcome might be inconsistent even if all individual judgment sets are consistent. The corresponding rule is called *majority rule* and always outputs a single judgment set where a formula $\varphi \in \Phi$ is contained if and only if more than half of the judges accept it. To illustrate the discursive dilemma, consider the following example by List and Pettit (2002).

Example 2.22 (discursive dilemma)**.** In the discursive dilemma, three judges have to decide whether a contract was valid (c), whether there was a breach of contract (b), and whether the defendant is liable (ℓ), which is only the case if the contract was valid and it was breached. The agenda is $\Phi = \{c, \neg c, b, \neg b, \ell, \neg \ell\}$ where $\ell = c \wedge b$, and the profile $\mathbf{J} \in \mathcal{J}(\Phi)^3$ can be seen in Table 2.9. Note that a 1 means that the corresponding issue was accepted (i.e., a 1 for an issue φ in J_i means that $\varphi \in J_i$ and $\neg\varphi \notin J_i$), whereas a 0 indicates that this issue was rejected (i.e., a 0 for an issue φ in J_i means that $\varphi \notin J_i$ and $\neg\varphi \in J_i$). Even though all judges have consistent judgment sets (as required), taking the majority decision for each issue leads to an inconsistent outcome: A majority of judges accept that the contract was valid, that it was breached, but that the defendant is not liable.

Therefore, the majority rule—arguably the most simple and natural judgment aggregation procedure—has a major drawback, which started the research on judgment aggregation. See the article by Mongin (2012) for a detailed comparison between the doctrinal paradox and the discursive dilemma.

As a way out of this dilemma, Dietrich and Mongin (2010) separate the issues into premises and conclusions and only aggregate the judgments on the premises in the premise-based approach (respectively, the conclusions in the conclusion-based approach). The premise-based procedure then infers the conclusions from the premises. Dietrich and List (2007b) generalize the majority rule by introducing quota rules where an issue is contained in the collective outcome if the number of judges that accept this issue is at least as high as a given respective quota for this issue. However, as with the majority rule, this approach does not guarantee to produce complete and at the same time consistent outcomes, unless it is paired with the premise-based approach (resulting in the premise-based quota rules) and certain assumptions about the agenda are made. Note that the term “uniform” denotes that the quota is equal for all issues. In this thesis, the following definition for uniform premise-based quota rules is used (see also Example 2.27).

Definition 2.23 (uniform premise-based quota rules)**.** Let $\Phi = \Phi_p \cup \Phi_c$ be the agenda partioned into a set of premises Φ_p and a set of conclusions Φ_c, and let $\Phi_p = \Phi_1 \cup \Phi_2$ be partioned into sets Φ_1 and Φ_2 where Φ_2 contains the complements of all $\varphi \in \Phi_1$. The *uniform premise-based quota rule with quota q* (denoted by $UPQR_q$) for a quota $0 \leq q < 1$ maps each profile $\mathbf{J} = (J_1, \ldots, J_n) \in \mathfrak{J}(\Phi)^n$ to the collective outcome $UPQR_q(\mathbf{J})$ containing

- the premises $\varphi \in \Phi_1$ that are contained in more than $n \cdot q$ judgment sets $J \in \mathbf{J}$,
- the premises $\varphi \in \Phi_2$ that are contained in at least $n \cdot (1-q)$ judgment sets $J \in \mathbf{J}$,
- and all conclusions $\varphi \in \Phi_c$ that can be derived from the premises in the collective outcome.

In contrast to this approach, List (2004) uses a slightly modified majority rule in a sequential context where—following a given ordering of the issues in the agenda—the majority quota for a formula only comes into play when the acceptance or rejection of an issue cannot already be inferred by the already determined part of the collective outcome, thus ensuring a complete and consistent outcome. Following the concept of scoring rules in voting (see Section 2.2.1), Dietrich (2014) introduces scoring rules in judgment aggregation that include an equivalent of Borda Count. Other counterparts of voting rules include rules based on minimization studied by Lang et al. (2011). Furthermore, Lang and Slavkovik (2013) investigate how the different judgment aggregation procedures relate to the established voting rules.

The family of distance-based judgment aggregation procedures as introduced by Pigozzi (2006) and Miller and Osherson (2009) consists of procedures that try to minimize the distance between the collective outcome and the individual judgment sets. In a similar approach, Botan et al. (2021) design an egalitarian rule where the difference in satisfaction with the collective outcome between the judges is minimized.

Another direction to circumvent the discursive dilemma is to analyze the agenda. An agenda is called safe for a judgment aggregation procedure $\mathcal{P}$ if $\mathcal{P}$ produces a consistent outcome regardless of the input profile. Several axioms for the agenda such as the median property were introduced to characterize safe agendas. See the article by Endriss et al. (2012) for an overview and a multitude of axioms regarding the safety of agendas. However, Endriss et al. (2012) also show that for all considered axioms and procedures, it is intractable to check whether a given agenda is safe. Endriss et al. (2015) strengthen these results for the majority rule by showing intractability for agenda safety when parameterized by, among others, the maximum formula size or the maximum variable degree, but also show fixed-parameter tractability when parameterizing by the size of the agenda.

There are quite a few properties that judgment aggregation procedures can satisfy. See, e.g., the article by Lang et al. (2017) for a list of judgment aggregation procedures and their properties. Corresponding to Arrow's theorem in preference aggregation (see Section 2.2.1), Dietrich and List (2007a) show that an analogue to this theorem holds for the field of judgment aggregation stating that a judgment aggregation procedure can only fulfill certain (deemed reasonable) properties if it is a dictatorship. This theorem strengthens earlier impossibility results, for example by List and Pettit (2002, 2004). Dietrich and List (2010) further weaken the requirements posed on the judgment aggregation procedure to be dictatorial. Following these results, there are several papers that investigate whether the aforementioned "reasonable properties" are in fact reasonable for a judgment aggregation procedure. A judgment aggregation procedure $\mathcal{P}$ satisfies *neutrality* if for each agenda Φ, each profile $\mathbf{J}$, and each pair of formulas $\varphi, \psi \in \Phi$ where $\varphi \in J \Leftrightarrow \psi \in J$ for all $J \in \mathbf{J}$, it holds that $\varphi \in \mathcal{P}(\mathbf{J})$ if and only if $\psi \in \mathcal{P}(\mathbf{J})$. Intuitively, neutrality requires that any two issues have to be treated equally. Slavkovik (2014) and Terzopoulou and Endriss (2019a, 2020) argue that there are several cases where the axiom of neutrality is too strong, and the latter propose weaker versions of neutrality. List (2003) suggests to implement domain restrictions to circumvent impossibility theorems and introduces the *unidimensional alignment* domain that is the judgment aggregation equivalent to the single-crossing domain in voting.

As in voting, it is important to take the computational complexity of winner determination into account when choosing the best judgment aggregation procedure for the given application. Endriss et al. (2012) define the first problem for winner determination as follows.

$\mathcal{P}$-WINNER-DETERMINATION	
Given:	An agenda Φ, a profile $\mathbf{J} \in \mathcal{J}(\Phi)^n$, and a subset of formulas $S \subseteq \Phi$.
Question:	Is there a $J \in \mathcal{P}(\mathbf{J})$ so that $S \subseteq J$?

Note that using the winner determination problem defined above, a winning judgment set can be computed with a polynomial number of queries: Start with the empty set S. For each $\varphi \in \Phi$, ask whether $S \cup \{\varphi\}$ is part of a judgment set in $\mathcal{P}(\mathbf{J})$, and set $S := S \cup \{\varphi\}$ if the answer is yes. Endriss et al. (2012) then show that the winner determination for quota rules and the premise-based rule is tractable, while the winner determination for the distance-based procedure that returns the judgment sets from $\mathcal{J}(\Phi)$ minimizing the sum of Hamming distances to the profile is intractable. Here, the *Hamming distance* $HD(S,T)$ between two complete and consistent judgment sets S and T is defined as the number of issues $\varphi \in \Phi$ on which both sets differ. De Haan and Slavkovik (2017) expand the aforementioned intractability result by considering several more procedures from the family of the distance-based procedures. Further, they show that winner determination for scoring rules in judgment aggregation is also intractable. Lang and Slavkovik (2014) study the computational complexity of winner determination for several majority-preserving rules, and Endriss and de Haan (2015) show the intractability of winner determination for several analogues of voting rules for judgment aggregation, e.g., analogues of the Kemeny rule defined on page 11 and the Slater rule (see the article by Endriss et al. (2020) for an overview). De Haan (2016) further studies the parameterized complexity of winner determination for the Kemeny procedure and shows that the computational complexity does not coincide in the respective frameworks. For the family of sequential quota rules, Baumeister et al. (2021a) show that winner determination is intractable.

Based on the concept of iterative voting in preference aggregation (see Section 2.2.4 and Chapter 5), Terzopoulou and Endriss (2018) introduce a model for iterative judgment aggregation where judges can update their individual judgment sets. Further, they study the convergence to equilibria.

Analogous to Arrow's theorem, the Gibbard-Satterthwaite Theorem introduced in Section 2.2.2 for voting can also be transferred to judgment aggregation. To this end, Dietrich and List (2007c) introduce the following notion of non-manipulability.

Definition 2.24 (non-manipulability)**.** A resolute judgment aggregation procedure $\mathcal{P}$ is *non-manipulable* if for each agenda Φ, each set of judges $N = \{1, \dots, n\}$, each profile $\mathbf{J} = (J_1, \dots, J_n) \in \mathcal{J}(\Phi)^n$, each formula $\varphi \in \Phi$, and each judge $i \in N$, it holds that if $\varphi \in J_i$, but $\varphi \notin \mathcal{P}(\mathbf{J})$, then $\varphi \notin \mathcal{P}(\mathbf{J}')$ for each $\mathbf{J}' = (J_1, \dots, J_{i-1}, J'_i, J_{i+1}, \dots, J_n)$ where $J'_i \in \mathcal{J}(\Phi)$.

They show that, given a certain agenda type, each resolute, non-imposing[12], non-manipulable judgment aggregation procedure that has no further restrictions on the input profile and always returns a complete and consistent collective outcome is a dictatorship of some judge i in the sense that $\mathcal{P}$ always returns the individual judgment set J_i. Furthermore, they define the following types of preferences over judgment sets to study whether a manipulator prefers a new outcome to the original collective outcome and therefore has an incentive to change her judgment set. Note that these preference types are only defined over complete and consistent judgment sets to simplify the presentation.

Definition 2.25 (preference types)**.** Let Φ be an agenda, and let $J \in \mathcal{J}(\Phi)$ be a judgment set.

1. The set U_J of *unrestricted preferences* contains all possible weak orders $\succsim$ over $\mathcal{J}(\Phi)$.
2. Define the set $TR_J \subseteq U_J$ of *top-respecting J-induced preferences* by
$$TR_J = \{\succsim \in U_J \mid J \succ Y \ \forall Y \in \mathcal{J}(\Phi) \setminus \{J\}\}.$$
3. Define the set $CR_J \subseteq TR_J$ of *closeness-respecting J-induced preferences* by
$$CR_J = \{\succsim \in U_J \mid X \succsim Y \ \forall\, X, Y \in \mathcal{J}(\Phi) \text{ where } X \cap J \supseteq Y \cap J\}.$$
4. Finally, define the Hamming-distance J-induced preference $HD_J \subseteq CR_J$ by
$$HD_J = \{\succsim \in U_J \mid (X \succsim Y \Leftrightarrow HD(X,J) \leq HD(Y,J)) \ \forall X, Y \in \mathcal{J}(\Phi)\}.$$

[12] A judgment aggregation procedure $\mathcal{P}$ is non-imposing if for each agenda Φ and each $\varphi \in \Phi$, there exist two profiles $\mathbf{J_1}, \mathbf{J_2}$ in the domain of $\mathcal{P}$ so that $\varphi \in \mathcal{P}(\mathbf{J_1})$ and $\varphi \notin \mathcal{P}(\mathbf{J_2})$.

		b	c	$b \wedge c$
$\mathbf{J}$:	J_1	1	1	1
	J_2	1	0	0
	J_3	0	1	0
	$UPQR_{1/2}$	1	1	1

(a) The original profile $\mathbf{J} = (J_1, J_2, J_3)$.

		b	c	$b \wedge c$
$\mathbf{J}'$:	J_1	1	1	1
	J_2	1	0	0
	J_3'	0	0	0
	$UPQR_{1/2}$	1	0	0

(b) The manipulated profile $\mathbf{J}' = (J_1, J_2, J_3')$.

Table 2.10: The original and the manipulated profile in Example 2.27.

Define strategy-proofness as follows.

Definition 2.26 (strategy-proofness)**.** A resolute judgment aggregation procedure $\mathcal{P}$ is strategy-proof for a preference type $\mathcal{C}$ if for each agenda Φ, each profile $\mathbf{J} = (J_1, \dots, J_n) \in \mathcal{J}(\Phi)^n$, each judge i, each profile $\mathbf{J}' = (J_1, \dots, J_{i-1}, J_i', J_{i+1}, \dots, J_n) \in \mathcal{J}(\Phi)^n$, and each $\succsim_i \in \mathcal{C}_{J_i}$, it holds that $\mathcal{P}(\mathbf{J}) \succsim_i \mathcal{P}(\mathbf{J}')$.

Analogous to the impossibility theorem for non-manipulability, Dietrich and List (2007c) further prove that, given a certain agenda type, each resolute and non-imposing judgment aggregation procedure that has no further restrictions on the input profile, always returns a complete and consistent collective outcome, and is strategy-proof for preference type $\mathcal{C}$ is a dictatorship of some judge i. However, as Terzopoulou and Endriss (2019b) show, this impossibility relies on the full information model, i.e., the model where a potential manipulator knows the complete profile of judgment sets. When the judgment sets of some judges are unknown, most impossibility results concerning strategy-proofness do not hold anymore.

Example 2.27 (non-manipulability and strategy-proofness)**.** Consider the uniform premise-based quota rule with quota $q = 1/2$ and the agenda and profile from Example 2.22, i.e., $\Phi = \{b, \neg b, c, \neg c, b \wedge c, \neg(b \wedge c)\}$ and the profile $\mathbf{J} \in \mathcal{J}(\Phi)^3$ in Table 2.10a. Define $\Phi_p = \{b, \neg b, c, \neg c\}$ as the premises with $\Phi_1 = \{b, c\}$ and define $\Phi_c = \{b \wedge c, \neg(b \wedge c)\}$ as the conclusions of the agenda. A respective majority accepts b and c, so that $b \wedge c$ is also part of the collective outcome ($UPQR_{1/2}(\mathbf{J}) = \{b, c, b \wedge c\}$). In particular, judge 3 disagrees with the collective outcome on issue $b \wedge c$. However, if she changes her judgment to $J_3' = \{\neg b, \neg c, \neg(b \wedge c)\}$ resulting in profile $\mathbf{J}'$ as seen in Table 2.10b, then the collective outcome $UPQR_{1/2}(\mathbf{J}') = \{b, \neg c, \neg(b \wedge c)\}$ agrees with J_3 on issue $b \wedge c$. Therefore, the uniform premise-based quota rule with quota $q = 1/2$ is manipulable.

Now consider the same scenario and assume that judge 3 has closeness-respecting (J_3-induced) preferences. There are four complete and consistent judgment sets for Φ, namely J_1, J_2, J_3, and J'_3. According to Definition 2.25, she prefers J_3 to all other sets. Furthermore, she prefers J'_3 to J_2 since $J'_3 \cap J_3 = \{b, \neg(b \wedge c)\} \supseteq \{\neg(b \wedge c)\} = J_2 \cap J_3$. The relation between J_1 and both J_2 and J'_3 is not fixed. Therefore,

$$\begin{aligned} CR_{J_3} = \{ & J_3 \succ J_1 \succ J'_3 \succ J_2, \\ & J_3 \succ J'_3 \succ J_1 \succ J_2, \\ & J_3 \succ J'_3 \succ J_2 \succ J_1 \} \end{aligned}$$

Since there is an $\succsim \in CR_{J_3}$ so that $UPQR_{1/2}(\mathbf{J}') \succ UPQR_{1/2}(\mathbf{J})$, it follows that the uniform premise-based quota rule with quota $q = 1/2$ is not strategyproof under closeness-respecting preferences.

Again, as with preference aggregation, there exist a multitude of papers investigating the complexity of strategic attacks in judgment aggregation. Manipulation includes manipulation by a single judge as defined above as well as manipulation by a coalition of judges. Endriss et al. (2012) study the complexity of manipulation for the premise-based procedure where the manipulator has Hamming-distance-induced preferences, whereas Baumeister et al. (2015b) consider the (uniform) premise based (quota) rules for the preference types in Definition 2.25 and for the case where the manipulator wants to include a set D as a subset of the manipulated collective outcome. Coalitional manipulation was first studied by Botan et al. (2016). The concept of bribery is closely related to lobbying (see, e.g., the papers by Christian et al. (2007) and Bredereck et al. (2014b) for lobbying in judgment aggregation) and was introduced to judgment aggregation by Baumeister et al. (2015b). Baumeister et al. (2015b,d) study the complexity of bribery for the (uniform) premise based (quota) rules and the preference types defined above. De Haan (2017) studies manipulation and bribery for the Kemeny procedure. Similar to control in preference aggregation, the concept of control in judgment aggregation can include adding or deleting judges (Baumeister et al., 2012, 2015d), bundling issues (Alon et al., 2013) and bundling judges (Baumeister et al., 2013), adding or deleting issues (Dietrich, 2016), or changing the order of issues in sequential rules (Bredereck et al., 2017). See the book chapter by Baumeister et al. (2017) for an overview of strategic behavior in judgment aggregation and see Chapter 6 for my contribution to the investigation of the complexity of strategic attacks in judgment aggregation.

De Haan (2018) argues that the intractability results for judgment aggregation are due to the used framework being overly expressive and proposes to use more limited languages that yield tractability results and are still able to model certain applications. De Haan and Slavkovik (2019) give an encoding for several procedures and problems in judgment aggregation into answer set programming.

Chapter 3

Minisum and Minimax Committee Election Rules for General Preference Types

This chapter deals with new types of ballots and new corresponding types of multiwinner voting rules, also called committee election rules. The corresponding publication is as follows.

Baumeister, D., Böhnlein, T., Rey, L., Schaudt, O., and Selker, A.-K. (2016). Minisum and minimax committee election rules for general preference types. In *Proceedings of the 22nd European Conference on Artificial Intelligence*, pages 1656–1657. IOS Press. Extended Abstract

Brams et al. (2007) introduce the minimax procedure for electing a committee. The goal is to select a committee that minimizes the maximum Hamming distance to a voter in the given approval-based profile. They also define a corresponding minisum procedure that aims to minimize the sum of Hamming distances to the voters' ballots. Baumeister and Dennisen (2015) modify these procedures to allow trichotomous ballots, complete linear orders, and incomplete linear orders. Alcantud and Laruelle (2014) characterize a new voting rule based on trichotomous ballots.

This chapter further extends these articles by introducing a type of ballot called ℓ-ballot, applying these ballots to minisum and minimax rules, and studying correspondingly modified axiomatic properties. In an ℓ-ballot, voters partition the candidates into ℓ (possibly empty) groups and then rank these groups. This can be seen as a compromise between dichotomous and trichotomous ballots on the one hand where voters lose the ability to express more fine-grained relations between the candidates, and linear orders on the other hand where voters are forced to express a strict preference between each pair of the (possibly large) candidate set. Note that a similar type of ballot was previously introduced by Obraztsova et al. (2015a, 2017). However, they study the ballots in a game-theoretic approach, whereas this chapter deals with the axiomatic properties of minisum and minimax voting rules using these kind of ballots. Furthermore, Balinski and Laraki (2011) obtain

ℓ-ballots in an intermediate stage of their majority judgment procedure where each judge (who correspond to voters in this context) first assigns an integer grade in a fixed interval to each contestant (or candidate) via a grade function, and the procedure then returns the contestant with the best median grade of all judges.

This chapter is organized as follows. Section 3.1 introduces the concept of ℓ-ballots and the corresponding multiwinner rules called ℓ-group rules, and illustrates these concepts in a detailed example. Section 3.2 then modifies axiomatic properties for rules using ordinal ballots to ones for rules using ℓ-ballots, and shows which of these properties are fulfilled by the ℓ-group rules. Next, Section 3.3 studies the complexity of winner determination for the new rules. Section 3.4 generalizes ℓ-ballots to cardinal ballots that allow for two independent values a and b for the dissatisfaction with a candidate being a member or not a member of a committee. Finally, Section 3.5 details my contribution to the findings of this chapter.

3.1 Minisum and Minimax ℓ-group rules

This section introduces a new type of ballot and corresponding multiwinner rules. For a more detailed overview on ballot types see Section 2.2, and for the basics of multiwinner elections see Section 2.2.3. Recall that $N = \{1,\dots,n\}$ is a set of voters, $C = \{c_1,\dots,c_m\}$ is a set of candidates, $k \in \mathbb{N}$ is a committee size, and $\mathcal{F}$ is a multiwinner voting rule.

Definition 3.1 (ℓ-ballots)**.** An *ℓ-ballot* v for an integer $\ell \geq 2$ over candidate set C is a partition of C into ℓ pairwise disjoint, possibly empty sets (called *groups*). Formally, $v = (G_1,\dots,G_\ell)$ so that $G_i \cap G_j = \emptyset$ for $1 \leq i < j \leq \ell$ and $\bigcup_{i=1}^{\ell} G_i = C$.

Note that 2-ballots correspond to approval ballots, whereas n-ballots do not correspond to linear orders or weak orders since groups might be empty. For each group G_j, j is called the *group number* of G_j. Intuitively, a voter prefers candidates with a low group number to ones with a higher one, and is indifferent between candidates with the same group number. Furthermore, let the candidates in C have a fixed ordering $c_1,\dots,c_m$ and let $v(k)$ denote the group number of a candidate c_k in ballot v. Here, $P_\ell = (v_1,\dots,v_n)$ denotes the profile of ℓ-ballots for a fixed $\ell \in \mathbb{N}$ where v_i is the ballot submitted by voter $i \in N$. A *committee election* is a tuple $\mathcal{E} = (C, P_\ell, k)$.

Definition 3.2 (Confirmed and potential committee members)**.** Let $\mathcal{E} = (C, P_\ell, k)$ be a committee election and let $\mathcal{F}$ be a multiwinner rule. A candidate $c \in C$ is a *confirmed committee member for* $\mathcal{E}$ (under $\mathcal{F}$) if $c \in W$ for all $W \in \mathcal{F}(P_\ell, k)$, and a *potential committee member for* $\mathcal{E}$ (under $\mathcal{F}$) if there exists a $W \in \mathcal{F}(P_\ell, k)$ so that $c \in W$.

In this chapter, the focus is on minimizing the dissatisfaction voters have with an elected committee W. Let $F_k(C) = \{S \mid S \subseteq C \wedge |S| = k\}$ be the set of all committees of size k over C. The dissatisfaction of a voter j—associated with ℓ-ballot v_j—with a committee $W \in F_k(C)$ is measured as

$$\delta_\ell(v_j, W) = \sum_{i=1}^{m} |v_j(i) - W(i)|,$$

where $W(i) = 1$ if $c_i \in W$, and $W(i) = \ell$ else. Note that $\delta_\ell(v_j, W)$ can be interpreted as the distance between ℓ-ballot v_j and committee W and generalizes the Hamming distance between two vectors.

Define the following two families of multiwinner rules tailored to ℓ-ballots and minimizing voters' dissatisfaction.

Definition 3.3 (minisum and minimax ℓ-group rules)**.** For each $\ell > 0$, let $\mathcal{E} = (C, P_\ell, k)$ be a committee election and let $F_k(C) = \{S \mid S \subseteq C \wedge |S| = k\}$.

- The *minisum ℓ-group rules* are functions f^ℓ_{sum} that return the committees minimizing the total dissatisfaction of the voters with the elected committees, i.e.,

$$f^\ell_{sum}(P_\ell, k) = \operatorname*{argmin}_{W \in F_k(C)} \sum_{v \in P_\ell} \delta_\ell(v, W).$$

- The *minimax ℓ-group rules* are functions f^ℓ_{max} that return the committees minimizing the dissatisfaction of the respective least satisfied voter with the elected committees, i.e.,

$$f^\ell_{max}(P_\ell, k) = \operatorname*{argmin}_{W \in F_k(C)} \max_{v \in P_\ell} \delta_\ell(v, W).$$

Define $\sum_{v \in P_\ell} v(i)$ as the *minisum score* of a candidate c_i for profile P_ℓ.

Claim 3.4. *For each committee election $\mathcal{E} = (C, P_\ell, k)$, an alternative way to compute the result of $f^\ell_{sum}(P_\ell, k)$ is to calculate the minisum score of each candidate for P_ℓ and then return the committees $W \in F_k(C)$ that contain all candidates with a minisum score lower than s, as well as contain only candidates with a minisum score lower or equal than s, where s denotes the k-lowest minisum score of a candidate.*

Proof. Assume that the claim is not true, i.e., there exists a committee election $\mathcal{E} = (C, P_\ell, k)$ so that either (1) for a winning committee W, there exists a candidate $c \notin W$ that has a lower minisum score than a member of W, or (2) a committee satisfying the requirements in the claim does not win.

Case (1): Let $W \in f^\ell_{sum}(P_\ell, k)$ be a winning committee, so that there exists a $c_i \in W$ and a $c_j \in C \setminus W$ so that $\sum_{v \in P_\ell} v(j) < \sum_{v \in P_\ell} v(i)$. Consider the committee $W' = (W \cup \{c_j\}) \setminus \{c_i\}$.

$$\begin{aligned}
&\sum_{v \in P_\ell} \delta_\ell(v, W) - \sum_{v \in P_\ell} \delta_\ell(v, W') \\
=&\sum_{v \in P_\ell} \big(v(i) - 1\big) + \sum_{v \in P_\ell} \big(\ell - v(j)\big) - \sum_{v \in P_\ell} \big(v(j) - 1\big) - \sum_{v \in P_\ell} \big(\ell - v(i)\big) \\
=&2 \cdot \sum_{v \in P_\ell} v(i) - 2 \cdot \sum_{v \in P_\ell} v(j) > 0 \,,
\end{aligned}$$

i.e., the total dissatisfaction of the voters with W is greater than with W', a contradiction to the fact that W is a winning committee.

Case (2): It follows from case (1) that all winning committees satisfy the requirements in the claim statement. Assume a committee $W \notin f^\ell_{sum}(P_\ell, k)$ satisfying the requirements does not win, whereas a committee $W' \in f^\ell_{sum}(P_\ell, k)$ does. It holds that

$$\sum_{v \in P_\ell} \Big(\sum_{\substack{c_i \in W \\ \wedge c_i \notin W'}} v(i) + \sum_{\substack{c_i \notin W \\ \wedge c_i \in W'}} v(i) \Big) = 2 \cdot |\{c \mid c \in W \wedge c \notin W'\}| \cdot s = \sum_{v \in P_\ell} \Big(\sum_{\substack{c_i \in W' \\ \wedge c_i \notin W}} v(i) + \sum_{\substack{c_i \notin W' \\ \wedge c_i \in W}} v(i) \Big),$$

and therefore

$$\begin{aligned}
&\sum_{v\in P_\ell} \delta_\ell(v,W)\\
&=\sum_{v\in P_\ell}\Bigg(\sum_{\substack{c_i\in W\\ \wedge c_i\in W'}}\big(v(i)-1\big)+\sum_{\substack{c_i\in W\\ \wedge c_i\notin W'}}\big(v(i)-1\big)+\sum_{\substack{c_i\notin W\\ \wedge c_i\in W'}}\big(\ell-v(i)\big)+\sum_{\substack{c_i\notin W\\ \wedge c_i\notin W'}}\big(\ell-v(i)\big)\Bigg)\\
&=\sum_{v\in P_\ell}\Bigg(\sum_{\substack{c_i\in W\\ \wedge c_i\in W'}}\big(v(i)-1\big)+\sum_{\substack{c_i\in W'\\ \wedge c_i\notin W}}\big(v(i)-1\big)+\sum_{\substack{c_i\notin W'\\ \wedge c_i\in W}}\big(\ell-v(i)\big)+\sum_{\substack{c_i\notin W\\ \wedge c_i\notin W'}}\big(\ell-v(i)\big)\Bigg)\\
&=\sum_{v\in P_\ell} \delta_\ell(v,W')
\end{aligned}$$

This is a contradiction to the fact that W' is a winning committee and W is not.

Since both cases end in a contradiction, the claim is true. □

The following example illustrates the minisum and minimax ℓ-group rules.

Example 3.5 (minisum and minimax ℓ-group rules)**.** Let $\mathcal{E}=(C,P_3,3)$ be a committee election where $C=\{a,b,c,d\}$ and P_3 is the profile of 3-ballots over C in Table 3.1.

		G_1	G_2	G_3	
P_3:	$v_1 = ($	$\{a\}$	$\{b,c\}$	$\{d\}$	$)$
	$v_2 = ($	$\{a,b,d\}$	$\{\}$	$\{c\}$	$)$
	$v_3 = ($	$\{c\}$	$\{a,b,d\}$	$\{\}$	$)$

Table 3.1: The profile P_3 in Example 3.5.

Start with the minisum 3-group rule f^3_{sum}. The minisum scores of the candidates are the total of the respective group numbers, see Table 3.2a. For example, candidate a is in G_1 for both voter 1 and 2, and in G_2 for voter 3, adding up to a score of $1+1+2=4$. For $k=3$, the third lowest score is 6. According to Claim 3.4, each winning committee has to contain the candidates with a score lower than 6, namely a and b, and can only contain candidates with a minisum score of at most 6. It follows that $f^3_{sum}(P_3,3)=\{\{a,b,c\},\{a,b,d\}\}$.

Next, turn to the minimax 3-group rule f^3_{max}. Recall that the dissatisfaction of a voter j with a given committee W is the total of $v_j(i)-1$ for each $c_i \in W$ and $\ell - v_j(i)$ for each

	a	b	c	d
v_1 :	1	2	2	3
v_2 :	1	1	3	1
v_3 :	2	2	1	2
$\sum$:	4	5	6	6

(a) Minisum scores

	$\{a,b,c\}$	$\{a,b,d\}$	$\{a,c,d\}$	$\{b,c,d\}$
v_1 :	2	4	4	6
v_2 :	4	0	4	4
v_3 :	3	5	3	3
max :	4	5	4	6

(b) Voters' dissatisfaction with each committee of size 3.

Table 3.2: The minisum scores of each candidate and the dissatisfaction of each voter for the minimax 3-group in Example 3.5.

$c_i \notin W$. See Table 3.2b for the dissatisfaction of each voter for each possible committee of size 3.

For example, voter 1 has a dissatisfaction with committee $\{b,c,d\}$ of $(3-1)+(2-1)+(2-1)+(3-1)=6$. The winning committees are then the committees with the lowest maximum of dissatisfaction. It follows that $f^3_{max}(P_3,3) = \{\{a,b,c\},\{a,c,d\}\}$ with a maximum dissatisfaction of 4.

3.2 Axiomatic Properties

As pointed out in Sections 2.2.1 and 2.2.3, the choice of using a specific (multiwinner) voting rule depends heavily on the axiomatic properties the rules satisfy. This section studies whether the minisum and minimax ℓ-group rules satisfy some well-known properties for singlewinner and/or multiwinner voting rules. Note that the considered properties were originally defined for linear or weak orders, so some properties have to be specifically adapted to ℓ-ballots. In this section, all proofs that I did not contribute are omitted.

The first (fundamental) property is non-imposition, that–similar to the singlewinner version defined on page 12–demands that each committee can win in some election (which also implies that no candidates are incompatible with each other and therefore cannot be part of the same committee). The property homogeneity asks that multiplying the given profile does not change the election result. The multiwinner adaptations of both properties are due to Elkind et al. (2017a).

Definition 3.6. A multiwinner voting rule $\mathcal{F}$ satisfies

- *non-imposition*, if for each set of candidates C and and each committee $W \in F_k(C)$ of size k there is a profile P so that $\mathcal{F}(P,k) = \{W\}$.
- *homogeneity*, if for each committee election $\mathcal{E} = (C,P,k)$ and each $t \in \mathbb{N}$ it holds that $\mathcal{F}(tP,k) = \mathcal{F}(P,k)$, where tP denotes the concatenation of t copies of P.

The families of minisum and minimax ℓ group rules satisfy both properties.

Theorem 3.7. *For each $\ell \in \mathbb{N}$, the minisum and minimax ℓ-group rules satisfy non-imposition and homogeneity.*

The property consistency (also called reinforcement) states that when a profile can be split in two so that the two parts agree on some winning committees, these committees should also win the election for the original profile. Again, this property was adapted from the singlewinner context to multiwinner rules by Elkind et al. (2017a).

Definition 3.8. Let $P_1 + P_2$ denote the concatenation of profiles P_1 and P_2. A multiwinner voting rule $\mathcal{F}$ satisfies *consistency*, if for each pair of profiles P_1, P_2 over the same candidate set C and for each committee size $k \leq |C|$, it holds that $\mathcal{F}(P_1 + P_2, k) = \mathcal{F}(P_1,k) \cap \mathcal{F}(P_2,k)$ whenever $\mathcal{F}(P_1,k) \cap \mathcal{F}(P_2,k) \neq \emptyset$.

Theorem 3.9. *For each $\ell \in \mathbb{N}$, the minisum ℓ-group rule satisfies consistency, while the minimax ℓ-group rule does not.*

A *clone* of a candidate c is another candidate c' that is very similar to c in the voter's eyes and will therefore be ranked close to the original candidate in the ballots. For example, a clone of a job applicant might be another applicant with the same qualifications and age. Ideally, adding a clone of candidate c to the election should not be detrimental to c, but for many contexts, the cloning might lead to a split vote. Consider for example plurality ballots, i.e., approval ballots where exactly one candidate can be approved. Some of the voters approving c might switch to approve the clone c' and thereby costing c votes to the benefit of another candidate. The idea of cloning was first introduced by Tideman (1987) who considers weak orders as ballots where clones of a candidate c are tied to or ranked directly above or below c. In the setting of ℓ-ballots, cloning a candidate c means adding an additional candidate c' to the set of candidates and placing c' in the same group as c for each ballot.

Definition 3.10. A multiwinner voting rule $\mathcal{F}$ is *independent of clones* if for each committee election $\mathcal{E} = (C, P, k)$ and each candidate $c \in C$, a candidate $d \in C$ that is not a potential committee member under $\mathcal{E}$ cannot be a potential committee member in a committee election when cloning c.

Note that the candidates c and d in the above definition do not have to be distinct, i.e., the cloning of candidate c should also not be beneficial to c.

Theorem 3.11. *For each ℓ, the minisum ℓ-group rule satisfies independence of clones, while the minimax ℓ-group rule does not.*

Next, consider variants of monotonicity starting with *committee monotonicity*. This property demands that members of winning committees should remain in winning committees even if the committee size is increased. For example, it is not reasonable to replace candidates on an interview shortlist just because an additional interview slot opened up. The following definition of *committee monotonicity* is due to Elkind et al. (2017a).

Definition 3.12. A multiwinner voting rule $\mathcal{F}$ satisfies *committee monotonicity* if for each committee election $\mathcal{E} = (C, P, k)$, and for each elected committee $W \in \mathcal{F}(P, k)$, the following holds:

- If $k < |C|$ then there exists a $W' \in \mathcal{F}(P, k+1)$ such that $W \subseteq W'$, and
- if $k > 1$ there exists a $W' \in \mathcal{F}(P, k-1)$ such that $W \supseteq W'$.

Theorem 3.13. *For each $\ell \in \mathbb{N}$, the minisum ℓ-group rule satisfies committee monotonicity, whereas the minimax ℓ-group rule does not for each $\ell \geq 2$.*

Proof. For an arbitrary positive integer ℓ, let $\mathcal{E} = (C, P_\ell, k)$ be a committee election. The minisum scores $\sum_{v \in P_\ell} v(j)$ of each candidate $c_j \in C$ are independent of the committee size k and therefore do not change by increasing or decreasing k. Since each minisum ℓ-group rule picks the committees that contain the candidates with the lowest minisum score, it satisfies committee monotonicity.

However, the minimax ℓ-group rule does not satisfy committee monotonicity for each $\ell \geq 2$. To prove this, consider $C = \{c_1, c_2, c_3\}$ and the following profile P_ℓ:

P_ℓ:	$v_1 =$	$(\{c_1\},$	$\{\},$	$\dots,$	$\{\},$	$\{c_2, c_3\})$
	$v_2 =$	$(\{c_2\},$	$\{\},$	$\dots,$	$\{\},$	$\{c_1, c_3\})$

It holds that $f^\ell_{max}(P_\ell, 1) = \{\{c_1\}, \{c_2\}, \{c_3\}\}$, whereas $f^\ell_{max}(P_\ell, 2) = \{\{c_1, c_2\}\}$. This violates the first condition of committee monotonicity since for $k = 1$, the winning committee $\{c_3\}$ is not a subset of any winning committee for $k = 2$. $\square$

The following notions of monotonicity aim at ensuring that additional support for a candidate c does not hurt the prospect of winning for c. In the example of the shortlist for a hiring process, a candidate should not lose an already secured interview spot just because an additional member of the hiring committee announces support for this candidate. The original multiwinner definition for *candidate monotonicity* and *monotonicity* for linear orders is again due to Elkind et al. (2017a) and states that an improvement of a potential committee member c should not be detrimental to c (candidate monotonicity) respectively to the winning committee of which c is a member (monotonicity). In the original (singlewinner) definition of *positive responsiveness* for linear orders that goes back to May (1952), an improvement of a winning candidate c in a ballot should even lead to c being a unique winner of the election (see Definition 3.14 for a restatement for multiwinner voting rules). However, their notion of *improvement*[1] of a candidate is tailored to linear orders and is therefore not suitable for ℓ-ballots. Definition 3.14 below therefore uses the following notion of improvement.

Given a profile $P_\ell = (v_1, \dots, v_n)$ of ℓ-ballots over candidate set C, let P'_ℓ denote the modified profile obtained by improving candidate c_i in some ballot v_j, i.e., shifting c_i into a better group in v_j while the rest of the ballot remains unchanged. Formally,

$$P'_\ell = (v_1, \dots, v_{j-1}, v'_j, v_{j+1}, \dots, v_n) \text{ where } v'_j(i) < v_j(i) \text{ and } v'_j(g) = v_j(g) \text{ for all } c_g \neq c_i.$$

Note that in contrast to the case of linear orders, no candidates are shifted backwards by improving c_j. This is to preserve the spirit of the original notion where the relation between candidates other than c_j remain unchanged. For example, let $C = \{c_1, c_2, c_3, c_4\}$ be the set of candidates and consider the 2-ballots $v = (\{c_3\}, \{c_1, c_2\})$ and $v' = (\{c_1\}, \{c_2, c_3\})$. The

[1] Improving a candidate c by one position in the linear order of voter i means swapping c and the candidate in position $pos_i(c) - 1$. An improvement of $t > 1$ steps can be carried out by improving c by one position t times.

ballot v' can be obtained by swapping the candidates c_1 and c_3, but this has consequences for the relation between c_2 and c_3.

Furthermore, it does not suffice to just require that all relative candidate relations barring the ones with c_i remain the same, but it is necessary to fix the group numbers of all candidates but c_i. In the 6-ballots

$$\begin{array}{rcccccccc} v = & (\{\}, & \{c_4\}, & \{c_2\}, & \{c_1\}, & \{c_3\}, & \{\}) & \text{and} \\ v' = & (\{c_4\}, & \{c_1, c_2\}, & \{\}, & \{\}, & \{\}, & \{c_3\}), & \end{array}$$

the candidate c_1 is improved in v' in relation to v, but c_2 and c_4 improved as well which is counterintuitive to an improvement of c_1. Additionally, the candidate c_3 is even shifted a group back without changing the relation to any other candidate including c_1, which is detrimental to c_3 and again counterintuitive. For a recent adaptation of the notion of monotonicity to approval ballots (respectively, to ballots using weak orders) see the article by Sánchez-Fernández and Fisteus (2019) (respectively, Aziz and Lee (2020)).

Definition 3.14. A multiwinner voting rule $\mathcal{F}$ satisfies

- *candidate monotonicity (resp., positive responsiveness)*, if for each committee election $\mathcal{E} = (C, P, k)$ and each $c \in C$, if c is a potential committee member for $\mathcal{E}$ under $\mathcal{F}$, then it holds that c is a potential (respectively, confirmed) committee member for (C, P', k), where P' is obtained from P by improving c in some ballot v.

- *monotonicity*, if for each committee election $\mathcal{E} = (C, P, k)$, each $W \in \mathcal{F}(P, k)$, and each $c \in W$, it holds that $W \in \mathcal{F}(P', k)$ for all P' that are obtained from P by improving c in some ballot v.

Note that monotonicity as well as positive responsiveness imply candidate monotonicity.

The following theorem shows that the minisum ℓ-group rules satisfy the notions of monotonicity defined above.

Theorem 3.15. *For each $\ell \in \mathbb{N}$, the minisum ℓ-group rule satisfies monotonicity, candidate monotonicity, and positive responsiveness.*

Proof. Start with the the property of monotonicity. Let $\ell \geq 2$ be an arbitrary integer and assume that the minisum ℓ-group rule does not satisfy monotonicity, i.e., there exists a committee election $\mathcal{E} = (C, P_\ell, k)$, a committee $W \in f^\ell_{sum}(P_\ell, k)$, and a candidate $c_i \in W$ so that $W \notin f^\ell_{sum}(P'_\ell, k)$, where P'_ℓ is obtained from P_ℓ by improving c_i in ballot v_j. In particular, that means that in the new committee election, there exists a committee W' that has a lower sum of dissatisfaction of the voters with W' than with the original winning committee W. Formally,

$$\sum_{v \in P_\ell} \delta_\ell(v, W') \geq \sum_{v \in P_\ell} \delta_\ell(v, W) \text{ and} \tag{3.1}$$

$$\sum_{v \in P'_\ell} \delta_\ell(v, W') < \sum_{v \in P'_\ell} \delta_\ell(v, W). \tag{3.2}$$

where Equation (3.1) holds due to $W \in f^\ell_{sum}(P_\ell, k)$ and Equation (3.2) due to $W \notin f^\ell_{sum}(P'_\ell, k)$. Recall that $c_i \in W$ and that only the group number of c_i differs in v_j and v'_j while the placement of all other candidates remains unchanged, and all other ballots in P_ℓ are identical in P'_ℓ. Consequently, for each committee S it holds that

$$\delta_\ell(v'_j, S) - \delta_\ell(v_j, S) = v'_j(i) - v_j(i). \tag{3.3}$$

Equation (3.2) can be transformed as follows.

$$\begin{aligned}
& \sum_{v \in P'_\ell} \delta_\ell(v, W') < \sum_{v \in P'_\ell} \delta_\ell(v, W) \\
\iff \quad & \Big(\sum_{v \in P'_\ell \setminus \{v'_j\}} \delta_\ell(v, W')\Big) + \delta_\ell(v'_j, W') < \Big(\sum_{v \in P'_\ell \setminus \{v'_j\}} \delta_\ell(v, W)\Big) + \delta_\ell(v'_j, W) \\
\iff \quad & \Big(\sum_{v \in P_\ell \setminus \{v_j\}} \delta_\ell(v, W')\Big) + \delta_\ell(v'_j, W') < \Big(\sum_{v \in P_\ell \setminus \{v_j\}} \delta_\ell(v, W)\Big) + \delta_\ell(v'_j, W) \\
\iff \quad & \Big(\sum_{v \in P_\ell} \delta_\ell(v, W')\Big) + \delta_\ell(v'_j, W') - \delta_\ell(v_j, W') < \Big(\sum_{v \in P_\ell} \delta_\ell(v, W)\Big) + \delta_\ell(v'_j, W) - \delta_\ell(v_j, W) \\
\iff \quad & \sum_{v \in P_\ell} \delta_\ell(v, W') - \sum_{v \in P_\ell} \delta_\ell(v, W) < \delta_\ell(v_j, W') - \delta_\ell(v'_j, W') + \underbrace{v'_j(i) - v_j(i)}_{\text{due to Eq. (3.3)}} \\
\iff \quad & \underbrace{\sum_{v \in P_\ell} \delta_\ell(v, W') - \sum_{v \in P_\ell} \delta_\ell(v, W)}_{\geq 0 \text{ due to Eq. (3.1)}} < \underbrace{v_j(i) - v'_j(i)}_{\text{due to Eq. (3.3)}} + v'_j(i) - v_j(i) = 0 \ \lightning
\end{aligned}$$

Due to the contradiction, it follows that the minisum ℓ-group rule satisfies monotonicity for each $\ell > 0$, which also implies that the minisum ℓ-group rule satisfies candidate monotonicity for each $\ell > 0$.

Next, assume that the minisum ℓ-group rule does not satisfy positive responsiveness, i.e., there exists a committee election $\mathcal{E} = (C, P_\ell, k)$, a potential committee member $c \in W$ for some $W \in f_{sum}^{\ell}(P_\ell, k)$, and a committee $W' \in f_{sum}^{\ell}(P'_\ell, k)$ so that $c \notin W'$, where P'_ℓ is obtained from P_ℓ by improving c in ballot v_j. Then $W \in f_{sum}^{\ell}(P'_\ell, k)$ since f_{sum}^{ℓ} satisfies monotonicity as proved above. The inequality

$$\sum_{v \in P_\ell} \delta_\ell(v, W') \overset{(1)}{<} \sum_{v \in P'_\ell} \delta_\ell(v, W') \overset{(2)}{=} \sum_{v \in P'_\ell} \delta_\ell(v, W) \overset{(3)}{<} \sum_{v \in P_\ell} \delta_\ell(v, W)$$

is a contradiction to the fact that W is a winning committee in the original election. Note that (1) is due to the fact that c is not a member of W' and an improvement of c therefore leads to a higher dissatisfaction of voter j with W'. Furthermore, (2) holds since both W and W' are winning committees in the new election, and (3) holds because an improvement of c leads to a lower dissatisfaction of voter j with W. Therefore, the minisum ℓ-group rule satisfies positive responsiveness for all $\ell > 0$. $\square$

Next, turn to the minimax ℓ-group rules.

Theorem 3.16. *For each $\ell \in \mathbb{N}$, the minimax ℓ-group rule satisfies candidate monotonicity. However, for each $\ell \geq 2$, the minimax ℓ-group rule does not satisfy positive responsiveness and monotonicity.*

Proof. First, show that for each $\ell > 0$, the minimax ℓ-group rule satisfies candidate monotonicity. For each committee election $\mathcal{E} = (C, P_\ell, k)$, and each profile P'_ℓ that is obtained by improving a potential committee member $c \in W$, $W \in f_{max}^{\ell}(P_\ell, k)$, in a ballot in P_ℓ, the following holds:

The dissatisfaction of a least satisfied voter does not change or even grows for all committees $W' \in F_k(C)$ where $c \notin W'$, so that the following holds.

$$\max_{v \in P'_\ell} \delta_\ell(v, W) \leq \max_{v \in P_\ell} \delta_\ell(v, W) \leq \max_{v \in P_\ell} \delta_\ell(v, W') \leq \max_{v \in P'_\ell} \delta_\ell(v, W')$$

P_ℓ :	$v_1 =$	$(\{c_1\},$	$\{c_2,c_3\},$	$\{\},$	$\dots)$
	$v_2 =$	$(\{c_3\},$	$\{c_1,c_2\},$	$\{\},$	$\dots)$
P'_ℓ :	$v'_1 =$	$(\{c_1,c_2\},$	$\{c_3\},$	$\{\},$	$\dots)$
	$v_2 =$	$(\{c_3\},$	$\{c_1,c_2\},$	$\{\},$	$\dots)$

Table 3.3: Profiles P_ℓ and P'_ℓ in the positive responsiveness proof of Theorem 3.16.

P_ℓ :	$v_1 =$	$(\{c_1\},$	$\{c_2,c_3\},$	$\{\},$	$\dots)$
	$v_2 =$	$(\{c_2\},$	$\{c_1,c_3\},$	$\{\},$	$\dots)$
	$v_3 =$	$(\{c_3\},$	$\{c_1,c_2\},$	$\{\},$	$\dots)$
P'_ℓ :	$v'_1 =$	$(\{c_1,c_2\},$	$\{c_3\},$	$\{\},$	$\dots)$
	$v_2 =$	$(\{c_2\},$	$\{c_1,c_3\},$	$\{\},$	$\dots)$
	$v_3 =$	$(\{c_3\},$	$\{c_1,c_2\},$	$\{\},$	$\dots)$

Table 3.4: Profiles P_ℓ and P'_ℓ in the monotonicity proof of Theorem 3.16.

Since $c \in W$, this means that there does not exist a committee W' not containing c that has a strictly lower dissatisfaction of the least satisfied voter in the new election than all committees containing c. It follows that c has to be a potential committee member in the new election.

However, such a candidate c does not have to be a confirmed committee member after an improvement, which the following example shows. Consider $C = \{c_1, c_2, c_3\}$ and the profiles P_ℓ and P'_ℓ in Table 3.3, where the latter is obtained by improving c_2 in the ballot v_1. It holds that $f^\ell_{max}(P_\ell, 1) = f^\ell_{max}(P'_\ell, 1) = \{\{c_1\},\{c_2\},\{c_3\}\}$, so c_2 is not a confirmed committee member in the new election despite being a potential committee member in the original election and being improved. This proves that the minimax ℓ-group rule does not satisfy positive responsiveness for each $\ell \geq 2$ because the improvement of a candidate c in a vote that is not the single vote with the minimal maximal dissatisfaction for any committee has no effect on the election result.

Furthermore, the minimax ℓ-group rule does not satisfy monotonicity for each $\ell \geq 2$. Consider for example $C = \{c_1, c_2, c_3\}$ and the profiles P_ℓ and P'_ℓ in Table 3.4, where the latter is obtained by improving c_2 in the ballot v_1. It holds that

$$f^{\ell}_{max}(P_\ell, 2) = \{\{c_1, c_2\}, \{c_2, c_3\}, \{c_1, c_3\}\}.$$

But the improvement of c_2 in ballot v_1 leads to a lower dissatisfaction of the single most dissatisfied voter—namely the first voter—with committee $\{c_2, c_3\}$, while the dissatisfaction of the most dissatisfied voter for the other winning committees does not change. Therefore, $\{c_1, c_2\} \notin f^{\ell}_{max}(P'_\ell, 2) = \{\{c_2, c_3\}\}$, although $c_2 \in \{c_1, c_2\}$ was improved in the original election. □

Elkind et al. (2017a) also consider the weaker notion of *non-crossing monotonicity* where only improvements for $c \in W$ are allowed that do not change the respective relations between all members of W. The results presented above still hold for this variant.

Next, consider Condorcet consistency. In a profile P_ℓ of ℓ-ballots, a candidate c_i beats a candidate c_j in a pairwise comparison if more voters prefer c_i to c_j than the other way round, i.e., $|\{v \in P_\ell \mid v(i) < v(j)\}| > |\{v \in P_\ell \mid v(j) < v(i)\}|$. Recall that a Condorcet winner $c \in C$ beats all other candidate in pairwise comparisons, and that a (singlewinner) voting rule is Condorcet consistent when the Condorcet winner is the unique winner of the election whenever it exists (see the definitions on pages 11 and 12). As in the case of monotonicity, one can define two notions of Condorcet consistency in the multiwinner context, namely one focusing on the candidates and one focusing on the committees. The candidate variant is due to Felsenthal and Maoz (1992) who adapted the singlewinner property of Condorcet consistency to resolute multiwinner voting rules by requiring that a Condorcet winner has to be a member of the winning committee. See Definition 3.17 for a restatement of this definition for irresolute multiwinner voting rules. The committee variant states that the committee Condorcet winner has to be the unique winning committee of the election if it exists, where the *committee Condorcet winner* of size k is a committee $W \in F_k(C)$ for a set of candidates C so that each $c \in W$ beats each $d \in C \setminus W$ in a pairwise comparison (Gehrlein, 1985). Note that a committee Condorcet winner is unique, and that it contains the Condorcet winner if such a candidate exists.

Definition 3.17. A multiwinner voting rule $\mathcal{F}$ satisfies

- *Condorcet consistency*, if for each committee election $\mathcal{E} = (C, P, k)$ where a Condorcet winner $c \in C$ exists, c is a confirmed committee member under $\mathcal{E}$.

- *committee Condorcet consistency*, if for each committee election $\mathcal{E} = (C, P, k)$ where a committee Condorcet winner W exists, it holds that $\{W\} = \mathcal{F}(P, k)$.

Note that committee Condorcet consistency implies Condorcet consistency.

Theorem 3.18. *Neither the minisum nor the minimax ℓ-group rules are Condorcet consistent or committee Condorcet consistent.*

Similar to the property defined for social welfare functions on page 12, a singlewinner voting rule is Pareto optimal (or satisfies the *Pareto criterion*) if only candidates that are not dominated by another candidate can win. A candidate c dominates a candidate d if c is preferred to d by all voters. However, this definition has to be slightly modified for the multiwinner context. For example, a dominated candidate might have to be part of the winning committee when there are not enough not dominated candidates left to fill the committee (as can be the case, e.g., in a *unanimous profile* where all voters prefer a candidate c to every other candidate). See, e.g., the work by Felsenthal and Maoz (1992).

Definition 3.19. Let $\mathcal{E} = (C, P, k)$ be a committee election. A multiwinner voting rule $\mathcal{F}$ satisfies the *Pareto criterion* if the following holds: If a candidate $c_i \in C$ is preferred to a candidate $c_j \in C$ by all voters in P, i.e., $v(i) < v(j)$ for all $v \in P$ in the context of ℓ-ballots, it holds that $c_j \in W$ for some $W \in \mathcal{F}(P, k)$ implies that $c_i \in W$.

Theorem 3.20. *For each $\ell \in \mathbb{N}$, both the minisum and the minimax ℓ-group rules satisfy the Pareto criterion.*

Proof. Le ℓ be an arbitrary positive integer. Assume that the minisum ℓ-group rule (respectively, the minimax ℓ-group rule) does not satisfy the Pareto criterion. Then there exists a committee election $\mathcal{E} = (C, P_\ell, k)$ and candidates $c_i, c_j \in C$, so that $v(i) < v(j)\ \forall v \in P_\ell$, but $c_j \in W$ and $c_i \notin W$ for some $W \in f^\ell_{sum}(P_\ell, k)$ (respectively, $W \in f^\ell_{max}(P_\ell, k)$). A voter's dissatisfaction with a committee only depends on the individual members. The dissatisfaction of each voter v with the committee $W' = (W \cup \{c_i\}) \setminus \{c_j\}$ is strictly less than the dissatisfaction with W:

$$\delta_\ell(v, W') - \delta_\ell(v, W) = v(i) - 1 + \ell - v(j) - \big(v(j) - 1 + \ell - v(i)\big) = 2v(i) - 2v(j) < 0,$$

and therefore

$$\sum_{v \in P_\ell} \delta_\ell(v, W') < \sum_{v \in P_\ell} \delta_\ell(v, W)$$

$$\Big(\text{respectively,} \qquad \max_{v \in P_\ell} \delta_\ell(v, W') < \max_{v \in P_\ell} \delta_\ell(v, W)\Big)$$

This is a contradiction to the fact that W is a winning committee under f^ℓ_{sum} (respectively, f^ℓ_{max}). It follows that both the minisum and the minimax ℓ-group rules satisfy the Pareto criterion for all $\ell > 0$. □

The following properties aim at ensuring proportionality by demanding that a committee W wins the election when a large enough group of voters prefers the members in W to the candidates outside W. They were introduced by Elkind et al. (2017a) and weaken an axiom proposed by Dummett (1984). Here, they are stated in the context of ℓ-ballots. See the axioms of justified representation and Gehrlein stability on page 29 for similar ideas.

Definition 3.21. A multiwinner voting rule $\mathcal{F}$ satisfies

- *solid coalitions* if for each committee election $\mathcal{E} = (C, P_\ell, k)$ with n voters where

  $$|\{v \in P_\ell \mid v(i) < v(j)\ \forall c_j \in C \setminus \{c_i\}\}| \geq n/k$$

  for a candidate $c_i \in C$ implies that c_i is a confirmed committee member,

- *consensus committee* if for each committee election $\mathcal{E} = (C, P_\ell, k)$ with n voters and each $W \in F_k(C)$ so that each voter ranks some member of W higher than all other candidates and each member of W is ranked higher than all other candidates by either $\lfloor n/k \rfloor$ or $\lceil n/k \rceil$ voters, it holds that $\mathcal{F}(P_\ell, k) = \{W\}$, and

- *strong (respectively, weak) unanimity* if for each committee election $\mathcal{E} = (C, P_\ell, k)$, where $v(i) < v(j)$ for all $v \in P_\ell$, $c_i \in W$, and $c_j \in C \setminus W$, it holds that $\{W\} = \mathcal{F}(P_\ell, k)$ (respectively, $W \in \mathcal{F}(P_\ell, k)$).

Note that committee Condorcet consistency implies strong unanimity, which implies weak unanimity.

Theorem 3.22. *For each $\ell \geq 2$, neither the minisum nor the minimax ℓ-group rules satisfy solid coalitions and consensus committee, but they satisfy strong unanimity.*

Table 3.5 summarizes the axiomatic property results in this section. Note that the notion of (candidate) monotonicity in the article Baumeister et al. (2016) corresponds to candidate monotonicity as defined in Definition 3.14. The notion of (crossing and non-crossing) monotonicity was not previously considered for ℓ-group rules.

Property	ℓ-group rules minisum		ℓ-group rules minimax	
Non-imposition	✓	Thm. 3.7	✓	Thm. 3.7
Homogeneity	✓	Thm. 3.7	✓	Thm. 3.7
Consistency	✓	Thm. 3.9	×	Thm. 3.9
Independence of clones	✓	Thm. 3.11	×	Thm. 3.11
Committee monotonicity	✓	Thm. 3.13	×	Thm. 3.13
Candidate monotonicity	✓	Thm. 3.15	✓	Thm. 3.16
Monotonicity	✓	Thm. 3.15	×	Thm. 3.16
Positive responsiveness	✓	Thm. 3.15	×	Thm. 3.16
Condorcet consistency	×	Thm. 3.18	×	Thm. 3.18
Committee Condorcet consistency	×	Thm. 3.18	×	Thm. 3.18
Pareto criterion	✓	Thm. 3.20	✓	Thm. 3.20
Solid coalitions	×	Thm. 3.22	×	Thm. 3.22
Consensus committee	×	Thm. 3.22	×	Thm. 3.22
Unanimity	strong	Thm. 3.22	strong	Thm. 3.22

Table 3.5: Property results for the minisum and minimax ℓ-group rules. The respective results hold for each applicable ℓ.

3.3 Computational Complexity of Winner Determination

This section studies the computational complexity of winner determination for the minisum and minimax ℓ-group rules, starting with the minisum ℓ-group rules. As stated in Claim 3.4, the winning committees for f^{ℓ}_{sum} can be determined by computing the minisum scores of all candidates, which is obviously possible in polynomial time.

Theorem 3.23. *For each $\ell \in \mathbb{N}$, all winning committees for the minisum ℓ-group rule can be computed in polynomial time.*

However, this is not the case for the minimax ℓ-group rules. Recall that 2-ballots correspond to approval ballots. Therefore, f^{2}_{max} corresponds to the original minimax rule introduced by Brams et al. (2007). LeGrand (2004) shows the NP-hardness of the following problem for $\ell = 2$.

MINIMAX-ℓ-SCORE	
Given:	A committee election $\mathcal{E} = (C, P_\ell, k)$, and a nonnegative integer d.
Question:	Is there a committee $W \in F_k(C)$ such that $\max_{v \in P_\ell} \delta(v, W) \leq d$?

Since this problem would be tractable if it were possible to compute a winning committee for f^2_{max} in polynomial time, it follows that the winner determination for f^2_{max} is intractable as well. Furthermore, Misra et al. (2015) show that MINIMAX-2-SCORE is W[2]-hard when parameterized by the size of the committee k. These results can be generalized to all values of ℓ. However, Misra et al. (2015) also show that there exists an fpt-algorithm for MINIMAX-2-SCORE when parameterized by d. The following theorem generalizes the result for each ℓ.

Theorem 3.24. *For each* ℓ*,* MINIMAX-ℓ-SCORE *is in* FPT *when parameterized by* d*.*

3.4 (a,b)-rules

The ℓ-ballots introduced in Section 3.1 are a compromise between ordinal and cardinal ballots in the way that voters can assign a dissatisfaction value to candidates where the values are bound by an underlying ranking of the candidates. There are some articles exploring cardinal preferences,[2] but cardinal ballots and corresponding voting rules are rarely considered in preference aggregation. One example is range voting for singlewinner elections introduced by Smith (2000). Voters assign each candidate a real number from the interval $[-1, 1]$ where a higher number implies a higher satisfaction with the respective candidate, and range voting then elects the candidate(s) with the highest sum of satisfaction.

[2] Ballester and Rey-Biel (2006) study a model where voters have cardinal preferences and have to translate these to the ballots allowed for the respective voting rules. In particular, they focus on approval voting which uses approval ballots and study whether the optimal strategy for voters with cardinal preferences is to submit sincere ballots, i.e., ballots correspond to their underlying preferences in a certain way. Procaccia and Rosenschein (2006) also focus on cardinal preferences and define a notion of "distortion" that occurs when translating cardinal preferences into linear orders.

This section introduces (a,b)-ballots that generalize the concept of ℓ-ballots to fully cardinal ballots. In contrast to range voting, voters assign not only a dissatisfaction value a for the case that a candidate is part of a committe, but also a separate, independent value b of dissatisfaction with a candidate not being a member of a committee. This models for example the case when candidates have different attributes and voters have a ranking for each attribute. A possible attribute for a candidate is, e.g., the party membership. Two independent dissatisfaction values are able to model cases where the conflicting preferences over candidates' attributes do not correspond to an underlying weak order.

Definition 3.25 ((a,b)-ballots)**.** An (a,b)-ballot v_j over candidate set C is a set

$$v_j = \{(c_i, a_i^j, b_i^j) \mid c_i \in C,\ a_i^j, b_i^j \in \mathbb{Q},\ 1 \leq i \leq |C|\}.$$

Then voter j is said to strictly prefer candidate c_1 to candidate c_2 if and only if

$$(a_1^j < a_2^j \text{ and } b_1^j \geq b_2^j) \text{ or } (a_1^j \leq a_2^j \text{ and } b_1^j > b_2^j),$$

and voter j is said to be indifferent between c_1 and c_2 if and only if $a_1^j = a_2^j$ and $b_1^j = b_2^j$. In the remaining cases, it remains unknown which candidate the voter prefers. Note that an ℓ-ballot is a special case of an (a,b)-ballot over the same candidate set C where $a_i^j = v_j(i) - 1$ and $b_i^j = \ell - v_j(i)$ for each candidate $c_i \in C$, but any (a,b)-ballot where the sum of both values is a constant can be interpreted as an ℓ-ballot for a corresponding ℓ.

Recall that $F_k(C) = \{S \mid S \subseteq C \wedge |S| = k\}$ is the set of all committees of size k over C. The dissatisfaction of a voter j—associated with (a,b)-ballot v_j—with a committee $W \in F_k(C)$ is measured as

$$\delta_{(a,b)}(v_j, W) = \sum_{c_i \in W} a_i^j + \sum_{c_i \notin W} b_i^j.$$

Analogous to the minisum and minimax ℓ-group rules, define the following multiwinner rules using (a,b)-ballots:

Definition 3.26 (minisum and minimax (a,b)-rules)**.** Let $\mathcal{E} = (C, P_{(a,b)}, k)$ be a committee election with a profile $P_{(a,b)}$ of (a,b)-ballots over C and let $F_k(C) = \{S \mid S \subseteq C \wedge |S| = k\}$ be the set of committees of size k.

- The *minisum* (a,b)*-rule* is a function $f_{sum}^{(a,b)}$ that returns the committees minimizing the total dissatisfaction of the voters with the elected committees, i.e.,

$$f_{sum}^{(a,b)}(P_{(a,b)},k) = \operatorname*{argmin}_{W \in F_k(C)} \sum_{v \in P_{(a,b)}} \delta_{(a,b)}(v,W).$$

- The *minimax* (a,b)*-rule* is a function $f_{max}^{(a,b)}$ that returns the committees minimizing the dissatisfaction of the respective least satisfied voter with the elected committees, i.e.,

$$f_{max}^{(a,b)}(P_{(a,b)},k) = \operatorname*{argmin}_{W \in F_k(C)} \max_{v \in P_{(a,b)}} \delta_{(a,b)}(v,W).$$

The following example illustrates (a,b)-ballots and the corresponding multiwinner rules.

Example 3.27. (a,b)-rules Let $\mathcal{E} = (C, P_{(a,b)}, 2)$ be a committee election where $C = \{c_1, c_2, c_3\}$ and $P_{(a,b)} = (v_1, v_2)$ is a profile of (a,b)-ballots over C with

$$\begin{aligned} v_1 =& \{(c_1,3,0),(c_2,2,2),(c_3,1,2)\}, \\ v_2 =& \{(c_1,1,2),(c_2,4,3),(c_3,0,4)\}. \end{aligned}$$

Voter 1 strictly prefers c_3 to both c_1 and c_2, since $(a_3^1 < a_1^1 \wedge b_3^1 \geq b_1^1)$ and $(a_3^1 < a_2^1 \wedge b_3^1 \geq b_2^1)$, and also strictly prefers c_2 to c_1. Voter 2 strictly prefers c_3 to c_2 and strictly prefers c_1 to c_2, but the ballot does not allow to draw conclusions about the voter's preference over c_1 and c_3.

Now turn to the minisum and the minimax (a,b)-rules. Recall that the dissatisfaction of a voter j with a committee W is the total of a_i^j for all $c_i \in W$ and b_i^j for all $c_i \notin W$. Table 3.6 shows the dissatisfaction of the voters with each committee of size 2. Therefore, $f_{sum}^{(a,b)}(P_{(a,b)},2) = \{\{c_2,c_3\}\}$ and $f_{max}^{(a,b)}(P_{(a,b)},2) = \{\{c_1,c_3\},\{c_2,c_3\}\}$.

The properties presented in Section 3.2 are defined for ballots that correspond to ordinal ballots. Note that in the restricted model where (a,b)-ballots have to correspond to an underlying weak order, the results in Section 3.2 also hold for the minisum and minimax (a,b)-rules. However, the results for winner determination differ slightly from the results presented in Section 3.3. My coauthors Böhnlein and Schaudt (personal communication, 2016) prove that while a winning committee for the (a,b)-minisum-rule can be computed

	$\{c_1, c_2\}$	$\{c_1, c_3\}$	$\{c_2, c_3\}$
v_1 :	$3+2+2=7$	$3+2+1=6$	$0+2+1=3$
v_2 :	$1+4+4=9$	$1+3+0=4$	$2+4+0=6$
Σ :	16	10	9
max :	9	6	6

Table 3.6: Dissatisfaction of voters with each committee of size 2 in Example 3.27.

in polynomial time, the auxiliary problem MINIMAX (a,b)-SCORE—the (a,b)-version of the problem defined in Section 3.3—is W[2]-hard when parameterized by d and k. However, when the (a,b)-ballots are restricted in a way that they correspond to ℓ-ballots after some normalization (where the value of ℓ may be different for each ballot), they show that Theorem 3.24 also holds for the (a,b)-minimax rule.

Unfortunately, the expressiveness of the presented model demands a great deal of the voters since the necessity to assign two values to each candidate becomes more and more infeasible as the number of candidates grows. Depending on the application, it might therefore be reasonable to restrict the possible values of a and b by, e.g., providing lower and upper bounds, fixing the total sum of a voter's a-values and of a voter's b-values, or by using ℓ-ballots.

3.5 My Contribution

In joint work with my coauthors, I developed the models of ℓ-ballots and (a,b)-ballots and the corresponding rules and modified the properties in Section 3.2—when necessary—for the context of ℓ-ballots. Furthermore, I contributed the proofs to Claim 3.4 and Theorems 3.13, 3.15, 3.16, 3.20, and 3.23. The writing of the article Baumeister et al. (2016) was done jointly with my coauthors.

Chapter 4

Complexity of Shift Bribery for Iterative Voting Rules

This chapter deals with shift bribery for iterative (positional) scoring rules as defined on page 23 and on page 10, respectively. The article Maushagen et al. (2021) was submitted to the *Annals of Mathematics and Artificial Intelligence* and is based on a preliminary conference version (Maushagen et al., 2018b). An earlier version was also presented at ISAIM'18 (Maushagen et al., 2018a).

Maushagen, C., Neveling, M., Rothe, J., and Selker, A.-K. (2021). Complexity of shift bribery for iterative voting rules. Submitted to *Annals of Mathematics and Artificial Intelligence*

Summary

My coauthors and I study the computational complexity of shift bribery for the following iterative scoring rules.

Hare (defined, e.g., in the textbook by Taylor (2005)) eliminates the candidates with the lowest plurality score.

Coombs (defined, e.g., by Levin and Nalebuff (1995)) eliminates the candidates with the lowest veto score.

Baldwin eliminates the candidates with the lowest Borda score (Baldwin, 1926).

Nanson eliminates all candidates with lower than average Borda score (Nanson, 1882).

Iterated plurality (defined, e.g., in the textbook by Taylor (2005)) eliminates the candidates that do not have the highest plurality score.

Iterated veto eliminates the candidates that do not have the highest veto score (see Example 2.12 on page 23).

Plurality with runoff (again defined, e.g., by Taylor (2005)) proceeds in two rounds. In the first round, all candidates that do not have the highest plurality score are eliminated, unless there is a unique plurality winner, then all candidates that do not have the highest or second-highest score are eliminated.

Veto with runoff is the veto variant of plurality with runoff (see Example 2.10 on page 19).

We show that shift bribery is NP-complete for all considered voting rules for both the constructive and the destructive variants. Our results hold for both the unique and the nonunique winner model: In the constructive case of the *unique winner model*, the attacker wants to make the target candidate the unique winner of the election, whereas in the *nonunique winner model*, the attack is successful when the target candidate is a member of the set of winning candidates. Further, as an example we state modified proofs for constructive shift bribery for the rules Hare and plurality with runoff that show that the computational complexity does not change when we allow the attacker to exploit nonmonotonicity. We conjecture that shift bribery for all our considered nonmonotonic rules does not become tractable in this setting.

My Contribution

The writing of the article Maushagen et al. (2021) was done jointly with my co-authors. I was responsible for Section 6 (the complexity results for iterated veto and veto with run-off including the proofs) and Example 2.

Chapter 5

Manipulation of Opinion Polls to Influence Iterative Elections

This chapter deals with the complexity of manipulation of opinion polls by a polling agency in the context of iterative elections. See Section 2.2.4 for a short overview on iterative elections and opinion polls. The corresponding publication is as follows.

Baumeister, D., Selker, A.-K., and Wilczynski, A. (2020b). Manipulation of opinion polls to influence iterative elections. In *Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems*, pages 132–140. IFAAMAS

Note that a version including the omitted proofs was published in the non-archival proceedings of COMSOC'21 (Baumeister et al., 2021b).

Summary

In iterative elections, voters repeatedly update their ballots to achieve a better outcome for them. In the model we use, the necessary information to compute a best response stems from an underlying social network where voters can see the ballots (and updates) of their neighbors, and from an opinion poll announced by a polling agency. Following the work by Wilczynski (2019), we study the manipulative power of the polling agency in iterative elections. First, we introduce a best-response variant for the voting rule veto that significantly differs from the already known best-response definition for plurality. This is due to the fact that under veto, changing who to veto might directly benefit a voter's most despised candidate since that candidate loses a veto from this voter. Second, as an addition to the already known poll manipulation problem without a parameter, we introduce two distance-restricted variants of poll manipulation, i.e., a poll-restricted variant that includes an upper bound on the distance between the sincere and the manipulated poll, and a voter-restricted variant where voters do not vote for candidates that need more than a given maximum number of swaps in their preferences to become their most preferred

candidates for plurality (respectively, their most despised candidate for veto). Third, we consider destructive manipulation in the context of poll manipulation.

To this end, we study the (parameterized) computational complexity of all three problem variants, for constructive and destructive manipulation, and for the voting rules plurality and veto, among others parameterized by the length of the longest path in the social network and by the allowed deviation distance. We show that all problems are NP-hard even in an acyclic social network, and for the constructive poll-restricted poll manipulation under plurality even when the social network does not contain any edges. Manipulation under veto is tractable for all considered problem variants, whereas for plurality we are only able to prove tractability in the destructive case and only for the unrestricted and the poll-restricted variant. The other three cases remain open.

Furthermore, we design efficient heuristics to compare both voting rules and come to the conclusion that—in our setting—manipulation is more successful under the veto rule and that destructive manipulation is more successful than constructive manipulation.

My Contribution

The writing of the article Baumeister et al. (2020b) was done jointly with my coauthors. I defined the poll-restricted and voter-restricted problem variants and contributed the complete results in Section 3, Theorem 4.1, and Proposition 4.2.

Chapter 6

Complexity of Control in Judgment Aggregation for Uniform Premise-Based Quota Rules

This chapter deals with the complexity of influencing the judgment aggregation procedures called uniform (constant) premise-based quota rules by adding, deleting, replacing, and bundling judges. See Section 2.3 for a short introduction to judgment aggregation.

The article Baumeister et al. (2020a) extends several conference versions, including an article I coauthored (Baumeister et al., 2015d).

Baumeister, D., Erdélyi, G., Erdélyi, O. J., Rothe, J., and Selker, A.-K. (2020a). Complexity of control in judgment aggregation for uniform premise-based quota rules. *Journal of Computer and System Sciences*, 112:13–33

Summary

In the article Baumeister et al. (2020a), my coauthors and I introduce the notions of adding, deleting, replacing and bundling judges in judgment aggregation. These concepts can, among others, be found in international arbitration procedures.

There are several ways to measure the success of a control action employed by the chair. Here, we ask whether the chair can achieve to include a given (possibly incomplete) judgment set called *desired set* in the new outcome. We call this problem *exact control by control type* $\mathcal{C}$. Further, we use the preference types introduced by Dietrich and List (2007c)—defined in Definition 2.25 in this thesis and adapted to incomplete desired sets by Baumeister et al. (2015b)—to identify better outcomes for the chair. This approach is based on the concept in preference aggregation where an attacker's preference over the candidates is part of the input. However, note that it is not feasible in judgment aggregation to state an explicit preference over different judgment sets since such a preference is exponential in the size of the agenda. Therefore, we assume that the chair's preference over the outcomes belongs to the set of one of the following four preference types: For the set of

- *unrestricted preferences*, we only know that the chair does not differentiate between two different outcomes that both include the desired set;
- *top-respecting preferences*, we additionally know that the chair prefers outcomes that include the desired set over outcomes that do not;
- *closeness-respecting preferences*, it holds that outcomes that contain a subset X of the desired set are preferred to outcomes that only contain a strict subset of X (and no other issues from the desired set);
- *Hamming-distance* preferences we know that the chair prefers outcomes that include more issues from the desired set over outcomes with fewer.

We obtain the following results for the uniform (constant) premise-based quota rules,[1] where we impose certain restrictions on the agenda's premises and conclusions to obtain complete and consistent versions of these judgment aggregation procedures. For the uniform *constant* premise-based quota rules, possible and necessary control by adding, deleting, and replacing judges is NP-complete for each admissible quota q and for exact control as well as nearly all preference types. The only exception is control under unrestricted preferences since the uniform constant premise-based quota rules are immune to control in this case. However, in the presence of a complete desired set, possible control under unrestricted and top-respecting preferences becomes tractable, whereas the complexity of exact control and necessary control under top-respecting preferences is unknown. We do not consider the complexity of bundling judges because this control type does not make sense for uniform constant premise-based quota rules. Further, note that since for control by replacing judges the number of judges remains constant throughout the control action, the results for control by replacing judges also hold for the case of uniform premise-based quota rules.

[1]Given n judges, a quota $q \in \mathbb{N}$, $0 \leq q < n$, and a partition of the agenda into premises and conclusions, the collective outcome under the uniform constant premise-based quota rule for quota q consists of the positive premises that more than q judges accept (and the negations of the remaining premises) as well as the conclusions derived from the included premises. The uniform premise-based quota rules are defined in Definition 2.23 on page 38.

In the case of uniform premise-based quota rules, the results are similar. Possible and necessary control by adding, deleting, and bundling judges is NP-complete for the quota $q = 1/2$ and for exact control as well as nearly all preference types. The only exceptions are again control under unrestricted preferences since the uniform premise-based quota rules are immune to control in this case, and possible control for unrestricted and top-respecting preferences since control is tractable in these cases. The results also hold for a complete desired set, but apart from the immunity and the tractability results, the exact complexity of control for different quotas remains an open problem.

My Contribution

The writing of the article Baumeister et al. (2020a) was done jointly with my coauthors. I was responsible for the examples (Example 2, Example 6, and Example 7), Definition 8, the results in Section 4 (i.e., Lemmas 9–11 and Propositions 12–14), Theorem 15, Theorem 16, Theorem 18, Theorem 19, Theorem 21, and Theorem 22.

Chapter 7

Conclusions and Future Work

This thesis covers several topics in the field of computational social choice. The goal is to further the understanding of axiomatic and complexity theoretic properties of decision making procedures.

First, in Chapter 3, my coauthors and I studied a new type of ballot called ℓ-ballot that combines the concept of ordinal and cardinal preferences, and defined two types of committee election rules tailored to these type of ballots. Our proposed minisum and minimax rules were designed to minimize the dissatisfaction that voters have with the winning committee. We then modified several existing properties of single- and multiwinner voting rules to fit our type of ballots and rules and studied the axiomatic properties of the minisum and minimax rules. Further, we were able to show that although the winner determination is NP-hard for the minimax rules, an auxiliary problem asking whether there exists a committee with a voter's maximum dissatisfaction of at most d is fixed-parameter tractable when parameterized by d. As an outlook, we proposed a type of ballot called (a,b)-ballot that is based on cardinal preferences and allows more flexibility for voters to express their underlying preferences. In contrast to existing cardinal-based ballots, here, voters can express their dissatisfaction for a candidate being in a winning committee (a) as well as not being in a winning committee (b), without the restriction that those two values need to be related (e.g., always add up to a fixed constant). Future work includes a characterization of our rules in the context of ℓ-group rules. Furthermore, it would be interesting to define fairness criteria to evaluate the outcome's quality for the voters. In regard to (a,b)-ballots, experiments are needed to determine optimal bounds for the values and study whether the added expressiveness of the model leads to a higher satisfaction of the voters with the election outcome. It would also be interesting to define a cardinal-based variant of the Chamberlin-Courant rules as introduced in Definition 2.14, or to focus on representation in this context, for example by modifying the axiom *justified representation* (Definition 2.16) to allow for cardinal ballots.

My respective coauthors and I were also able to identify several barriers to strategic behavior in voting and judgment aggregation. In Chapter 4, we closed a gap by showing that shift bribery is also hard for several iterative scoring rules, i.e., scoring rules that proceed in rounds where in each round, candidates are eliminated. By allowing the campaign manager to exploit the nonmonotonicity of most of our considered rules, we further showed by using Hare and plurality with runoff as an example that this hardness does not result from restricting the briber to shift the designated candidate forwards in the constructive case (respectively, backwards in the destructive case). Based on our results, Zhou and Guo (2020) started the study of parameterized complexity for the iterative scoring rules considered in this thesis for the parameters number of voters, number of candidates, and budget. For future work, we propose to extend the study of parameterized complexity and to investigate the effect of exploiting nonmonotonicity in-depth. We conjecture that the complexity of shift bribery for all nonmonotonic rules considered by us remains unchanged, but it would be interesting to identify a rule for which shift bribery becomes tractable in these circumstances. Furthermore, domain restrictions as defined on page 13 might also lead to a complexity shift for iterative scoring rules.

For iterative voting, i.e., voting where voters are allowed to update their ballots repeatedly, we studied the manipulative power of the polling agency that announces a dishonest opinion poll to reach a desired outcome of the election. Chapter 5 extended the research on manipulation by the polling agency by introducing a best-response model for the voting rule veto, by studying destructive manipulation, by conducting experiments on efficient heuristics, and most importantly by introducing distance-based problem variants and providing parameterized tractability and intractability results. In particular, we showed that manipulation is para-NP-hard for all considered problems even for very restricted underlying social networks. However, we were able to show that all considered problem variants for veto are tractable when the social network contains no edges, which can be seen as a case where voters are not influenced by their neighbors. Here, future work includes completing the complexity results for plurality in the case of a social network without edges and further the study of parameterized complexity for more natural parameters, for example parameters that describe the underlying social network. We also propose to define best-response dynamics for voting rules that require more complex ballots than plurality and veto and experimentally study how these dynamics affect the quality of the outcome and the possibilities to manipulate by the polling agency. Currently, the polling agency has complete information over the voters' preferences, the voters have no memory and are myopic, they only deviate in the rare cases that they are pivotal, voters trust the polling agency, and they communicate truthfully to their neighbors which candidate they currently vote for. It would therefore be interesting to incorporate changes regarding these aspects into the model.

Finally, in Chapter 6, my coauthors and I introduced control in judgment aggregation. We defined the concepts of control by adding, deleting, replacing, and bundling judges, and proved that these types of control are intractable for the uniform (constant) premise-based quota rules and for several types of the chair's preferences. Our results hold for each rational quota in the case of the uniform constant premise-based quota rules, and for the quota $q = 1/2$ in the case of the uniform premise-based quota rules. Future work includes completing the classical complexity results for all quotas, studying parameterized complexity in our context, and considering new rules. The question whether restricted domains such as the domain of unidimensionally aligned profiles, a variant of the single-crossing domain in preference aggregation, have an impact of the complexity of control in judgment aggregation is also still an open problem. Further, it would be interesting to define new types of control in judgment aggregation, especially types that influence the agenda or the aggregation rule itself instead of the set of judges.

Bibliography

Ågotnes, T., van der Hoek, W., and Wooldridge, M. (2011). On the logic of preference and judgment aggregation. *Journal of Autonomous Agents and Multi-Agent Systems*, 22(1):4–30.

Alcantud, J. C. R. and Laruelle, A. (2014). Dis&approval voting: A characterization. *Social Choice and Welfare*, 43(1):1–10.

Alon, N., Falik, D., Meir, R., and Tennenholtz, M. (2013). Bundling attacks in judgment aggregation. In *Proceedings of the 27th AAAI Conference on Artificial Intelligence*, pages 39–45. AAAI Press.

Arora, S. and Barak, B. (2009). *Computational Complexity: A Modern Approach*. Cambridge University Press.

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

Arrow, K. J. (1950). A difficulty in the concept of social welfare. *Journal of Political Economy*, 58(4):328–346.

Aziz, H., Brill, M., Conitzer, V., Elkind, E., Freeman, R., and Walsh, T. (2017a). Justified representation in approval-based committee voting. *Social Choice and Welfare*, 48(2):461–485.

Aziz, H., Elkind, E., Faliszewski, P., Lackner, M., and Skowron, P. (2017b). The Condorcet principle for multiwinner elections: From shortlisting to proportionality. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence*, pages 84–90. IJCAI.

Aziz, H., Elkind, E., Huang, S., Lackner, M., Sánchez-Fernández, L., and Skowron, P. (2018). On the complexity of extended and proportional justified representation. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence*, pages 902–909. AAAI Press.

Aziz, H. and Lee, B. E. (2020). The expanding approvals rule: Improving proportional representation and monotonicity. *Social Choice and Welfare*, 54(1):1–45.

Baldwin, J. (1926). The technique of the Nanson preferential majority system of election. *Transactions and Proceedings of the Royal Society of Victoria*, 39:42–52.

Balinski, M. and Laraki, R. (2011). *Majority judgment*. Cambridge/Mass.

Ballester, M. and Rey-Biel, P. (2006). Sincere voting with cardinal preferences: Approval voting. Manuscript.

Barberà, S., Bossert, W., and Pattanaik, P. K. (2004). Ranking sets of objects. In *Handbook of Utility Theory*, pages 893–977. Springer.

Barberà, S. and Coelho, D. (2008). How to choose a non-controversial list with k names. *Social Choice and Welfare*, 31(1):79–96.

Bartholdi, III, J. and Trick, M. A. (1986). Stable matching with preferences derived from a psychological model. *Operations Research Letters*, 5(4):165–169.

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

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

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

Bartholdi III, J., Tovey, C. A., and Trick, M. A. (1989b). Voting schemes for which it can be difficult to tell who won the election. *Social Choice and Welfare*, 6(2):157–165.

Baumeister, D., Boes, L., and Weishaupt, R. (2021a). Complexity of sequential rules in judgment aggregation. In *Proceedings of the 20th International Conference on Autonomous Agents and Multiagent Systems*, pages 187–195. IFAAMAS.

Baumeister, D., Böhnlein, T., Rey, L., Schaudt, O., and Selker, A.-K. (2016). Minisum and minimax committee election rules for general preference types. In *Proceedings of the 22nd European Conference on Artificial Intelligence*, pages 1656–1657. IOS Press. Extended Abstract.

Baumeister, D. and Dennisen, S. (2015). Voter dissatisfaction in committee elections. In *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems*. IFAAMAS. Extended abstract.

Baumeister, D., Dennisen, S., and Rey, L. (2015a). Winner determination and manipulation in minisum and minimax committee elections. In *Proceedings of the 4th International Conference on Algorithm Decision Theory*, pages 469–485. Springer-Verlag *Lecture Notes in Computer Science #9346*.

Baumeister, D., Erdélyi, G., Erdélyi, O., and Rothe, J. (2012). Control in judgment aggregation. In *Proceedings of the 6th European Starting AI Researcher Symposium*, pages 23–34. IOS Press.

Baumeister, D., Erdélyi, G., Erdélyi, O., and Rothe, J. (2013). Computational aspects of manipulation and control in judgment aggregation. In *Proceedings of the 3rd International Conference on Algorithmic Decision Theory*, pages 71–85. Springer-Verlag *Lecture Notes in Artificial Intelligence #8176*.

Baumeister, D., Erdélyi, G., Erdélyi, O. J., and Rothe, J. (2015b). Complexity of manipulation and bribery in judgment aggregation for uniform premise-based quota rules. *Mathematical Social Sciences*, 76:19–30.

Baumeister, D., Erdélyi, G., Erdélyi, O. J., Rothe, J., and Selker, A.-K. (2020a). Complexity of control in judgment aggregation for uniform premise-based quota rules. *Journal of Computer and System Sciences*, 112:13–33.

Baumeister, D., Erdélyi, G., and Rothe, J. (2015c). Judgment aggregation. In Rothe, J., editor, *Economics and Computation. An Introduction to Algorithmic Game Theory, Computational Social Choice, and Fair Division*, chapter 6, pages 361–391. Springer-Verlag.

Baumeister, D., Hogrebe, T., and Rey, L. (2019). Generalized distance bribery. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence*, pages 1764–1771. AAAI Press.

Baumeister, D. and Rothe, J. (2015). Preference aggregation by voting. In Rothe, J., editor, *Economics and Computation. An Introduction to Algorithmic Game Theory,* *Computational Social Choice, and Fair Division*, chapter 4, pages 197–325. Springer-Verlag.

Baumeister, D., Rothe, J., and Selker, A.-K. (2015d). Complexity of bribery and control for uniform premise-based quota rules under various preference types. In *Proceedings of the 4th International Conference on Algorithmic Decision Theory*, pages 432–448. Springer-Verlag *Lecture Notes in Artificial Intelligence #9346*.

Baumeister, D., Rothe, J., and Selker, A.-K. (2017). Strategic behavior in judgment aggregation. In Endriss, U., editor, *Trends in Computational Social Choice*, chapter 8, pages 145–168. AI Access.

Baumeister, D., Selker, A.-K., and Wilczynski, A. (2020b). Manipulation of opinion polls to influence iterative elections. In *Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems*, pages 132–140. IFAAMAS.

Baumeister, D., Selker, A.-K., and Wilczynski, A. (2021b). Manipulation of opinion polls to influence iterative elections. In *Proceedings of the 8th International Workshop on Computational Social Choice*.

Betzler, N., Slinko, A., and Uhlmann, J. (2013). On the computation of fully proportional representation. *Journal of Artificial Intelligence Research*, 47:475–519.

Black, D. (1958). *The Theory of Committees and Elections*. Cambridge University Press.

Botan, S., de Haan, R., Slavkovik, M., and Terzopoulou, Z. (2021). Egalitarian judgment aggregation. In *Proceedings of the 20th International Conference on Autonomous Agents and Multiagent Systems*, pages 214–222. IFAAMAS.

Botan, S., Novaro, A., and Endriss, U. (2016). Group manipulation in judgment aggregation. In *Proceedings of the 15th International Conference on Autonomous Agents and Multiagent Systems*, pages 411–419. IFAAMAS.

Brams, S., Kilgour, D., and Sanver, R. (2007). A minimax procedure for electing committees. *Public Choice*, 132:401–420.

Brandt, F., Brill, M., Hemaspaandra, E., and Hemaspaandra, L. A. (2015). Bypassing combinatorial protections: Polynomial-time algorithms for single-peaked electorates. *Journal of Autonomous Agents and Multi-Agent Systems*, 53:439–496.

Brandt, F., Conitzer, V., Endriss, U., Lang, J., and Procaccia, A. D., editors (2016). *Handbook of Computational Social Choice*. Cambridge University Press.

Brânzei, S., Caragiannis, I., Morgenstern, J., and Procaccia, A. D. (2013). How bad is selfish voting? In *Proceedings of the 27th AAAI Conference on Artificial Intelligence*, pages 138–144. AAAI Press.

Bredereck, R., Chen, J., Faliszewski, P., Guo, J., Niedermeier, R., and Woeginger, G. J. (2014a). Parameterized algorithmics for computational social choice: Nine research challenges. *Tsinghua Science and Technology*, 19(4):358–373.

Bredereck, R., Chen, J., Faliszewski, P., Nichterlein, A., and Niedermeier, R. (2016a). Prices matter for the parameterized complexity of shift bribery. *Information and Computation*, 251:140–164.

Bredereck, R., Chen, J., Hartung, S., Kratsch, S., Niedermeier, R., Suchy, O., and Woeginger, G. J. (2014b). A multivariate complexity analysis of lobbying in multiple referenda. *Journal of Artificial Intelligence Research*, 50:409–446.

Bredereck, R., Chen, J., Niedermeier, R., and Walsh, T. (2017). Parliamentary voting procedures: Agenda control, manipulation, and uncertainty. *Journal of Artificial Intelligence Research*, 59:133–173.

Bredereck, R., Chen, J., and Woeginger, G. J. (2013). A characterization of the single-crossing domain. *Social Choice and Welfare*, 41(4):989–998.

Bredereck, R., Faliszewski, P., Igarashi, A., Lackner, M., and Skowron, P. (2018). Multi-winner elections with diversity constraints. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence*, pages 933–940. AAAI Press.

Bredereck, R., Faliszewski, P., Kaczmarczyk, A., Niedermeier, R., Skowron, P., and Talmon, N. (2021). Robustness among multiwinner voting rules. *Artificial Intelligence*, 290:103403.

Bredereck, R., Faliszewski, P., Niedermeier, R., Skowron, P., and Talmon, N. (2015). Elections with few candidates: Prices, weights, and covering problems. In *Proceedings of the 4th International Conference on Algorithmic Decision Theory*, pages 414–431. Springer-Verlag *Lecture Notes in Artificial Intelligence #9346*.

Bredereck, R., Faliszewski, P., Niedermeier, R., and Talmon, N. (2016b). Complexity of shift bribery in committee elections. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence*, pages 2452–2458. AAAI Press.

Bredereck, R., Faliszewski, P., Niedermeier, R., and Talmon, N. (2016c). Large-scale election campaigns: Combinatorial shift bribery. *Journal of Artificial Intelligence Research*, 55:603–652.

Brill, M., Faliszewski, P., Sommer, F., and Talmon, N. (2019). Approximation algorithms for balancedCC multiwinner rules. In *Proceedings of the 18th International Conference on Autonomous Agents and Multiagent Systems*, pages 494–502. IFAAMAS.

Chamberlin, J. R. and Courant, P. N. (1983). Representative deliberations and representative decisions: Proportional representation and the Borda rule. *American Political Science Review*, 77(3):718–733.

Chen, J., Faliszewski, P., Niedermeier, R., and Talmon, N. (2017). Elections with few voters: Candidate control can be easy. *Journal of Artificial Intelligence Research*, 60:937–1002.

Christian, R., Fellows, M., Rosamond, F., and Slinko, A. (2007). On complexity of lobbying in multiple referenda. *Review of Economic Design*, 11(3):217–224.

Condorcet, M. d. (1785). *Essay sur l'application de l'analyse de la probabilité des decisions: Redues et pluralité des voix*. l'Imprimerie Royale. Translated by McLean and Urken (1995).

Conitzer, V. and Sandholm, T. (2003). Universal voting protocol tweaks to make manipulation hard. In *Proceedings of the 18th International Joint Conference on Artificial Intelligence*, pages 781–788.

Conitzer, V., Sandholm, T., and Lang, J. (2007). When are elections with few candidates hard to manipulate? *Journal of the ACM*, 54(3):Article 14.

Conitzer, V. and Walsh, T. (2016). Barriers to manipulation in voting. In Brandt, F., Conitzer, V., Endriss, U., Lang, J., and Procaccia, A., editors, *Handbook of Computational Social Choice*, chapter 6, pages 127–145. Cambridge University Press.

Conitzer, V., Walsh, T., and Xia, L. (2011). Dominating manipulations in voting with partial information. In *Proceedings of the 25th AAAI Conference on Artificial Intelligence*, pages 638–643. AAAI Press.

Cook, S. A. (1971). The complexity of theorem-proving procedures. In *Proceedings of the 3rd Annual ACM Symposium on Theory of Computing*, pages 151–158. ACM Press.

Copeland, A. H. (1951). A reasonable social welfare function. Technical report, Mimeo, University of Michigan.

Cygan, M., Fomin, F. V., Kowalik, Ł., Lokshtanov, D., Marx, D., Pilipczuk, M., Pilipczuk, M., and Saurabh, S. (2015). *Parameterized algorithms*. Springer.

Cygan, M., Kowalik, Ł., Socała, A., and Sornat, K. (2018). Approximation and parameterized complexity of minimax approval voting. *Journal of Artificial Intelligence Research*, 63:495–513.

de Haan, R. (2016). Parameterized complexity results for the Kemeny rule in judgment aggregation. In *Proceedings of the 22nd European Conference on Artificial Intelligence*, pages 1502–1510. IOS Press.

de Haan, R. (2017). Complexity results for manipulation, bribery and control of the Kemeny judgment aggregation procedure. In *Proceedings of the 16th International Conference on Autonomous Agents and Multiagent Systems*, pages 1151–1159. IFAAMAS.

de Haan, R. (2018). Hunting for tractable languages for judgment aggregation. In *Proceedings of the 16th International Conference on Principles of Knowledge Representation and Reasoning*, pages 194–203. AAAI Press.

de Haan, R. and Slavkovik, M. (2017). Complexity results for aggregating judgments using scoring or distance-based procedures. In *Proceedings of the 16th International Conference on Autonomous Agents and Multiagent Systems*, pages 952–961. IFAAMAS.

de Haan, R. and Slavkovik, M. (2019). Answer set programming for judgment aggregation. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence*, pages 1668–1674. IJCAI.

Debord, B. (1992). An axiomatic characterization of Borda's k-choice function. *Social Choice and Welfare*, 9(4):337–343.

Dey, P. (2021). Local distance constrained bribery in voting. *Theoretical Computer Science*, 849:1–21.

Dey, P., Misra, N., and Narahari, Y. (2017). Frugal bribery in voting. *Theoretical Computer Science*, 676:15–32.

Dey, P., Misra, N., and Narahari, Y. (2018). Complexity of manipulation with partial information in voting. *Theoretical Computer Science*, 726:78–99.

Dey, P. and Narahari, Y. (2015). Estimating the margin of victory of an election using sampling. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence*, pages 1120–1126. AAAI Press/IJCAI.

Dietrich, F. (2007). A generalised model of judgment aggregation. *Social Choice and Welfare*, 28(4):529–565.

Dietrich, F. (2014). Scoring rules for judgment aggregation. *Social Choice and Welfare*, 42(4):873–911.

Dietrich, F. (2016). Judgment aggregation and agenda manipulation. *Games and Economic Behavior*, 95:113–136.

Dietrich, F. and List, C. (2007a). Arrow's theorem in judgment aggregation. *Social Choice and Welfare*, 29(1):19–33.

Dietrich, F. and List, C. (2007b). Judgment aggregation by quota rules: Majority voting generalized. *Journal of Theoretical Politics*, 19(4):391–424.

Dietrich, F. and List, C. (2007c). Strategy-proof judgment aggregation. *Economics and Philosophy*, 23(3):269–300.

Dietrich, F. and List, C. (2010). The impossibility of unbiased judgment aggregation. *Theory and Decision*, 68(3):281–299.

Dietrich, F. and Mongin, P. (2010). The premiss-based approach to judgment aggregation. *Journal of Economic Theory*, 145(2):562–582.

Dodgson, C. (1876). A method of taking votes on more than two issues. *The Theory of Committees and Elections*. Reprinted by McLean and Urken (1995).

Dokow, E. and Holzman, R. (2010). Aggregation of binary evaluations. *Journal of Economic Theory*, 145(2):495–511.

Dorn, B. and Schlotter, I. (2012). Multivariate complexity analysis of swap bribery. *Algorithmica*, 64(1):126–151.

Downey, R. G. and Fellows, M. R. (1995a). Fixed-parameter tractability and completeness I: Basic results. *SIAM Journal on Computing*, 24(4):873–921.

Downey, R. G. and Fellows, M. R. (1995b). Fixed-parameter tractability and completeness II: On completeness for W[1]. *Theoretical Computer Science*, 141(1-2):109–131.

Dummett, M. (1984). *Voting procedures*. Oxford University Press.

Elkind, E. and Faliszewski, P. (2010). Approximation algorithms for campaign management. In *International Workshop on Internet and Network Economics*, pages 473–482. Springer.

Elkind, E., Faliszewski, P., Gupta, S., and Roy, S. (2020a). Algorithms for swap and shift bribery in structured elections. In *Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems*, pages 366–374. IFAAMAS.

Elkind, E., Faliszewski, P., and Skowron, P. (2020b). A characterization of the single-peaked single-crossing domain. *Social Choice and Welfare*, 54(1):167–181.

Elkind, E., Faliszewski, P., Skowron, P., and Slinko, A. (2017a). Properties of multiwinner voting rules. *Social Choice and Welfare*, 48(3):599–632.

Elkind, E., Faliszewski, P., and Slinko, A. (2009). Swap bribery. In *Proceedings of the 2nd International Symposium on Algorithmic Game Theory*, pages 299–310. Springer-Verlag *Lecture Notes in Computer Science #5814*.

Elkind, E., Faliszewski, P., and Slinko, A. (2011). Cloning in elections: Finding the possible winners. *Journal of Artificial Intelligence Research*, 42:529–573.

Elkind, E., Faliszewski, P., and Slinko, A. (2012). Clone structures in voters' preferences. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, pages 496–513.

Elkind, E. and Lackner, M. (2015). Structure in dichotomous preferences. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence*, pages 2019–2025. AAAI Press/IJCAI.

Elkind, E., Lackner, M., and Peters, D. (2017b). Structured preferences. In Endriss, U., editor, *Trends in Computational Social Choice*, chapter 10, pages 187–207. AI Access.

Elkind, E., Lang, J., and Saffidine, A. (2015). Condorcet winning sets. *Social Choice and Welfare*, 44(3):493–517.

Endriss, U. (2016). Judgment aggregation. In *Handbook of Computational Social Choice*, chapter 17, pages 399–426. Cambridge University Press.

Endriss, U., editor (2017). *Trends in Computational Social Choice*. AI Access.

Endriss, U. and de Haan, R. (2015). Complexity of the winner determination problem in judgment aggregation: Kemeny, Slater, Tideman, Young. In *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems*, pages 117–125. IFAAMAS.

Endriss, U., de Haan, R., Lang, J., and Slavkovik, M. (2020). The complexity landscape of outcome determination in judgment aggregation. *Journal of Artificial Intelligence Research*, 69:687–731.

Endriss, U., de Haan, R., and Szeider, S. (2015). Parameterized complexity results for agenda safety in judgment aggregation. In *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems*, pages 127–136. IFAAMAS.

Endriss, U., Grandi, U., de Haan, R., and Lang, J. (2016a). Succinctness of languages for judgment aggregation. In *Proceedings of the 15th International Conference on Principles of Knowledge Representation and Reasoning*, pages 176–185. AAAI Press.

Endriss, U., Grandi, U., and Porello, D. (2012). Complexity of judgment aggregation. *Journal of Artificial Intelligence Research*, 45:481–514.

Endriss, U., Obraztsova, S., Polukarov, M., and Rosenschein, J. S. (2016b). Strategic voting with incomplete information. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence*, pages 236–242. AAAI Press/IJCAI.

Fairstein, R., Lauz, A., Meir, R., and Gal, K. (2019). Modeling people's voting behavior with poll information. In *Proceedings of the 18th International Conference on Autonomous Agents and Multiagent Systems*, pages 1422–1430. IFAAMAS.

Faliszewski, P., Hemaspaandra, E., and Hemaspaandra, L. (2009a). How hard is bribery in elections? *Journal of Artificial Intelligence Research*, 35:485–532.

Faliszewski, P., Hemaspaandra, E., Hemaspaandra, L. A., and Rothe, J. (2009b). Llull and Copeland voting computationally resist bribery and constructive control. *Journal of Artificial Intelligence Research*, 35:275–341.

Faliszewski, P., Hemaspaandra, E., Hemaspaandra, L. A., and Rothe, J. (2011). The shield that never was: Societies with single-peaked preferences are more open to manipulation and control. *Information and Computation*, 209(2):89–107.

Faliszewski, P., Manurangsi, P., and Sornat, K. (2021). Approximation and hardness of shift-bribery. *Artificial Intelligence*, 298:103520.

Faliszewski, P. and Procaccia, A. D. (2010). AI's war on manipulation: Are we winning? *AI Magazine*, 31(4):53–64.

Faliszewski, P., Rothe, I., and Rothe, J. (2016). Noncooperative game theory. In *Economics and Computation*, pages 41–134. Springer.

Faliszewski, P. and Rothe, J. (2016). Control and bribery in voting. In Brandt, F., Conitzer, V., Endriss, U., Lang, J., and Procaccia, A., editors, *Handbook of Computational Social Choice*, chapter 7, pages 146–168. Cambridge University Press.

Faliszewski, P., Skowron, P., Slinko, A., and Talmon, N. (2017a). Multiwinner rules on paths from k–Borda to Chamberlin–Courant. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence*, pages 192–198. IJCAI.

Faliszewski, P., Skowron, P., Slinko, A., and Talmon, N. (2017b). Multiwinner voting: A new challenge for social choice theory. In Endriss, U., editor, *Trends in Computational Social Choice*, chapter 2, pages 27–47. AI Access.

Faliszewski, P., Skowron, P., Slinko, A., and Talmon, N. (2019). Committee scoring rules: Axiomatic characterization and hierarchy. *ACM Transactions on Economics and Computation*, 7(1):1–39.

Faliszewski, P., Skowron, P., and Talmon, N. (2017c). Bribery as a measure of candidate success: Complexity results for approval-based multiwinner rules. In *Proceedings of the 16th International Conference on Autonomous Agents and Multiagent Systems*, pages 6–14. IFAAMAS.

Faliszewski, P., Slinko, A., and Talmon, N. (2020). Multiwinner rules with variable number of winners. In *Proceedings of the 24th European Conference on Artificial Intelligence*, pages 67–74. IOS Press.

Felsenthal, D. and Maoz, Z. (1992). Normative properties of four single-stage multi-winner electoral procedures. *Behavioral Science*, 37(2):109–127.

Fishburn, P. C. (1974). Paradoxes of voting. *The American Political Science Review*, 68(2):537–546.

Flum, J. and Grohe, M. (2003). Describing parameterized complexity classes. *Information and Computation*, 187(2):291–319.

Flum, J. and Grohe, M. (2006). *Parameterized Complexity Theory*. EATCS Texts in Theoretical Computer Science. Springer.

Garey, M. R. and Johnson, D. S. (1979). *Computers and intractability*, volume 174. freeman San Francisco.

Gehrlein, W. (1985). The Condorcet criterion and committee selection. *Mathematical Social Sciences*, 10(3):199–209.

Gibbard, A. (1973). Manipulation of voting schemes: A general result. *Econometrica*, 41(4):587–601.

Gourvès, L., Lesca, J., and Wilczynski, A. (2016). Strategic voting in a social context: Considerate equilibria. In *Proceedings of the 22nd European Conference on Artificial Intelligence*, pages 1423–1431. IOS Press.

Grandi, U., Loreggia, A., Rossi, F., Venable, K. B., and Walsh, T. (2013). Restricted manipulation in iterative voting: Condorcet efficiency and Borda score. In *Proceedings of the 3rd International Conference on Algorithmic Decision Theory*, pages 181–192. Springer-Verlag *Lecture Notes in Artificial Intelligence #8176*.

Gupta, S., Jain, P., Roy, S., Saurabh, S., and Zehavi, M. (2019). Gehrlein stability in committee selection: Parameterized hardness and algorithms. In *Proceedings of the 18th International Conference on Autonomous Agents and Multiagent Systems*, pages 511–519. IFAAMAS.

Hartmanis, J. and Stearns, R. E. (1965). On the computational complexity of algorithms. *Transactions of the American Mathematical Society*, 117:285–306.

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

Hemaspaandra, E., Hemaspaandra, L. A., and Rothe, J. (2007). Anyone but him: The complexity of precluding an alternative. *Artificial Intelligence*, 171(5-6):255–285.

Hemaspaandra, E., Spakowski, H., and Vogel, J. (2005). The complexity of Kemeny elections. *Theoretical Computer Science*, 349(3):382–391.

Izsak, R., Talmon, N., and Woeginger, G. (2018). Committee selection with intraclass and interclass synergies. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence*, pages 1071–1078. AAAI Press.

Johnson, D. S. (1974). Approximation algorithms for combinatorial problems. *Journal of Computer and System Sciences*, 9(3):256–278.

Kaczmarczyk, A. and Faliszewski, P. (2019). Algorithms for destructive shift bribery. *Journal of Autonomous Agents and Multi-Agent Systems*, 33(3):275–297.

Kagita, V. R., Pujari, A. K., Padmanabhan, V., Aziz, H., and Kumar, V. (2021). Committee selection using attribute approvals. In *Proceedings of the 20th International Conference on Autonomous Agents and Multiagent Systems*, pages 683–691. IFAAMAS.

Karp, R. M. (1972). Reducibility among combinatorial problems. In *Complexity of Computer Computations*, pages 85–103. Springer.

Kemeny, J. (1959). Mathematics without numbers. *Daedalus*, 88(4):577–591.

Kendall, M. G. (1938). A new measure of rank correlation. *Biometrika*, 30(1/2):81–93.

Kilgour, D., Brams, S., and Sanver, R. (2006). How to elect a representative committee using approval balloting. In Pukelsheim, F. and Simeone, B., editors, *Mathematics and Democracy: Recent Advances in Voting Systems and Collective Choice*, pages 83–95. Springer.

Kilgour, D. M. (2016). Approval elections with a variable number of winners. *Theory and Decision*, 81(2):199–211.

Knop, D., Koutecký, M., and Mnich, M. (2020). Voting and bribing in single-exponential time. *ACM Transactions on Economics and Computation*, 8(3):1–28.

Kocot, M., Kolonko, A., Elkind, E., Faliszewski, P., and Talmon, N. (2019). Multigoal committee selection. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence*, pages 385–391. IJCAI.

Kornhauser, L. and Sager, L. (1986). Unpacking the court. *Yale Law Journal*, 96(1):82–117.

Lackner, M. and Maly, J. (2021). Approval-based shortlisting. In *Proceedings of the 20th International Conference on Autonomous Agents and Multiagent Systems*, pages 737–745. IFAAMAS.

Lackner, M. and Skowron, P. (2018). Approval-based multi-winner rules and strategic voting. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence*, pages 340–346. IJCAI.

Lackner, M. and Skowron, P. (2019). A quantitative analysis of multi-winner rules. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence*, pages 407–413. IJCAI.

Lackner, M. and Skowron, P. (2021). Consistent approval-based multi-winner rules. *Journal of Economic Theory*, 192:105173.

Lang, J., Pigozzi, G., Slavkovik, M., and Torre, J. (2011). Judgment aggregation rules based on minimization. In *Proceedings of the 13th Conference on Theoretical Aspects of Rationality and Knowledge*, pages 238–246. ACM Press.

Lang, J., Pigozzi, G., Slavkovik, M., van der Torre, L., and Vesic, S. (2017). A partial taxonomy of judgment aggregation rules and their properties. *Social Choice and Welfare*, 48(2):327–356.

Lang, J. and Slavkovik, M. (2013). Judgment aggregation rules and voting rules. In *Proceedings of the 3rd International Conference on Algorithmic Decision Theory*, pages 230–243. Springer-Verlag *Lecture Notes in Artificial Intelligence #8176*.

Lang, J. and Slavkovik, M. (2014). How hard is it to compute majority-preserving judgment aggregation rules? In *Proceedings of the 21st European Conference on Artificial Intelligence*, pages 501–506. IOS Press.

Laslier, J.-F. and Sanver, M. R. (2010). *Handbook on Approval Voting*. Springer Science & Business Media.

LeGrand, R. (2004). Analysis of the minimax procedure. Technical Report WUCSE-2004-67, Department of Computer Science and Engineering, Washington University, St. Louis, Missouri.

Lev, O. and Rosenschein, J. S. (2016). Convergence of iterative scoring rules. *Journal of Artificial Intelligence Research*, 57:573–591.

Levin, J. and Nalebuff, B. (1995). An introduction to vote-counting schemes. *The Journal of Economic Perspectives*, 9(1):3–26.

List, C. (2003). A possibility theorem on aggregation over multiple interconnected propositions. *Mathematical Social Sciences*, 45(1):1–13.

List, C. (2004). A model of path-dependence in decisions over multiple propositions. *American Political Science Review*, 93(3):495–513.

List, C. and Pettit, P. (2002). Aggregating sets of judgments: An impossibility result. *Economics and Philosophy*, 18(1):89–110.

List, C. and Pettit, P. (2004). Aggregating sets of judgments: Two impossibility results compared. *Synthese*, 140(1–2):207–235.

List, C. and Puppe, C. (2009). Judgment aggregation. In Anand, P., Pattanaik, P., and Puppe, C., editors, *The Handbook of Rational and Social Choice*, chapter 19, pages 457–482. Oxford University Press.

Liu, H. and Guo, J. (2016). Parameterized complexity of winner determination in minimax committee elections. In *Proceedings of the 15th International Conference on Autonomous Agents and Multiagent Systems*, pages 341–349. IFAAMAS.

Magiera, K. and Faliszewski, P. (2017). How hard is control in single-crossing elections? *Journal of Autonomous Agents and Multi-Agent Systems*, 31(3):606–627.

Maushagen, C., Neveling, M., Rothe, J., and Selker, A.-K. (2018a). Complexity of shift bribery in hare, coombs, baldwin, and nanson elections. In *Proceedings of the 15th International Symposium on Artificial Intelligence and Mathematics*.

Maushagen, C., Neveling, M., Rothe, J., and Selker, A.-K. (2018b). Complexity of shift bribery in iterative elections. In *Proceedings of the 17th International Conference on Autonomous Agents and Multiagent Systems*, pages 1567–1575. IFAAMAS.

Maushagen, C., Neveling, M., Rothe, J., and Selker, A.-K. (2021). Complexity of shift bribery for iterative voting rules. Submitted to *Annals of Mathematics and Artificial Intelligence*.

May, K. (1952). A set of independent, necessary and sufficient conditions for simple majority décisions. *Econometrica*, 20:680–684.

McLean, I. and Urken, A. B. (1995). *Classics of Social Choice*. University of Michigan Press.

Meir, R. (2017). Iterative voting. In Endriss, U., editor, *Trends in Computational Social Choice*, chapter 4, pages 69–86. AI Access.

Meir, R., Lev, O., and Rosenschein, J. S. (2014). A local-dominance theory of voting equilibria. In *Proceedings of the 15th ACM Conference on Economics and Computation*, pages 313–330. ACM Press.

Meir, R., Polukarov, M., Rosenschein, J. S., and Jennings, N. R. (2010). Convergence to equilibria in plurality voting. In *Proceedings of the 24th AAAI Conference on Artificial Intelligence*, pages 823–828. AAAI Press.

Meir, R., Procaccia, A., Rosenschein, J., and Zohar, A. (2008). Complexity of strategic behaviour in multi-winner elections. *Journal of Artificial Intelligence Research*, 33:149–178.

Miasko, T. and Faliszewski, P. (2016). The complexity of priced control in elections. *Annals of Mathematics and Artificial Intelligence*, 77(3-4):225–250.

Miller, M. and Osherson, D. (2009). Methods for distance-based judgment aggregation. *Social Choice and Welfare*, 32(4):575–601.

Mirrlees, J. A. (1971). An exploration in the theory of optimum income taxation. *The Review of Economic Studies*, 38(2):175–208.

Misra, N., Nabeel, A., and Singh, H. (2015). On the parameterized complexity of minimax approval voting. In *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems*, pages 97–105. IFAAMAS.

Mongin, P. (2012). The doctrinal paradox, the discursive dilemma, and logical aggregation theory. *Theory and Decision*, 73(3):315–355.

Monroe, B. L. (1995). Fully proportional representation. *American Political Science Review*, 89(4):925–940.

Nanson, E. (1882). Methods of election. *Transactions and Proceedings of the Royal Society of Victoria*, 19:197–240.

Obraztsova, S. and Elkind, E. (2012). Optimal manipulation of voting rules. In *Proceedings of the 26th AAAI Conference on Artificial Intelligence*, pages 2141–2147. AAAI Press.

Obraztsova, S., Elkind, E., Polukarov, M., and Rabinovich, Z. (2015a). Doodle poll games. In *Proceedings of the First IJCAI-Workshop on Algorithmic Game Theory*.

Obraztsova, S., Lev, O., Polukarov, M., Rabinovich, Z., and Rosenschein, J. S. (2016). Nonmyopic voting dynamics: An optimistic approach. In *Proceedings of the 10th Multidisciplinary Workshop on Advances in Preference Handling*.

Obraztsova, S., Markakis, E., Polukarov, M., Rabinovich, Z., and Jennings, N. R. (2015b). On the convergence of iterative voting: How restrictive should restricted dynamics be? In *Proceedings of the 29th AAAI Conference on Artificial Intelligence*, pages 993–999. AAAI Press.

Obraztsova, S., Polukarov, M., Rabinovich, Z., and Elkind, E. (2017). Doodle poll games. In *Proceedings of the 16th International Conference on Autonomous Agents and Multiagent Systems*, pages 876–884. IFAAMAS.

Obraztsova, S., Zick, Y., and Elkind, E. (2013). On manipulation in multi-winner elections based on scoring rules. In *Proceedings of the 12th International Conference on Autonomous Agents and Multiagent Systems*, pages 359–366. IFAAMAS.

Peters, D. (2018a). Proportionality and strategyproofness in multiwinner elections. In *Proceedings of the 17th International Conference on Autonomous Agents and Multiagent Systems*, pages 1549–1557. IFAAMAS.

Peters, D. (2018b). Single-peakedness and total unimodularity: New polynomial-time algorithms for multi-winner elections. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence*, pages 1169–1176. AAAI Press.

Pettit, P. (2001). Deliberative democracy and the discursive dilemma. *Philosophical Issues*, 11(1):268–299.

Pigozzi, G. (2006). Belief merging and the discursive dilemma: An argument-based account of paradoxes of judgment. *Synthese*, 152(2):285–298.

Procaccia, A. and Rosenschein, J. (2006). The distortion of cardinal preferences in voting. In *Proceedings of the 10th International Workshop on Cooperative Information Agents*, pages 317–331. Springer-Verlag *Lecture Notes in Computer Science #4149*.

Procaccia, A., Rosenschein, J., and Zohar, A. (2007). Multi-winner elections: Complexity of manipulation, control, and winner-determination. In *Proceedings of the 20th International Joint Conference on Artificial Intelligence*, pages 1476–1481.

Procaccia, A. D. and Rosenschein, J. S. (2007). Average-case tractability of manipulation in voting via the fraction of manipulators. In *Proceedings of the 6th International Conference on Autonomous Agents and Multiagent Systems*, pages 718–720. IFAAMAS.

Ratliff, T. C. (2003). Some startling inconsistencies when electing committees. *Social Choice and Welfare*, 21(3):433–454.

Reijngoud, A. and Endriss, U. (2012). Voter response to iterated poll information. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems*, pages 635–644, Valencia, Spain. IFAAMAS.

Reisch, Y., Rothe, J., and Schend, L. (2014). The margin of victory in Schulze, cup, and Copeland elections: Complexity of the regular and exact variants. In *Proceedings of the 7th European Starting AI Researcher Symposium*, pages 250–259. IOS Press.

Reyhani, R. and Wilson, M. C. (2012). Best reply dynamics for scoring rules. In *Proceedings of the 20th European Conference on Artificial Intelligence*, pages 672–677. IOS Press.

Rothe, J. and Schend, L. (2013). Challenges to complexity shields that are supposed to protect elections against manipulation and control: A survey. *Annals of Mathematics and Artificial Intelligence*, 68(1–3):161–193.

Sánchez-Fernández, L., Elkind, E., Lackner, M., Fernández, N., Fisteus, J. A., Val, P. B., and Skowron, P. (2017). Proportional justified representation. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence*, pages 670–676. AAAI Press.

Sánchez-Fernández, L. and Fisteus, J. A. (2019). Monotonicity axioms in approval-based multi-winner voting rules. In *Proceedings of the 18th International Conference on Autonomous Agents and Multiagent Systems*, pages 485–493. IFAAMAS.

Saporiti, A. and Tohmé, F. (2006). Single-crossing, strategic voting and the median choice rule. *Social Choice and Welfare*, 26(2):363–383.

Satterthwaite, M. (1975). Strategy-proofness and Arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, 10(2):187–217.

Schlotter, I., Faliszewski, P., and Elkind, E. (2017). Campaign management under approval-driven voting rules. *Algorithmica*, 77(1):84–115.

Shiryaev, D., Yu, L., and Elkind, E. (2013). On elections with robust winners. In *Proceedings of the 12th International Conference on Autonomous Agents and Multiagent Systems*, pages 415–422. IFAAMAS.

Sina, S., Hazon, N., Hassidim, A., and Kraus, S. (2015). Adapting the social network to affect elections. In *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems*, pages 705–713. IFAAMAS.

Skowron, P. (2018). Proportionality degree of multiwinner rules. Technical Report arXiv: 1810.08799v2 [cs.GT], ACM Computing Research Repository (CoRR).

Skowron, P. and Faliszewski, P. (2017). Chamberlin–Courant rule with approval ballots: Approximating the MaxCover problem with bounded frequencies in FPT time. *Journal of Artificial Intelligence Research*, 60:687–716.

Skowron, P., Faliszewski, P., and Lang, J. (2016). Finding a collective set of items: From proportional multirepresentation to group recommendation. *Artificial Intelligence*, 241:191–216.

Skowron, P., Faliszewski, P., and Slinko, A. (2015a). Achieving fully proportional representation: Approximability results. *Artificial Intelligence*, 222:67–103.

Skowron, P., Faliszewski, P., and Slinko, A. (2019). Axiomatic characterization of committee scoring rules. *Journal of Economic Theory*, 180:244–273.

Skowron, P., Yu, L., Faliszewski, P., and Elkind, E. (2015b). The complexity of fully proportional representation for single-crossing electorates. *Theoretical Computer Science*, 569:43–57.

Slavkovik, M. (2014). Not all judgment aggregation should be neutral. In *Proceedings of the European Conference on Social Intelligence*, pages 198–211.

Smith, W. D. (2000). Range voting. http://scorevoting.net/WarrenSmithPages/homepage/rangevote.pdf.

Taylor, A. D. (2005). *Social Choice and the Mathematics of Manipulation*. Cambridge University Press.

Terzopoulou, Z. and Endriss, U. (2018). Modelling iterative judgment aggregation. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence*, pages 1234–1241. AAAI Press.

Terzopoulou, Z. and Endriss, U. (2019a). Rethinking the neutrality axiom in judgment aggregation. In *Proceedings of the 18th International Conference on Autonomous Agents and Multiagent Systems*, pages 2212–2214. IFAAMAS.

Terzopoulou, Z. and Endriss, U. (2019b). Strategyproof judgment aggregation under partial information. *Social Choice and Welfare*, 53(3):415–442.

Terzopoulou, Z. and Endriss, U. (2020). Neutrality and relative acceptability in judgment aggregation. *Social Choice and Welfare*, 55(1):25–49.

Terzopoulou, Z., Endriss, U., and de Haan, R. (2018). Aggregating incomplete judgments: Axiomatisations for scoring rules. In *Proceedings of the 7th International Workshop on Computational Social Choice*.

Tideman, T. (1987). Independence of clones as a criterion for voting rules. *Social Choice and Welfare*, 4(3):185–206.

Tsang, A. and Larson, K. (2016). The echo chamber: Strategic voting and homophily in social networks. In *Proceedings of the 15th International Conference on Autonomous Agents and Multiagent Systems*, pages 368–375. IFAAMAS.

Vazirani, V. V. (2013). *Approximation Algorithms*. Springer Science & Business Media.

Wilczynski, A. (2019). Poll-confident voters in iterative voting. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence*, pages 2205–2212. AAAI Press.

Xia, L. (2012). Computing the margin of victory for various voting rules. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, pages 982–999. ACM Press.

Yang, Y. (2019). Complexity of manipulating and controlling approval-based multiwinner voting. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence*, pages 637–643. IJCAI.

Yang, Y. (2020). On the complexity of destructive bribery in approval-based multi-winner voting. In *Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems*, pages 1584–1592. IFAAMAS.

Zhou, A. and Guo, J. (2020). Parameterized complexity of shift bribery in iterative elections. In *Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems*, pages 1665–1673. IFAAMAS.

Zwicker, W. S. (2016). Introduction to the theory of voting. In Brandt, F., Conitzer, V., Endriss, U., Lang, J., and Procaccia, A., editors, *Handbook of Computational Social Choice*, chapter 1, pages 23–56. Cambridge University Press.

Eidesstattliche Erklärung

entsprechend §5 der Promotionsordnung vom 15.06.2018

Ich versichere an Eides Statt, dass die Dissertation von mir selbständig und ohne unzulässige fremde Hilfe unter Beachtung der „Grundsätze zur Sicherung guter wissenschaftlicher Praxis an der Heinrich-Heine-Universität Düsseldorf“ erstellt worden ist.

Desweiteren erkläre ich, dass ich eine Dissertation in der vorliegenden oder in ähnlicher Form noch bei keiner anderen Institution eingereicht habe.

Teile dieser Dissertation wurden bereits in Form folgender Zeitschriftenartikel und Konferenzberichte veröffentlicht oder zur Begutachtung eingereicht und sind entsprechend gekennzeichnet: Baumeister et al. (2016), Maushagen et al. (2021), Baumeister et al. (2020b), Baumeister et al. (2020a)

Ort, Datum

Ann-Kathrin Selker