

CHAPTER 2

Multiwinner Voting: A New Challenge for Social Choice Theory

Piotr Faliszewski, Piotr Skowron, Arkadii Slinko, and
Nimrod Talmon

2.1 Introduction

There are many reasons why societies run elections. For example, a given society may need to select its leader (e.g., a president), members of a team may need to find an appropriate meeting time, or referees may need to decide which candidate should receive an award in a contest. Each of these settings may call for a different type of election and a different voting rule. For example, Plurality with Runoff is used for presidential elections in France and Poland, Approval is used by Doodle, a popular website for scheduling meetings, and a rule very similar to Borda is used to select winners in the Eurovision song contest, in ski-jumping competitions, and in Formula 1 racing. Nonetheless, the general goal of finding a single candidate that is judged as highly as possible by as many people as possible is the same in each of these settings. The differences stem from a tension between the desire to select a candidate judged “as highly as possible” and supported “by as many people as possible” (e.g., in presidential election the focus is on the former requirement and it is considered acceptable that large minorities are dissatisfied with the elected president; in the scheduling example the focus is on the latter and it is perfectly fine to have a meeting time that is not optimal for anyone, provided that a large number of team members can attend). Other differences between the settings can be explained through practical considerations (e.g., choosing a president and choosing an award recipient are similar in spirit, but the latter carries much less weight for the society, is not constrained by laws, and so societies are willing to experiment with more sophisticated voting rules). Nonetheless, a rule that is very good for one of the above settings will likely do well for the others (for example, Laslier and Van der Straeten (2008) have shown feasibility of using approval voting for presidential elections).

However, there is also another family of elections, where instead of choosing a single best candidate, the goal is to choose a group of candidates, i.e., a committee. Such elections are even more ubiquitous than the single-winner ones, and include parliamentary elections, various business decisions (e.g., an Internet store has to decide which products to show on its homepage), or shortlisting

tasks (prior to deciding who should receive an award, typically there is a procedure that finds the finalists). These elections are far more varied than the single-winner ones, and different scenarios may require rules which follow different principles. Indeed, a rule that is good for shortlisting would, likely, select a poorly representative parliament, if a society were to use it for that purpose.

Following Elkind et al. (2017b), we distinguish the following three main types of multiwinner elections:¹

Shortlisting. In this case the voters correspond to experts acting as judges, referees, or reviewers. They have their opinions on the quality of the candidates (or on their suitability for the position that candidates have applied for) and the goal is to shortlist the “finalists”. The finalists are then evaluated far more accurately (e.g., invited for an interview) and, say, a single one of them is eventually chosen; making this final choice is beyond the scope of shortlisting. Thus, multiwinner rules for shortlisting tasks should simply pick the individuals of the highest quality, independently and without regard to any interactions between them. For example, two very similar candidates should either both be shortlisted or should both be rejected (with possible exceptions for boundary cases). As a consequence, shortlisting scenarios closely resemble single-winner elections.

Selecting a Diverse Committee. Consider the task of selecting locations for a number of facilities, such as fire stations in a city. Even though a location in the city center minimizes the average driving time to all other points in the city, and thus this is objectively the best location if we were to build just a single fire station, we do not want to build all fire stations in the central area; rather, we would prefer to distribute them more uniformly, so that each point in a city is in a close proximity to some fire station.² Similarly, consider an Internet store that has to choose what products to display on its homepage. One of the best strategies is to present a set of options which is as diverse as possible, keeping in mind that each customer should see something appealing to him or her. Selecting a diverse committee is guided by very different principles than shortlisting. It is no longer possible to evaluate the candidates separately and, e.g., if there are two similar candidates then we may either select one of them or neither of them (if there are better options), but we should not select them both.

Proportional Representation. Parliamentary elections are perhaps the best known type of multiwinner elections. In this case the goal is to select a committee (say, members of the parliament) in such a way that the views of the society are represented proportionally. Thus the main objective of

¹Some elements of this classification existed, of course, prior to the work of Elkind et al. (2017b). For example, Barberà and Coelho (2008) considered shortlisting tasks, and Chamberlin and Courant (1983) and Monroe (1995) (and many others) have considered committees providing optimal proportional representation.

²The facility location problem is often studied without regard to multiwinner elections and with somewhat different assumptions (e.g., optimizing locations in Euclidean spaces and not with respect to preferences of potential users). We point the reader to the book of Farahani and Hekmatfar (2009) for detailed discussion of facility location.

proportional representation is to find a committee of, say, k representatives, each associated with an equally sized constituency of approximately n/k voters (where n is the total number of voters). Importantly, this constituency may be territorial or virtual (i.e., depends on either geography or preferences). The requirement of constituencies of equal size is the incarnation of ‘one man, one vote’ principle for representative democracies, but sometimes it precludes electing the most diverse committee possible.

Naturally, there are also other, often more involved, settings where multiwinner elections are useful, but discussing them is beyond the scope of this chapter.

Multiwinner elections lead to a number of challenges, of which we discuss two in this chapter. The first one pertains to the problem of choosing a voting rule for a particular election type. How can we predict if a given rule would provide good results for a given setting? One approach, which we pursue, is to seek axiomatic properties useful for judging the suitability of a multiwinner rule for a particular application, and to analyze different rules with respect to these properties. For instance, we may check whether a rule that is meant for shortlisting extends the selected committee (without removing anyone from it) when we increase the target committee size (Elkind et al., 2017b; Barberà and Coelho, 2008), or we may check whether a rule for finding a proportional committee satisfies Dummett’s proportionality (Dummett, 1984), the Droop Proportionality Criterion (Woodall, 1994), as well as other similar notions (Elkind et al., 2017b; Aziz et al., 2017a). Other approaches, which we mostly omit due to space restrictions, include considering what various rules do on certain simpler domains, where their behavior can be interpreted intuitively (Elkind et al., 2017a; Brill et al., 2017), and various types of other theoretical and experimental evaluations (Diss and Doghmi, 2016; Caragiannis et al., 2016).

The second challenge regards our ability to compute the results of multiwinner elections. In the single-winner setting, almost all prominent voting rules are polynomial-time computable (although there are important exceptions, such as the rules of Dodgson, Young, and Kemeny (Bartholdi et al., 1989; Hemaspaandra et al., 1997; Rothe et al., 2003; Hemaspaandra et al., 2005)). For the multiwinner setting, the situation is much more complex. There is a number of polynomial-time computable rules, but many interesting ones are NP-hard. There are several ways in which we can deal with this problem. For elections of small enough size, we may be able to compute a winning committee either through FPT winner-determination algorithms (which are efficient when certain parameters, such as the number of voters or the number of candidates, are small), or through fast heuristics.³ If this approach is infeasible, then we may use (deterministic or randomized) approximation algorithms. Such algorithms can be viewed as new, easy to compute, rules, which even sometimes correspond to previously known voting rules. Thus, we study axiomatic properties of the rules defined by such approximation algorithms, just as we do for the original voting rules.

This chapter is organized as follows. First, in Section 2.2, we introduce formal

³Fortunately, “small enough” does not need to mean “impractically small”. For example, Elkind et al. (2017a) routinely compute results for several NP-hard rules for elections with 200 candidates and 200 voters each.

voter	ordinal ballot	approval ballot
v_1 :	$a \succ b \succ c \succ d \succ e$	$\{a, b, c\}$
v_2 :	$e \succ a \succ b \succ d \succ c$	$\{a, e\}$
v_3 :	$d \succ a \succ b \succ c \succ e$	$\{d\}$
v_4 :	$c \succ b \succ d \succ e \succ a$	$\{b, c, d\}$
v_5 :	$c \succ b \succ e \succ a \succ d$	$\{b, c\}$
v_6 :	$b \succ c \succ d \succ e \succ a$	$\{b\}$

Table 2.1: Two sample elections for candidate set $A = \{a, b, c, d, e\}$ and 6 voters, one with ordinal ballots and one with approval ballots (the approval ballots are formed by taking the top-ranked candidates from the ordinal ballots, for each voter choosing individually how many candidates to approve).

notions regarding the theory of multiwinner elections and discuss three important groups of multiwinner rules: committee scoring rules, approval-based rules, and rules based on the Condorcet principle. In Section 2.3 we discuss rules from these families, as well as some other relevant rules, for our three main tasks: shortlisting, selecting a diverse committee, and finding a committee that represents the voters proportionally. We conclude in Section 2.4, where we mention some further challenges regarding multiwinner voting.

2.2 Preliminaries

An election is a pair (A, R) , where A is a set of candidates and R is a profile of the voters' preferences. In the ordinal model, R consists of linear orders \succ_v , one for each voter v (order \succ_v ranks all the candidates and is often referred to as the preference order or the ordinal ballot of voter v). In the approval (or dichotomous) model, the profile contains, for each voter v , the set A_v of those candidates that this voter approves of (often referred to as the approval ballot of this voter). We show an example of both types of elections in Table 2.1.

A single-winner voting rule is a function that, given an election (A, R) , returns a set of candidates that tie as winners. For example, the Plurality rule selects those candidates that are ranked first by the largest number of voters (formally, we assume the voters have ordinal preferences; in practice, each voter provides its top candidate only). Analogously, a multiwinner voting rule is a function f that, given an election (A, R) and a positive integer k , $1 \leq k \leq |A|$, returns a nonempty family of size- k subsets of A , referred to as the winning committees. In practice, there always is some tie-breaking scheme that selects a single winning committee, but for simplicity we will disregard this issue. Unless specified otherwise, we assume the parallel universes tie-breaking model (Conitzer et al., 2009), where a voting rule outputs all the committees that could end up winning for some way of resolving ties that occur while executing the rule. If a rule always selects a single committee (e.g., because it is already combined with some tie-breaking scheme), then we say that it is resolute.

One of the most famous examples of multiwinner voting rules is the single transferable vote rule (STV) for ordinal elections, defined next.

Single Transferable Vote (STV) Rule. Consider an election with m candidates, n voters, and with the target committee size k . STV proceeds in rounds, until k candidates are elected. A single round proceeds as follows: We check if there is a candidate ranked first by at least $q = \lceil \frac{n}{k+1} \rceil + 1$ voters. If so, then such a candidate is included in the winning committee, q voters that rank him or her first are removed from the election, and he or she is removed from all the remaining preference orders. If such a candidate does not exist, then a candidate that is ranked first by the smallest number of voters is removed. (Note that this description strongly relies on parallel universes tie-breaking).

Example 2.1. Consider the ordinal election from Table 2.1 with the target committee size $k = 2$. STV uses the quota value $q = \lceil \frac{6}{3} \rceil + 1 = 3$. No candidate is ranked first by at least three voters so in the first round STV removes one candidate from $\{a, b, d, e\}$ (each of whom is ranked first only once, whereas c is ranked first twice). If we remove a , then in the next round still no candidate is ranked first by at least three voters and we need to remove either d or e . Say that we remove d . Then, in the next round b is ranked first by three voters (v_1, v_3 , and v_6), so we add b to the committee, remove b from the election, and remove these three voters. In the next two rounds we first remove e from the election and then add c to the committee. Thus $\{b, c\}$ is among the winning committees for this election under STV.

In what follows, we describe several families of multiwinner rules, starting with multiwinner analogues of single-winner scoring rules, through rules for approval elections, to rules based on the Condorcet principle. For a positive integer t , we write $[t]$ to denote the set $\{1, \dots, t\}$.

2.2.1 Committee Scoring Rules

Let us consider a setting with a set A of m candidates and with ordinal ballots. For a preference order \succ and candidate c , we write $\text{pos}_{\succ}(c)$ to denote the position of c in \succ (candidate ranked first has position 1, candidate ranked last has position m). A single-winner scoring function $\gamma_m, \gamma_m: [m] \rightarrow \mathbb{R}$, is a function that associates each position in a vote with a number of points, such that if $i < j$ then $\gamma_m(i) \geq \gamma_m(j)$. Two best-known examples of single-winner scoring functions are the Borda scoring function, $\beta_m(i) = m - i$, and the t -Approval family of scoring functions (where t is a positive integer; 1-Approval is known as Plurality):

$$\alpha_t(i) = \begin{cases} 1 & \text{if } i \leq t, \\ 0 & \text{otherwise.} \end{cases}$$

A family $\gamma = (\gamma_m)_{m \in \mathbb{N}}$ of single-winner scoring functions defines a rule f_γ as follows. The score of candidate c in an election $E = (A, R)$, where $R = (\succ_1, \dots, \succ_n)$, is $\text{score}(c, E) = \sum_{i=1}^n \gamma_{|A|}(\text{pos}_{\succ_i}(c))$. The rule selects the candidate(s) with the highest score (for example, the Borda rule uses the scoring functions β_m while the t -Approval rule uses α_t). Committee scoring rules are defined analogously, but for an extended notion of *position*.

Let S be a size- k committee and let \succ be a preference order. By the *position* of S in \succ , denoted $\text{pos}_\succ(S)$, we mean the sequence of positions of the members of S sorted in an increasing order; we write $[m]_k$ to denote the set of all size- k increasing sequences of elements from $[m]$. For two committee positions $I = (i_1, \dots, i_k)$ and $J = (j_1, \dots, j_k)$ from $[m]_k$, we say that I dominates J (denoted $I \succ J$) if for each t we have that $i_t \leq j_t$.

Elkind et al. (2017b) defined committee scoring rules as follows. A *committee scoring function* $\gamma_{m,k}: [m]_k \rightarrow \mathbb{R}$ for m candidates and committee size k , is a function that associates each committee position with a score in such a way that if $I, J \in [m]_k$ are two committee positions such that $I \succ J$, then $\gamma_{m,k}(I) \geq \gamma_{m,k}(J)$.

Definition 2.1 (Elkind et al. (2017b)). Let $\gamma = (\gamma_{m,k})_{k \leq m}$ be a family of committee scoring functions (one for each number m of candidates and committee size k). A committee scoring rule f_γ is a multiwinner rule that for election $E = (A, R)$, with $R = (\succ_1, \dots, \succ_n)$, and committee size k outputs those committees W for which $\text{score}(W, E) = \sum_{i=1}^n \gamma_{|A|,k}(\text{pos}_{\succ_i}(W))$ is the highest.

Many well-known multiwinner rules are, in fact, committee scoring rules:

Single Non-Transferable Vote (SNTV). Under SNTV, a committee receives a point from a voter if this committee contains the voters' most preferred candidate. That is, SNTV uses the scoring functions $\gamma_{m,k}^{\text{SNTV}}(i_1, \dots, i_k) = \alpha_1(i_1)$.

Bloc. Under Bloc, each voter names his or her k favorite candidates and the winning committee consists of those mentioned most frequently. In other words, Bloc uses the scoring functions $\gamma_{m,k}^{\text{Bloc}}(i_1, \dots, i_k) = \sum_{t=1}^k \alpha_k(i_t)$.

k -Borda. k -Borda outputs committee(s) that consist of k candidates with the highest (individual) Borda scores. That is, k -Borda uses the scoring functions $\gamma_{m,k}^{k\text{-Borda}}(i_1, \dots, i_k) = \sum_{t=1}^k \beta_m(i_t)$.

Chamberlin–Courant (β -CC). The Chamberlin–Courant rule (β -CC) uses the scoring functions $\gamma_{m,k}^{\beta\text{-CC}}(i_1, \dots, i_k) = \beta_m(i_1)$. This means that the score that a committee receives from a voter is the Borda score of the committee member that the voter ranks highest (among all the committee members). One possible interpretation is that each voter chooses a representative from the committee (clearly, a voter chooses the candidate that he or she likes the most) and gives the committee the Borda score of his or her representative. The rule was introduced by Chamberlin and Courant (1983).

Example 2.2. Let us again consider the ordinal election from Table 2.1. Under SNTV, every winning committee contains the candidate c and one other candidate. Under Bloc, the two winning committees are $\{a, b\}$ and $\{b, c\}$. Under k -Borda, the winning committee is $\{b, c\}$. The winning committee under β -CC is $\{a, c\}$, with a representing the voters v_1, v_2, v_3 and c representing the voters v_4, v_5, v_6 (it is a coincidence that each candidate represents the same number of voters).

Naturally, there are many other interesting committee scoring rules. For an overview of the internal structure of such rules we point the reader to the works of Faliszewski et al. (2016a,b); axiomatic characterization of these rules is due to Skowron et al. (2016b).

2.2.2 Approval-Based Rules

Let us now consider the approval model of elections. For the single-winner case, the approval rule simply selects those candidates that are approved by the largest number of voters. For the multiwinner setting, Aziz et al. (2017a) defined the following class of rules (which generalizes many previously studied ones; see the overview of Kilgour (2010) for more details regarding approval-based multiwinner rules, and the work of Aziz et al. (2015) for a computational perspective). Let A be a set of m candidates, let k be the committee size, and let $w^{(k)} = (w_1^{(k)}, \dots, w_k^{(k)})$ be a vector of k real numbers. The $w^{(k)}$ -AV score that a voter with approval ballot A_i assigns to a committee S is $\sum_{i=j}^{|S \cap A_i|} w_j^{(k)}$.

Definition 2.2 (Aziz et al. (2017a)). *Let $w = (w^{(i)})_{i \in \mathbb{N}}$ be a sequence of real-valued vectors (where each w^i has i coordinates). Given an election (A, R) and a committee size k , the w -AV rule outputs those committees for which the sum of the $w^{(k)}$ -AV scores assigned by the voters is the highest.*

Examples of w -AV rules include the following rules.

Approval Voting (AV). AV uses vectors $w^{(k)}$ of the form $(1, \dots, 1)$. That is, AV outputs committees of those k candidates that are approved most frequently.

Approval-Based Chamberlin–Courant rule (α -CC). Under the α -CC rule we use vectors of the form $(1, 0, \dots, 0)$. As in the case of the ordinal-based Chamberlin–Courant rule (β -CC), a possible interpretation is that each voter chooses a representative from the committee and, thus, increases the score of the committee by one if there is at least one committee member that this voter approves.

Proportional Approval Voting (PAV). The PAV rule uses vectors of the form $(1, 1/2, 1/3, \dots, 1/k)$. This rule indeed satisfies strong axioms pertaining to the proportionality of election results. We discuss this in more detail in Section 2.3.3. Quite amazingly, this rule was defined at the end of the nineteenth century (Thiele, 1895).

Example 2.3. *Let us consider the approval election from Table 2.1. The AV rule selects the committee $\{b, c\}$ (b is approved four times, c is approved three times, each other candidate is approved at most twice). The winning committee under α -CC is $\{a, b\}$ (with score five, where only v_3 does not approve any committee member), and the winning committees under PAV are $\{a, b\}$, $\{b, c\}$, and $\{b, d\}$, each obtaining 5.5 points (e.g., $\{a, b\}$ receives 1.5 points from v_1 and one point from each of the other voters except v_3 , who assigns zero points to this committee).*

There is a relation between w -AV rules and committee scoring rules (for the ordinal election model). For example, given a preference order and a committee size k , we might say that a voter approves his or her top k candidates. Then, a w -AV rule generates the following family of committee scoring functions,

$$\gamma_k^{w\text{-AV}}(i_1, \dots, i_k) = w_1 \alpha_k(i_1) + w_2 \alpha_k(i_2) + \dots + w_k \alpha_k(i_k),$$

and, thus, the corresponding committee scoring rule. For example, AV generates the Bloc rule. The choice of the approval threshold, k in this case, is quite arbitrary, but Faliszewski et al. (2016b) suggest reasons why it is natural: they refer to this committee scoring rules as top- k -counting rules and argue that only rules of such form can have certain axiomatic properties.

There are many multiwinner rules for the approval setting that are based on other principles than the w -AV rules. While the discussion of those is beyond the scope of this chapter, we do mention the Minimax approval voting rule (Brams et al., 2007), which, together with its generalizations (Amanatidis et al., 2015), received substantial attention from the research community.

2.2.3 Condorcet Committees and Related Rules

One of the most important notions regarding single-winner elections (in the ordinal model) is that of a Condorcet winner. A candidate c is a Condorcet winner if, for every other candidate d , a majority of the voters prefer c to d . A single-winner rule is Condorcet-consistent if it selects the Condorcet winner whenever one exists. Two prominent examples of Condorcet-consistent rules include the Copeland rule and the Maximin rule, defined next.

Consider an election $E = (A, R)$. For each two candidates c and d , we define $N_E(c, d)$ to be the number of voters that prefer c to d . The Copeland score of candidate c is the number of candidates d such that $N_E(c, d) > N_E(d, c)$ (i.e., the number of candidates that c defeats in a head-to-head majority contest⁴), whereas the Maximin score of c is defined as $\min_{d \in A \setminus \{c\}} N_E(c, d)$. The Copeland rule selects the candidates with the highest Copeland score and the Maximin rule selects those with the highest Maximin score.

The notion of the Condorcet winner was adapted to the multiwinner setting by Fishburn (1981a,b) as follows: A committee C is a Condorcet committee if for every other committee D (of the same size) a majority of voters prefers C to D . However, for this definition to be meaningful one has to either assume that the voters have explicit preferences over the committees, or that there is an accepted mechanism for lifting preferences over candidates to those over committees. For example, Fishburn considered the latter possibility for approval elections (he assumed that a voter prefers committee C over committee D if it contains more approved candidates). Recently, Darmann (2013) considered Condorcet committees for ordinal elections, where voters use Borda scores to compare committees (i.e., a voter prefers committee C to committee D if the sum of the Borda scores that the voter assigns to the members of C is greater than that of the members of D). Darmann (2013) showed computational hardness of deciding whether a given set is a Condorcet committee (both for the approval and ordinal settings; in fact, under some preference extensions even computing a Pareto optimal committee may be hard (Aziz et al., 2016)).

Gehrlein (1985) and Ratliff (2003) provided another interpretation of Condorcet consistency for the case of multiwinner elections, based directly on the

⁴If it happens that for some two candidates c and d we have $N_E(c, d) = N_E(d, c)$ then, typically, each of them receives some $\alpha \in [0, 1]$ points. Values 0, 0.5, and 1 are the most typical ones.

preferences over the candidates (Kaymak and Sanver (2003) showed that their notion can be understood in terms of Fishburn's Condorcet committees as well).

Definition 2.3 (Gehrlein (1985); Ratliff (2003)). *Let (A, R) be an election, let k be a committee size, and let S be some committee of size k . We say that S is a (weak) Condorcet set if for every candidate c in S and every candidate d in $A \setminus S$ it holds that more than half (at least half) of the voters prefer c to d .*

Following Barberà and Coelho (2008), we say that a multiwinner rule is stable if it outputs a weak Condorcet set of a given size k whenever such a set exists. For example, Coelho (2004) proposed the following weakly stable rules.

Number of External Defeats (NED). Under the NED rule, the score of a committee S is the number of pairs (c, d) of candidates such that $c \in S$, $d \in A \setminus S$, and at least half of the voters prefer c to d . The committee(s) with the highest score are the winners.

Minimum Size of External Opposition (SEO). Under the SEO rule, the score of a committee S in an election $E = (A, R)$ is defined as $\min_{c \in S, d \in A \setminus S} N_E(c, d)$ (i.e., the score of a committee S is the smallest number of voters that prefer some committee member to a committee nonmember). The committee(s) with the highest score are the winners.

These rules are natural analogues of the Copeland and Maximin rules (for a particular way of handling the cases where $N_E(c, d) = N_E(d, c)$ under the Copeland rule). Other single-winner Condorcet-consistent rules were adapted to the multiwinner setting by Ratliff (2003) and Kamwa (2016).

Example 2.4. *In the (ordinal) election from Table 2.1, the committee $\{b, c\}$ is a weak Condorcet set of size two. Indeed, exactly half of the voters prefer b to a , half of the voters prefer c to a , and strict majorities of the voters prefer each of b and c to each of d and e . In fact, $\{b, c\}$ is the unique weak Condorcet set of size two for this election and, so, is the unique winning committee under both NED and SEO.*

A completely different idea for extending the notion of a Condorcet winner to the multiwinner setting was introduced by Elkind et al. (2015). Briefly put, they said that committee S is a θ -winning set if for every candidate d not in S , more than a θ -fraction of the voters prefer some member of S to d ; they refer to $1/2$ -winning sets as Condorcet winning sets. Unfortunately, Condorcet winning sets cannot be easily interpreted as Condorcet committees in the sense of Fishburn (specifically, Elkind et al. considered several standard means of extending preferences over candidates to preferences over committees and under neither of them Condorcet winning sets turned out to be Fishburn's Condorcet committees). Nonetheless, the notion of a θ -winning set leads to an interesting multiwinner rule: Elkind et al. propose to output those committees S (of a given committee size k) that are θ -winning sets for the largest value of θ . For $k = 1$ this rule degenerates to the Maximin rule.

Example 2.5. *Under the rule of Elkind et al. (2015), the unique size-two winning committee for the (ordinal) election from Table 2.1 is $\{a, c\}$. For each candidate x from the set $\{b, d, e\}$, exactly five voters prefer either a or c to x ; e.g., v_1, v_2 and v_3 prefer a to b , v_4 and v_5 prefer c to b , and only v_6 prefers b to both a and c .*

2.3 Three Main Types of Multiwinner Elections

We now discuss the three main types of multiwinner elections mentioned in the introduction. For each of them, we consider formal properties that multiwinner rules for these elections should satisfy, mention rules that do satisfy these properties (and sometimes those that fail them), and discuss the computational complexity of identifying the winning committees under these rules.

2.3.1 Shortlisting

Shortlisting rules (also called screening rules by Barberà and Coelho (2008)) are those multiwinner rules that can be thought of as preliminary selection of candidates for the subsequent ultimate choice of, say, a single candidate; since only one candidate will be ultimately selected, it must be the ‘best’ one and any dependencies or similarities between candidates should not matter.⁵ It is implicitly assumed that the final choice can be made by other voters and will be based on other principles so any candidate from the selection can be ultimately chosen. Thus the main normative principle which any shortlisting rule should satisfy is the *committee monotonicity* (or *enlargement consistency* (Barberà and Coelho, 2008)). For simplicity, throughout the discussion of shortlisting we assume that our voting rules are resolute, i.e., that $f(E, k)$ is a singleton for each E and k .

Definition 2.4 (Elkind et al. (2017b); Barberà and Coelho (2008)). *Let f be a multiwinner voting rule. It is said to be committee monotone if for any election $E = (A, R)$ and any size of the target committee $k < |A|$ we have $f(E, k) \subset f(E, k+1)$.*

The idea is that if a candidate was good enough to be included in the list of k best ones, then it should be good enough to be included in the list of $k + 1$ best ones. For some rules committee monotonicity follows from their definition as for the following rule.

Sequential Plurality (Barberà and Coelho, 2008). We proceed in rounds. The first selected candidate is the Plurality winner (i.e., the candidate ranked first by the largest number of voters). Then this candidate is removed and the procedure is repeated. This is done k times.

Each committee monotone rule f produces a ranking of the candidates. Let us consider some election $E = (A, R)$ and take the convention that $f(E, 0) = \emptyset$. If for each k we let $\{a_k\} = f(E, k) \setminus f(E, k-1)$, then we obtain the ranking $a_1 \geq a_2 \geq \dots \geq a_m$. In other words, a committee monotone rule f generates a social welfare function F which, given an election E , produces the ranking $F(E)$ constructed above. Moreover, $f(E, k)$ is the set of top k elements of $F(E)$ (relative to some tie-breaking mechanism). Analogously, if F is a social welfare function and E is

⁵Some shortlisting scenarios are more complicated. For example, when shortlisting a group of people considered for a job, it may be necessary to maintain a certain level of diversity of the committee, to ensure that minorities are not discriminated against. In this chapter, we assume that this need is understood by the voters and is already reflected in their preference orders. If one cannot make such an assumption, then one may have to seek other rules and principles than those discussed here.

an election, then we can define a multiwinner voting rule f by setting $f(E, k)$ to be the top k candidates of $F(E)$ (relative to some fixed tie-breaking rule). Elkind et al. (2017b) refer to such rules as *best- k rules*. Some examples follow.

Best- k rules for positional scoring SWFs. Let $\gamma = (\gamma_m)_{m \in \mathbb{N}}$ be a single-winner scoring function. The social welfare function associated with γ ranks the candidates (in a given election) according to their γ scores. For example, k -Borda is a best- k rule from this family.

Best- k rules based on non-positional scoring SWFs. Sometimes scores of candidates come from other sources. For example, a social welfare function can output a ranking of candidates according to their Maximin scores. This leads to a best- k rule that we call k -Maximin.

Best- k rules based on the majority relation. Suppose for simplicity that n is odd and, given an election $E = (A, R)$ with $R = (\succ_1, \dots, \succ_n)$, define the majority relation $>_E$ as:

$$a >_E b \iff |\{i \in [n] \mid a \succ_i b\}| > |\{i \in [n] \mid b \succ_i a\}|.$$

Notice that this majority relation is a tournament. We can now define the score of a candidate c as the outdegree of c (considered as a vertex in this tournament). This score is, in fact, the Copeland score of c and, so, we refer to the corresponding best- k rule as k -Copeland.

Barberà and Coelho (2008) noticed that no committee monotone (shortlisting) rule can be stable (see Section 2.2.3); to this end, they presented a simple profile which possesses a unique Condorcet set with two elements and a disjoint unique Condorcet set with three elements. This is disappointing because an unstable shortlisting rule can produce a committee that contains some candidate c such that a majority of the voters prefers to it another candidate d who is not in the committee. Another consequence of this result is that the NED rule is different from k -Copeland and the SEO rule is different from k -Maximin. Indeed, the former two rules are NP-hard to compute (as all stable rules (Aziz et al., 2017b)), whereas the latter two are polynomial-time computable.

On the other hand, Elkind et al. (2017b) identified a subclass of committee scoring rules that are committee monotone.

Definition 2.5 (Elkind et al. (2017b)). *A committee scoring rule f is separable if there exists a family of committee scoring functions $\gamma = (\gamma_{m,k})_{k \leq m}$ and a family of single-winner scoring functions $\delta = (\delta_m)_{m \in \mathbb{N}}$ such that $f = f_\gamma$ and for each m, k ($k \leq m$), and a committee position $I = (i_1, \dots, i_k) \in [m]_k$ we have that*

$$\gamma_{m,k}(i_1, \dots, i_k) = \delta^m(i_1) + \dots + \delta^m(i_k).$$

For example, k -Borda is a separable committee scoring rule, whereas Bloc is not (while at first it seems to be defined in an appropriate way, the single-winner scoring functions used in its definition depend on k and this is not allowed in separable committee scoring rules). In particular, Bloc is not committee monotone (Staring, 1986).

Theorem 2.1 (Elkind et al. (2017b)). *Every separable committee scoring rule is committee monotone.*

It also holds that every separable committee scoring rule is polynomial-time computable, provided that its underlying single-winner scoring functions are.

2.3.2 Selecting a Diverse Committee

In the introduction we provided examples of settings where a diverse committee is a desirable outcome of a voting rule. Throughout this section we will focus on yet another one, due to Elkind et al. (2017b), considering an airline which designs the content of its in-flight entertainment system for the airplanes. There are numerous movies, TV programs, and sports competitions to choose from and, due to technical and financial reasons, only a small selection can be chosen. The airline would like to maximize the satisfaction of the passengers and, thus, a diversity among the selected entertainment items is highly desirable.

Specifically, we assume that each passenger chooses a single movie⁶ (the one that he or she likes best among the available ones).⁷ If every passenger has only a single favorite movie and does not wish to watch anything else (as might be the case for a group of small children), then it is natural for the airline to use the SNTV rule. This way, the largest number of passengers will get their favorite movie (while the rest will be left dissatisfied). On the other hand, if each passenger has a set of good movies and is satisfied if at least one of these movies is available, then it is natural to model the problem as an approval election and to use the α -CC rule. Finally, if every passenger has a ranking of the movies and the appreciation that a passenger has for a movie decreases linearly as its position in the ranking grows,⁸ then β -CC is our rule of choice.

The above rules are either committee scoring rules (SNTV and β -CC) or can be interpreted as such (recall the discussion below Example 2.3). Elkind et al. (2017b) refer to committee scoring rules where the score depends only on the position of the most preferred candidate as *representation-focused* rules.

Definition 2.6 (Elkind et al. (2017b)). *A committee scoring rule f is representation-focused if there exists a family of committee scoring functions $\gamma = (\gamma_{m,k})_{k \leq m}$ and a family of single-winner scoring function $\delta = (\delta_{m,k})_{k \leq m}$ such that $f = f_\gamma$ and for each m, k ($k \leq m$), and a committee position $I = (i_1, \dots, i_k) \in [m]_k$ we have $\gamma_{m,k}(i_1, \dots, i_k) = \delta_{m,k}(i_1)$.*

Let us now consider which axiomatic properties should be satisfied by rules that are appropriate for selecting diverse committees (we focus on the ordinal setting). Somewhat surprisingly, the literature does not offer many choices. Firstly,

⁶For simplicity, we speak only of movies, omitting other types of entertainment.

⁷It would also be quite natural to assume that every passenger chooses two best movies, or that he or she watches the best movie with some high probability, the second best with a lower probability, the third best with even lower one, and so on. Skowron et al. (2016a) study such settings and identify an interesting class of rules based on ordered weighted average (OWA) operators (these rules can also be interpreted as committee scoring rules).

⁸This is a very idealized assumption. In practice, no passenger can possibly have an opinion about all movies.

such a rule must satisfy the following criterion which is a straightforward adaptation of the notion of a *consensus committee* of Elkind et al. (2017b).

Definition 2.7. *A voting rule f satisfies the narrow-top criterion if for each election $E = (A, R)$ and each positive integer $k \leq |A|$ the following holds: if there exists a committee W of size k such that each voter ranks some member of W on top, then $W \in f(E, k)$.*

Secondly, the following condition requires that if a rule selects some committee W then this committee should still win if any voter shifts his or her most preferred member of W forward.

Definition 2.8 (Faliszewski et al. (2016a)). *We say that a voting rule f is top-member monotone if for every election E , positive integer k , committee $W \in f(E, k)$, and election E' obtained from E by shifting forward in some vote the top ranked member of W , it holds that $W \in f(E', k)$.*

All representation-focused committee scoring rules satisfy the narrow-top criterion and are top-member monotone (Faliszewski et al., 2016a).

Unfortunately, among the three rules that we discussed here only SNTV is polynomial-time computable (and this rule suffers from being dependant on each voter's first choice only). As for the other rules, Procaccia et al. (2008) showed that both α -CC and β -CC are NP-hard to compute. On the other hand, Betzler et al. (2013) used the framework of parameterized complexity to show that winner determination for these rules can be solved efficiently for elections with few voters or with few alternatives. They also showed that these rules are polynomial-time computable for single-peaked elections, whereas Skowron et al. (2015b) have shown the same for single-crossing elections.

There are approximation algorithms which efficiently find committees whose score is close to the optimal one. For example, the greedy algorithm of Lu and Boutilier (2011) executes k greedy iterations, in each selecting a candidate whose inclusion brings the greatest marginal increase to the total committee score; this algorithm achieves approximation ratio of $1 - 1/e$ (this holds for both α -CC and β -CC; unless $P = NP$, this is the best possible polynomial-time approximation for α -CC (Skowron and Faliszewski, 2015)). Skowron et al. (2015a) describe several other approximation algorithms for β -CC, all of which are somewhat based on the greedy approach, that achieve better approximation guarantees in certain situations, including a polynomial-time approximation scheme (PTAS). Skowron and Faliszewski (2015) give an FPT approximation scheme for α -CC (parameterized by the committee size). While using approximation algorithms for computing outcomes of voting rules in political elections may be controversial (but see the discussion of Faliszewski et al. (2016c)), in any business related application of voting rules the use of approximation algorithms is fully justified.

In practice, for up to medium-sized elections, finding a winning committee under α -CC and β -CC can be done by solving a certain integer linear program (ILP), as described by Lu and Boutilier (2011). Currently the best heuristic solution is to use a clustering algorithm by Faliszewski et al. (2016c).

2.3.3 Proportional Representation

Black (1958) defines proportionality of a voting rule as the ability to reflect “all shades of political opinion” of a society within the winning committee. Commonly, parliaments—or any other committees that are meant to represent voters proportionally—are elected using the first-past-the-post (FPTP) voting system, where the voters and candidates are divided into electoral districts, and a representative of each district is elected via Plurality voting. This is practical because typically it is easier for voters to compare candidates from their districts only, but it might lead to large disproportionality. For example, if there are two main opposing political views, X and Y , and 49% of the voters in each district support view X while 51% support view Y , then each district elects a Y supporter, and nearly half of the population is not represented.

Under SNTV each voter also votes for a single person, but the voters are not divided into electoral districts. If a committee of size k is to be elected, then the k candidates with the best plurality scores form it. Both under FPTP and SNTV the voters only reveal their top-preferred candidates, yet, often the preferences of the voters are much more complex and they are rarely apathetic about the candidates different from their top choice. Thus, it is natural and important to study forms of proportionality which take into account full preferences of the voters; this idea is often referred to as *fully proportional representation*. Dummett (1984) was among the first to initiate such study for the case of ordinal preferences, formulating the following axiom.

Definition 2.9 (Dummett (1984)). *Consider a setting with n voters, where we want to select a committee of size k . If there exists some $\ell \in [k]$ and a group of $\ell \cdot n/k$ voters who all rank the same ℓ candidates on top of their preference orders, then these ℓ candidates should all belong to all the winning committees.*

For $\ell = 1$, Elkind et al. (2017b) refer to this property as the *solid coalitions* property and show that both STV and SNTV, among others, satisfy it. There is also a variant of Dummett’s proportionality which uses the Droop quota (i.e., $\lfloor n/k \rfloor + 1$) instead of the value n/k (Woodall, 1994). A variant of the STV rule which is used, e.g., for electing the Australian senate satisfies this version of Dummett’s proportionality. Indeed, STV is often considered to be very well-suited for tasks that require proportional representation (Tideman and Richardson, 2000; Elkind et al., 2017b,a).

Monroe (1995) suggested another interesting rule that takes full ordinal ballots as input and aims at achieving proportional representation.

Monroe. Consider an election $E = (A, R)$, with $R = (\succ_1, \dots, \succ_n)$, and let k be the size of committee to be elected. For a committee S , an *assignment* is a function $\Phi: [n] \rightarrow S$ that maps voters to committee members. We interpret $\Phi(i)$ as the member of S that represents voter i (under the assignment Φ). We say that Φ is *balanced* if for each $c \in S$ we have $\lfloor n/k \rfloor \leq |\Phi^{-1}(c)| \leq \lceil n/k \rceil$. We define the score of assignment Φ as $\text{score}(\Phi) = \sum_{i=1}^n \beta(\text{pos}_{\succ_i}(\Phi(i)))$, i.e., as the total Borda score of the voters’ representatives. The score of a committee S is the score of the best balanced assignment of voters to the members of S . The Monroe rule selects the committee(s) with the highest score.

Monroe's rule resembles the Chamberlin and Courant rule, which also implicitly defines an assignment of voters to their representatives in a winning committee, and both are based on the concept of satisfaction which both rules maximize. However, Monroe's rule additionally requires that each committee member represents roughly the same number of voters. This makes a meaningful difference—the Monroe's rule is proportional while the Chamberlin–Courant's is not. Unfortunately, finding winners according to the Monroe rule is computationally hard (Procaccia et al., 2008), even when certain natural parameters of the election are small (Betzler et al., 2013) or when preferences of the voters are single-crossing (Skowron et al., 2015b) (hardness for single-peaked elections is known only for a more general variant of the rule (Betzler et al., 2013)). Yet, recently, Skowron et al. (2015a) proposed a greedy variant of this rule:

Greedy Monroe (Skowron et al., 2015a). The rule executes k iterations as follows. In iteration i , we find a group V_i of n/k voters and a candidate c for which the total Borda score that the voters from V_i assign to c is maximal. Then, we add c to the winning committee, assigns c as a representative to the voters from V_i , and remove these voters from further consideration.

The Greedy Monroe rule can be viewed as an approximation algorithm for the original rule, but it also exhibits some new interesting properties. For example, it satisfies the solid coalitions property, whereas the original Monroe rule does not (Elkind et al., 2017b).

To conclude the discussion of proportional representation in the ordinal election model, let us recall that Elkind et al. (2015) introduced the concept of θ -winning sets (see Section 2.2.3 for the definition) which combines the ideas behind proportional representation and the Condorcet principle.

Now, let us move to the rules which take approval ballots as input. We start by considering the following illustrative example.

Example 2.6. Consider an approval election where the set of 30 standing candidates can be split into three disjoint sets, C_1 , C_2 , and C_3 of equal size, such that 50 voters approve all candidates in C_1 , 30 voters—all candidates in C_2 , and 20 voters—all candidates in C_3 . If our goal is to select a committee of size $k = 10$, then we would expect any proportional rule to choose 5 candidates from C_1 , 3 candidates from C_2 , and 2 candidates from C_3 .

Of course, usually we cannot hope for such a nice structure of the voters' preferences, but Example 2.6 is helpful in understanding the behavior of approval-based voting rules. Let us consider RAV, the greedy variant of the PAV rule:

Reweight Approval Voting (RAV). Consider an election with n voters, where the i -th voter approves candidates in the set A_i . RAV starts with an empty committee S and executes k rounds. In each round it adds to S a candidate c with the maximal value of $\sum_{i: c \in A_i} \frac{1}{|S \cap A_i| + 1}$, i.e., a candidate c which maximizes the PAV score of $S \cup \{c\}$.

Let us discuss how RAV works for the election from Example 2.6. Before the first round S is empty, so adding a candidate from C_1 to S would increase the

total PAV score of S by 50; adding a candidate from C_2 and C_3 would increase the total score by 30 and 20, respectively. Thus, in the first round a candidate from C_1 is selected. The following rounds proceed analogously. Eventually, after 7 rounds, S contains 4 candidates from C_1 , 2 candidates from C_2 and 1 candidate from C_3 . In the eighth step, 50 voters have already 4 representatives so adding a candidate from C_1 to S (which would become their fifth approved candidate) would increase the PAV score of each of them by $1/5$, increasing the total score by 10. Similarly, adding a candidate from C_2 or from C_3 to S would also increase the total score of S by 10. We see that in the next three steps RAV selects one candidate from each of the sets C_1 , C_2 , and C_3 , forming a proportional committee.

Interestingly, the harmonic sequence of weights $w^{(k)} = (1, 1/2, \dots, 1/k)$ is the unique sequence which results in proportionality on such nicely structured preferences as in Example 2.6. This was formalized by Aziz et al. (2017a) and Brill et al. (2017). In particular, Aziz et al. (2017a) defined two properties, called justified representation and extended justified representation, defined next.

Definition 2.10 (Aziz et al. (2017a)). *A rule satisfies extended justified representation (EJR) if for each approval election with n voters, each committee size k , and each $\ell \in [k]$, the following holds: There is no group of $\lceil \ell \cdot n/k \rceil$ voters that all approve at least ℓ common candidates, but neither of whom approves ℓ or more members of each winning committee. A rule satisfies justified representation (JR) if it satisfies EJR for $\ell = 1$.*

Intuitively, justified representation requires that, if there is a group of at least n/k voters whose approval ballots have at least one candidate in common, then it cannot be the case that neither of these voters is represented in the committee. EJR extends this reasoning to larger groups of voters and larger sets of jointly approved candidates. Aziz et al. (2017a) showed that PAV is the only w -AV rule which satisfies EJR. Brill et al. (2017), on the other hand, discussed a relation between multiwinner voting rules and methods of apportionment, which allows to view PAV and RAV as extensions of the d'Hondt method of apportionment to the multiwinner setting (see Chapter 3 of this book for more details on seat allocations). Similarly, the Monroe rule can be adapted to work on approval ballots—such variant of the Monroe rule can be viewed as a generalization of the Hamilton method. Unfortunately, finding winners according to PAV is NP-hard (Aziz et al., 2015; Skowron et al., 2016a). Yet, RAV can be viewed as a good approximation algorithm for PAV (Skowron et al., 2016a) which can be even better approximated when certain natural parameters are low (Skowron, 2016).

So far, we only referred to “linear proportionality”. There exist other interesting concepts, such as degressive proportionality (Koriyama et al., 2013) which says that smaller groups of voters should be given more representatives than the traditional proportionality suggests. Thus, degressive proportionality recommends taking a step from traditional proportionality towards diversity. Also, we only discussed proportionality with respect to voters’ preferences. Other forms of proportional representation can be considered as well—for instance, where different candidates have different attributes (e.g., gender, age, nationality, affiliation), and where our goal is to select a representative committee with respect to each of the attributes (Lang and Skowron, 2016).

2.4 Further Challenges

We discussed axiomatic and algorithmic properties of various multiwinner rules for our three main tasks. Yet, these are not the only challenges regarding electing committees. For instance, many voting rules require full preference rankings provided by voters, and with a large number of candidates obtaining such information might be infeasible. It is thus natural to study multiwinner voting for the case where only partial preference information is available. Other challenges include the problem of convincing societies to adopt new rules, the problem of modeling political parties (Brill et al. (2017) provide some very initial studies in this respect), the problem of presenting the election results (it is easy to tell who won, but candidates may wish to know how well they did even if they lost), and many others. These are very important and we believe that addressing them will at least partially shape future studies of multiwinner voting.

Acknowledgments

Piotr Faliszewski was supported by NCN project DEC-2016/21/B/ST6/01509. Piotr Skowron was supported by ERC-StG 639945 (ACCORD). Arkadii Slinko was supported by the Royal Society of NZ Marsden Fund 3706352.

Bibliography

- G. Amanatidis, N. Barrot, J. Lang, E. Markakis, and B. Ries. Multiple referenda and multiwinner elections using Hamming distances: Complexity and manipulability. In *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems*, pages 715–723, 2015.
- H. Aziz, S. Gaspers, J. Gudmundsson, S. Mackenzie, N. Mattei, and T. Walsh. Computational aspects of multi-winner approval voting. In *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems*, pages 107–115, 2015.
- H. Aziz, J. Lang, and J. Monnot. Computing Pareto optimal committees. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence*, pages 60–66, 2016.
- H. Aziz, M. Brill, V. Conitzer, E. Elkind, R. Freeman, and T. Walsh. Justified representation in approval-based committee voting. *Social Choice and Welfare*, 2017a. To appear.
- H. Aziz, E. Elkind, P. Faliszewski, M. Lackner, and P. Skowron. The Condorcet principle for multiwinner elections: From shortlisting to proportionality. Technical Report arXiv:1701.08023 [cs.GT], Jan. 2017b.
- S. Barberà and D. Coelho. How to choose a non-controversial list with k names. *Social Choice and Welfare*, 31(1):79–96, 2008.

- J. Bartholdi, III, C. Tovey, and M. Trick. Voting schemes for which it can be difficult to tell who won the election. *Social Choice and Welfare*, 6(2):157–165, 1989.
- N. Betzler, A. Slinko, and J. Uhlmann. On the computation of fully proportional representation. *Journal of Artificial Intelligence Research*, 47:475–519, 2013.
- D. Black. *The Theory of Committees and Elections*. Cambridge University Press, 1958.
- S. Brams, M. Kilgour, and R. Sanver. A minimax procedure for electing committees. *Public Choice*, 132(3–4):401–420, 2007.
- M. Brill, J. Laslier, and P. Skowron. Multiwinner approval rules as apportionment methods. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence*, 2017. To appear.
- I. Caragiannis, S. Nath, A. D. Procaccia, and N. Shah. Subset selection via implicit utilitarian voting. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence*, pages 151–157, 2016.
- B. Chamberlin and P. Courant. Representative deliberations and representative decisions: Proportional representation and the Borda rule. *American Political Science Review*, 77(3):718–733, 1983.
- D. Coelho. *Understanding, Evaluating and Selecting Voting Rules Through Games and Axioms*. PhD thesis, Universitat Autònoma de Barcelona, 2004.
- V. Conitzer, M. Rognlie, and L. Xia. Preference functions that score rankings and maximum likelihood estimation. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence*, pages 109–115. AAAI Press, July 2009.
- A. Darmann. How hard is it to tell which is a Condorcet committee? *Mathematical Social Sciences*, 66(3):282–292, 2013.
- M. Diss and A. Doghmi. Multi-winner scoring election methods: Condorcet consistency and paradoxes. Technical Report WP 1613, GATE Lyon Saint-Étienne, Mar. 2016.
- M. Dummett. *Voting Procedures*. Oxford University Press, 1984.
- E. Elkind, J. Lang, and A. Saffidine. Condorcet winning sets. *Social Choice and Welfare*, 44(3):493–517, 2015.
- E. Elkind, P. Faliszewski, J. Laslier, P. Skowron, A. Slinko, and N. Talmon. What do multiwinner voting rules do? An experiment over the two-dimensional euclidean domain. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence*, 2017a. To appear.
- E. Elkind, P. Faliszewski, P. Skowron, and A. Slinko. Properties of multiwinner voting rules. *Social Choice and Welfare*, 2017b. To appear.

- P. Faliszewski, P. Skowron, A. Slinko, and N. Talmon. Committee scoring rules: Axiomatic classification and hierarchy. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence*, pages 250–256, 2016a.
- P. Faliszewski, P. Skowron, A. Slinko, and N. Talmon. Multiwinner analogues of the plurality rule: Axiomatic and algorithmic views. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence*, pages 482–488, 2016b.
- P. Faliszewski, A. Slinko, K. Stahl, and N. Talmon. Achieving fully proportional representation by clustering voters. In *Proceedings of the 15th International Conference on Autonomous Agents and Multiagent Systems*, pages 296–304, 2016c.
- F. Z. Farahani and M. Hekmatfar, editors. *Facility Location: Concepts, Models, and Case Studies*. Springer, 2009.
- P. Fishburn. Majority committees. *Journal of Economic Theory*, 25(2):255–268, 1981a.
- P. Fishburn. An analysis of simple voting systems for electing committees. *SIAM Journal on Applied Mathematics*, 41(3):499–502, 1981b.
- W. Gehrlein. The Condorcet criterion and committee selection. *Mathematical Social Sciences*, 10(3):199–209, 1985.
- E. Hemaspaandra, L. Hemaspaandra, and J. Rothe. 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, 1997.
- E. Hemaspaandra, H. Spakowski, and J. Vogel. The complexity of Kemeny elections. *Theoretical Computer Science*, 349(3):382–391, 2005.
- E. Kamwa. Stable rules for electing committees and divergence on outcomes. *Group Decision and Negotiation*, pages 1–18, 2016.
- B. Kaymak and R. Sanver. Sets of alternatives as Condorcet winners. *Social Choice and Welfare*, 20(3):477–494, 2003.
- M. Kilgour. Approval balloting for multi-winner elections. In *Handbook on Approval Voting*. Springer, 2010. Chapter 6.
- Y. Koriyama, J. F. Laslier, A. Macé, and R. Treibich. Optimal Apportionment. *Journal of Political Economy*, 121(3):584–608, 2013.
- J. Lang and P. Skowron. Multi-attribute proportional representation. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence*, pages 530–536, 2016.
- J. Laslier and K. Van der Straeten. A live experiment on approval voting. *Experimental Economics*, 11(1):97–105, 2008.

- T. Lu and C. Boutilier. Budgeted social choice: From consensus to personalized decision making. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence*, pages 280–286, 2011.
- B. Monroe. Fully proportional representation. *American Political Science Review*, 89(4):925–940, 1995.
- A. Procaccia, J. Rosenschein, and A. Zohar. On the complexity of achieving proportional representation. *Social Choice and Welfare*, 30(3):353–362, 2008.
- T. Ratliff. Some startling inconsistencies when electing committees. *Social Choice and Welfare*, 21(3):433–454, 2003.
- J. Rothe, H. Spakowski, and J. Vogel. Exact complexity of the winner problem for Young elections. *Theory of Computing Systems*, 36(4):375–386, 2003.
- P. Skowron. FPT approximation schemes for maximizing submodular functions. In *Proceedings of the 12th Conference on Web and Internet Economics*, pages 324–338, 2016.
- P. Skowron and P. Faliszewski. Fully proportional representation with approval ballots: Approximating the MaxCover problem with bounded frequencies in FPT time. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence*, pages 2124–2130, 2015.
- P. Skowron, P. Faliszewski, and A. Slinko. Achieving fully proportional representation: Approximability result. *Artificial Intelligence*, 222:67–103, 2015a.
- P. Skowron, L. Yu, P. Faliszewski, and E. Elkind. The complexity of fully proportional representation for single-crossing electorates. *Theoretical Computer Science*, 569:43–57, 2015b.
- P. Skowron, P. Faliszewski, and J. Lang. Finding a collective set of items: From proportional multirepresentation to group recommendation. *Artificial Intelligence*, 241:191–216, 2016a.
- P. Skowron, P. Faliszewski, and A. Slinko. Axiomatic characterization of committee scoring rules. Technical Report arXiv:1604.01529 [cs.GT], Apr. 2016b.
- M. Starling. Two paradoxes of committee elections. *Mathematics Magazine*, 59: 158–159, 1986.
- T. N. Thiele. Om flerfoldsvvalg. In *Oversigt over det Kongelige Danske Videnskaberne Selskabs Forhandlinger*, pages 415–441. 1895.
- N. Tideman and D. Richardson. Better voting methods through technology: The refinement-manageability trade-off in the Single Transferable Vote. *Public Choice*, 103(1–2):13–34, 2000.
- D. Woodall. Properties of preferential election rules. *Voting Matters*, 3:Paper 4, 1994.